Severity-Aware Optimization of UAV-Based Emergency Medical Services with AI-Driven Prioritization

by

Habiba Yeasmin

A thesis presented to the Lakehead University in partial
fulfillment of the thesis requirement for the degree of

Master of Science
in
Electrical and Computer Engineering

Lakehead University
Barrie STEM Hub, Ontario, Canada
April 2026

Examining Committee Membership

The following served on the Examining Committee for this thesis.

Supervisor:	Dr. Waleed Ejaz Associate Professor, Dept. of Electrical and Computer Engineering, Lakehead University
Co-supervisor:	Dr. Faria Khandaker Assistant Professor, Dept. of Computer Science and Technology, Algoma University
Examiner:	Dr. Shafiqul Hai Assistant Professor, Dept. of Electrical and Computer Engineering, Lakehead University
Examiner:	Dr. Maysa Yaseen Assistant Professor, Dept. of Electrical and Computer Engineering, Lakehead University

Author's Declaration

Abstract

Rapid emergency medical response following disasters is often hindered by damaged infrastructure, limited situational awareness, and the difficulty of rapidly assessing and prioritizing victims using conventional emergency medical service (EMS) systems. Although uncrewed aerial vehicles (UAVs) have shown promise for aerial reconnaissance and disaster monitoring, existing UAV-assisted emergency-response frameworks typically focus either on victim detection or on logistics-oriented resource allocation in isolation, with limited integration between aerial perception and downstream dispatch decision making. Consequently, current systems do not adequately support severity-aware UAV-assisted EMS allocation in which dispatch decisions are informed by the inferred condition or urgency of observed victims. To address this problem, this thesis proposes an integrated UAV-assisted emergency medical response framework that links aerial victim detection, visual criticality estimation, and optimization-based UAV dispatch within a unified perception-to-decision pipeline. UAV-acquired disaster imagery is first processed using a YOLOv8-based human detection model, a deep learning–based real-time object detection algorithm, to localize affected individuals. Detected victims are then analyzed using a binary criticality classifier trained on aerial disaster imagery from the C2A dataset, augmented with posture-based criticality annotations to distinguish higher-risk victims from less urgent cases. These outputs are combined within a triage-inspired scoring framework to generate severity and priority estimates for spatial demand regions. The resulting perception-derived severity and priority information is incorporated into a tailored mixed integer linear programming (MILP) model for UAV-enabled EMS dispatch and facility-allocation optimization that jointly considers travel time, operational cost, severity coverage, and priority coverage. Unlike conventional cost-focused UAV-assisted EMS baseline, which assumes homogeneous demand, the proposed model explicitly incorporates perception-derived triage information into dispatch decisions. Experimental evaluation demonstrates that incorporating perception-derived severity and priority information enables the proposed framework to allocate UAV resources in a manner more aligned with victim criticality than conventional cost-focused dispatch strategies. These results demonstrate the feasibility of integrating aerial perception with optimization-based dispatch to support severity-aware UAV-assisted EMS planning and provide a foundation for future perception-driven emergency-response systems.

Acknowledgements

Firstly, I would like to express my sincere gratitude to my advisors Dr. Waleed Ejaz and Dr. Faria Khandaker for the continuous support of my MASc study and related research, for their patience, motivation, and immense knowledge. Their guidance helped me in all the time of research and writing my thesis. I could not have imagined having better advisors and mentors for my MSc study.

Besides my advisors, I would like to thank the rest of my thesis committee for their insightful comments, which encouraged me to broaden my research from multiple perspectives.

Last but not the least, I would like to thank all my research group members, my parents and family for always encouraging me.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

Acronyms	Description
AI	Artificial intelligence
BS	Base station
C2A	Combination to application
CNN	Convolutional neural network
CPU	Central processing unit
DL	Deep learning
DoF	Degrees of freedom
EMS	Emergency medical service
GNSS	Global navigation satellite system
I/O	Input or output
IoT	Internet of things
MAP	Mean average precision
MEC	Mobile edge computing
MILP	Mixed integer linear programming
MIMO	Multiple input multiple output
RF	Random forest
SAR	Search and rescue
SVM	Support vector machine
UAV	Uncrewed aerial vehicle
UEMS	Uncrewed aerial vehicle for the emergency medical service
UGV	Unmanned ground vehicle
ULAP	UEMS location-allocation problem
WSN	Wireless sensor network
YOLO	You only look once

List of Symbols

Symbol	**Description**
I	Set of UAV facility locations.
J	Set of demand points (victims or tiles).
$i \in I$	Index for facilities.
$j \in J$	Index for demand points.
x_{ij}	Binary variable; 1 if facility i serves demand j, otherwise 0.
y_i	Facility activation variable; 1 if facility i is opened.
u_i	Number of UAVs assigned to facility i.
z_j	Service indicator; 1 if demand point j is served.
T_{ij}	Travel time from facility i to demand point j.
$T_{\max}$	Maximum allowable UAV flight time.
f_i	Fixed cost to open facility i.
p_i	Operating cost per UAV at facility i.
N_i^{UAV}	Maximum UAV capacity that facility i can host.
$N_{\max}$	Maximum number of facilities allowed to open.
$B_{\max}$	Maximum permissible budget for UAV operations.
S_j	Severity score at point j
U_j	Urgency score t point j
D_j	Demand load at point j
p_j	Priority score t point j
L_j	Number of affected individuals
T_j	Time since the incident report at demand point j.
E_j	Proximity-to-epicenter metric.
a_1, a_2, a_3, a_4	Weights for severity scoring
$\theta_1, \theta_2, \theta_3$	Weights for urgency formulation.
β_1, β_2	Weights for demand formulation.
λ	Weight controlling demand–urgency trade-off in priority scoring.
w_1, w_2, w_3, w_4	Optimization parameter weights .
K	Total number of UAVs available

Chapter 1

Introduction

1.1 Background

Large-scale disasters, such as earthquakes, floods, wildfires, and industrial accidents continue to reveal significant weaknesses in conventional emergency response systems [1–3]. Ground-based emergency medical services (EMS) often rely on road infrastructure, stable communication networks, and the physical reach of human teams. When roads are blocked, areas are inaccessible, or communication infrastructure is damaged, response times increase , and situational awareness [4] deteriorates . The first hours after a disaster are often the most critical for survival, yet they are also the most chaotic, uncertain, and resource-constrained. The authors of [5] have explored integrated uncrewed aerial vehicles (UAVs)-Internet of things (IoT) disaster management systems in which UAVs not only collect heterogeneous disaster scene data but also support downstream analytics such as infrastructure assessment, blocked-road identification, and health-risk estimation to guide first responders toward high-priority areas.

UAVs have emerged as a key technology for enhancing disaster management [6]. They can rapidly survey large areas, access unsafe or unreachable locations for ground teams, and provide real-time imagery and sensor data. An early and influential system-level perspective on UAV-assisted disaster management highlights the potential of UAV networks, combined with wireless sensor networks (WSNs) and cellular infrastructure, to support disaster prediction, assessment, and response throughout multiple phases of the disaster lifecycle [7]. It emphasizes that UAVs extend beyond mobile cameras, acting as integral nodes in a sensing and communication infrastructure [8], supporting search and rescue (SAR) [9, 10], damage assessment, and connectivity restoration [11].

While this system's perspective is important, operational effectiveness ultimately depends on quality of information UAVs can gather from the environment and how that information is utilized. Recent work has focused on enhancing capabilities such as robust human detection [12], small-object recognition [13], infrared imaging, pose estimation, and cooperative perception. At the same time, another area of research addresses trajectory optimization, multi-agent coordination, and edge/cloud-based decision making for UAV swarms during emergencies. The authors of [14] have explored UAV-assisted mobile edge computing (MEC) frameworks for disaster scenarios, optimizing UAV placement, task offloading, user association, and computational resource allocation to improve communication and computational efficiency under infrastructure-constrained emergency conditions. However, these developments have generally occurred in isolation; perception systems are evaluated primarily in terms of detection accuracy, while optimization frameworks concentrate on trajectory cost or coverage, often neglecting the incorporation of medically relevant severity scores derived from visual evidence.

Despite significant progress in UAV perception and disaster response optimization, the link between these fields remains limited. Current UAV-based perception systems primarily prioritize detecting and locating victims, whereas optimization strategies for UAV coordination typically allocate resources based on coverage, trajectory cost, or demand volume. Consequently, detected victims are often treated as uniform targets, without regard for the variations in medical severity among them. In real disaster response scenarios, victims can have widely varying levels of injury and urgency, necessitating prioritization in resource allocation.

Thus, a key challenge in UAV-assisted disaster response is determining how to allocate limited UAV resources across affected areas while prioritizing victims who require the most urgent medical attention. This thesis addresses the identified challenge by proposing a severity-aware UAV-based

EMS framework integrating visual perception, severity inference, and optimization-based dispatch.

1.2 UAVs for Disaster Response and Human Detection

The first major focus of this thesis is on UAV-assisted disaster response and human detection. The authors of [7] have categorized the roles of UAVs across three stages of disaster management: pre-disaster preparedness, real-time disaster assessment, and post-disaster response and recovery. They showed that UAVs can enhance early warning systems, provide live situational awareness, and establish communication networks. However, they did not specify how to prioritize individual victims once they are detected.

To transition from generic situational awareness to victim-centered triage, several studies propose specialized human detection pipelines for UAV imagery. The authors of [15] have proposed an efficient artificial intelligence (AI)-based ensemble model for human detection and rescue during disasters. They integrated multiple computer vision techniques into an ensemble, enhancing robustness in cluttered and degraded environments. It specifically targeted challenging disaster scenarios where victims may be partially occluded, lying on the ground, or surrounded by debris. By combining outputs from different detectors, the ensemble approach reduced false negatives and improved the reliability of victim localization.

In addition, previous research has proposed a comprehensive framework for real-time human detection using UAVs in disaster management, presenting a fully integrated perception system designed for SAR missions [16]. This framework focused on real-time processing, deployment constraints, and practical aspects of integrating UAV-based human detection into field operations. It underscored that the effectiveness of these systems is measured not only by per-frame accuracy but also by their capacity to facilitate timely and coordinated actions during an evolving disaster. These contributions indicated a shift in the community from general UAV imagery to application-specific human detection. However, in both cases, detected humans are typically treated as indistinguishable targets. There is a lack of explicit modeling of medical severity or triage categories based on observable cues such as pose, posture, and density. Addressing this gap is central to the motivation for this thesis.

1.3 UAV Perception: Small and Infrared Object Detection

Disaster environments often involve small, distant, partially occluded targets, because people or objects may occupy only a few pixels in high-altitude imagery. Standard object detectors tend to underperform in these situations. To address this issue, the authors of [17] have introduced a high-frequency enhancement mechanism specifically designed for small-scale targets while maintaining real-time inference capabilities. This enhancement is particularly beneficial for high-altitude or wide-area surveys, where individual victims may appear as very small blobs in the scene. Meanwhile, the authors of [18] proposed a real-time network for UAV-based infrared object detection, focusing on challenging low-visibility conditions such as night-time operations or smoke-filled environments. They utilized infrared imaging and enhanced feature aggregation to effectively detect warm-bodied targets in situations where red green blue (RGB) imagery was degraded or unreliable. Hence, these detection-oriented works demonstrated that specialized architectures and modalities (e.g., infrared) can significantly enhance the robustness of UAV perception in disaster scenarios. In the context of a severity-aware EMS framework, this indicates that (i) victims can be reliably detected under more varied conditions and (ii) detection confidence and target size can be utilized as components of a severity score. However, [17], [18] did not translate the perception outputs into explicit triage categories or optimization-relevant demand metrics.

1.4 UAV Pose Estimation and Cooperative Vision

A second major theme in the supporting literature is precise pose estimation [19] and collaborative visual perception [20] for UAVs. These studies focus less on victim detection and more on understanding the geometric characteristics of the environment and the UAV swarm itself. This understanding is crucial for stable and reliable operations in complex disaster scenes. The authors of [21] proposed a multiview landmark-assisted method for 6-degree of freedom (DoF) UAV swarm pose estimation at the swarm level. They integrated resonant beam ranging, visual-inertial odometry (VIO), and multiview landmark constraints to achieve high-precision global pose optimization for multiple UAVs. The back-end optimization employed factor graphs and sliding windows to jointly optimize the poses of all UAVs, yielding significant improvements in localization accuracy compared to VIO alone. This kind of cooperative localization is highly relevant for multi-UAV search patterns over large disaster areas, where consistent coverage and accurate inter-UAV alignment are essential.

At the individual UAV level, the authors of [22] have developed a robust 6-DoF pose estimation method based on hybrid features. They combined keypoint-based and pose-object-based representations, utilizing both semantic and geometric cues for high-precision UAV pose estimation. Their method was designed to be transferable across different UAV shapes and sizes, while also operating at real-time speeds. This robustness is critical when the same perception framework must function on various UAV platforms in the field. Similarly, the authors of [23] have advanced the idea of aerial collaborative stereo, introducing a real-time cross-camera feature association and relative pose estimation framework for UAVs. Their dual-channel association strategy, combined with a relative multistate-constrained Kalman filter, achieved real-time performance on resource-constrained onboard hardware. The key insight was that multi-UAV systems can dynamically form variable-baseline stereo setups, enabling flexible, large-scale depth perception, even when global navigation satellite system (GNSS) signals are unreliable or degraded.

The authors of [24] proposed a globally optimal relative pose estimation solver that used affine correspondences under known vertical direction constraints. They transformed the relative pose estimation problem into a reduced DoF polynomial system and solved it optimally. Although not specifically for disaster scenarios, this approach achieved highly accurate relative poses between views. This method is directly applicable to UAV platforms integrating cameras with inertial measurement units (IMUs), which enhances robustness through multimodal perception. Similarly, the authors of [25] adopted a geolocation-based approach, where UAV pose estimation was performed using IMU assistance and satellite imagery matching. The combination of onboard IMU data and appearance-based matching to satellite basemaps enables reliable localization in GPS-challenged environments, an important requirement in large-scale disaster scenarios. Furthermore, the authors of [26] presented a vision-based estimation and tracking framework for fixed-wing multi-UAV systems that integrated you only look once (YOLO)-based detection with a bearing-angle estimators and tracking modules. Their approach demonstrated that real-time visual tracking of multiple UAVs in formation can be achieved through careful integration of perception and state estimation.

Overall, these works establish that (i) cooperative visual perception among UAVs is feasible and (ii) high-precision pose estimation can be maintained even in GNSS-degraded environments. For a severity-aware EMS system, this means that multi-UAV swarms can reliably execute coordinated search patterns, revisit critical tiles, and support consistent mapping of severity over time. However, these contributions have focused primarily on pose estimation and cooperative vision without defining how the resulting geometric information should be integrated with medical triage priorities or dispatch optimization.

1.5 Computational Infrastructure for Real-Time UAV Decision Support

Real-time deployment of UAV-assisted disaster response systems requires a robust computational infrastructure to support perception, optimization, and decision-making in time-sensitive operational conditions. Prior studies have investigated computational architectures and system frameworks that enable real-time processing and optimization of UAV data in disaster response scenarios [14]. The authors of [27] have proposed GPU-accelerated edge computing frameworks that offload computationally intensive UAV optimization tasks to nearby edge servers, thus improving responsiveness in emergencies. Similarly, the authors of [28] have proposed an integrated uncrewed ground vehicle (UGV)–UAV disaster response platform that allowed heterogeneous autonomous agents to coordinate through communication and computation frameworks for multi-platform disaster operations. Furthermore, the authors of [29] and [30] have proposed a cloud and IoT-based disaster management framework that collected, processed, and disseminated large-scale sensor and UAV data streams through a distributed infrastructure to support operational decision-making. Collectively, these works highlight the practical feasibility of implementing computationally intensive perception and optimization pipelines in disaster response environments, provided they are supported by suitable edge, cloud, or distributed computing infrastructures. While these systems primarily focus on computational architecture and operational feasibility rather than triage-aware optimization, they offer valuable context for the real-time deployment assumptions underlying the proposed UAV EMS framework.

1.6 Thesis Motivation

The central motivation of this thesis is to bridge the gap between UAV-based visual perception and EMS resource allocation by introducing severity-aware prioritization into UAV dispatch optimization. On one hand, recent advances in UAV perception have significantly improved human detection accuracy, robustness under challenging conditions such as small-object and infrared imaging, and cooperative pose estimation [15–18, 21–26]. On the other hand, trajectory planning, multi-agent coordination, and distributed optimization frameworks have demonstrated the feasibility of near real-time decision-making for UAV swarms and heterogeneous emergency response systems [7, 27, 28, 30]. However, the connection between perception and decision-making is typically loose. The existing works typically treat perception outputs such as detected human location as end results, rather than actionable inputs for resource allocation. Consequently, there is no principled mechanism for translating visual information into clinically meaningful indicators of urgency. In practice, this leads to decision frameworks where all detected victims are implicitly treated as equally important, ignoring the substantial variation in injury severity and medical need. This limitation becomes particularly critical in disaster scenarios, where resource constraints necessitate prioritization: responders must decide not only where victims are located, but also who should be assisted first, when regions should be revisited, and how limited UAV resources should be allocated under competing demands.

In this context, the thesis adopts severity-aware coverage as a fundamental organizing principle for UAV-based EMS systems. This principle requires moving beyond mere detection to inferring the criticality of each detected individual from visual cues, enabling its integration into downstream decision-making. This shift transforms perception outputs into structured, optimization-relevant inputs that directly guide resource allocation decisions. Motivated by this perspective, the thesis develops an integrated framework, where different aspects of the thesis are presented as follows:

- The visual descriptors extracted from perception are transformed into severity, urgency, demand, and priority scores through a spatial triage modeling layer inspired by EMS principles.
- These scores are then used to parameterize EMS-inspired demand, urgency, and priority models that guide decision-making.
- A multi-objective optimization model subsequently allocates UAV missions to enable severity-aware coverage while satisfying operational constraints.

The proposed framework is evaluated through simulation experiments and compared against severity-agnostic allocation strategies. This allows a systematic analysis of the impact of incorporating severity-aware prioritization on UAV dispatch decisions, resource utilization, and overall emergency response effectiveness.

1.7 Thesis Contributions

Guided by the above-mentioned motivation and the literature review, this thesis positions itself at the intersection of perception-driven analysis and optimization-based decision-making in UAV-assisted emergency response systems. It aims to build a severity-aware UAV-based dynamic EMS framework in which visual perception, severity inference, demand modeling, and optimization are tightly coupled. The focus is not only on finding people but also on quantifying how critical their situation is and using that information to drive mathematically grounded dispatch decisions. Hence, the thesis pursues the following main contributions:

- We develop a pipeline that converts raw UAV imagery into per-person severity indicators and tile-level severity and priority scores. This is achieved by selecting and adapting an aerial human-detection architecture and defining severity-related features derived from visual evidence.
- We develop a tiling and aggregation scheme that represents each disaster region through estimated severity, victim density, and priority values, enabling direct integration of spatial triage information into optimization models.
- We develop a mathematical optimization model that assigns UAVs to disaster tiles and optimizes dispatch decisions under capacity, travel-time, and budget constraints while explicitly prioritizing severity-aware coverage and rapid response to higher-severity regions. The proposed formulation incorporates perception-derived severity prioritization directly into the UAV dispatch process.
- We integrate the perception, triage, and optimization components into a unified end-to-end framework and evaluate its performance under realistic simulation scenarios to illustrate how severity-aware formulations influence dispatch decisions relative to severity-agnostic baseline approaches. A severity-aware approach incorporates pose-based binary criticality classification and aggregated severity scores to prioritize high-risk regions, whereas severity-agnostic approaches treat all demand points equally.

1.8 Thesis Organization

The remainder of the thesis is organized as follows: **Chapter 2** presents a comprehensive literature review and analyzes each prior work in detail. It examines their assumptions, evaluation protocols, and limitations, while positioning the proposed framework within the broader context of UAV, EMS, and optimization research. **Chapter 3** describes the proposed severity-aware UAV perception model using real-time AI prioritization. This chapter explains the detection backbone, feature extraction strategy, spatial tiling, and the definition of severity and priority scores derived from UAV imagery. **Chapter 4** presents the experimental setup and evaluation results, while **Chapter 5** concludes the thesis with a summary of contributions, limitations, and potential directions for future work.

Chapter 2

Literature Review

This chapter reviews the existing literature related to uncrewed aerial vehicle (UAV)-assisted emergency medical response systems and provides the research background for the proposed framework in this thesis. Since the proposed system combines aerial victim perception, severity-aware assessment, and optimization-based UAV dispatch, the literature is organized into three main areas. First, studies on UAV-based victim detection and aerial perception are reviewed to examine current approaches for understanding disaster scenes from aerial imagery. Second, artificial intelligence (AI)-based severity and triage assessment methods are discussed to explore how victim criticality and emergency priority can be estimated in healthcare and emergency response settings. Third, existing research on UAV-enabled emergency medical services (EMS) optimization and resource allocation is reviewed to analyze current approaches for UAV dispatch, facility placement, and medical resource planning. Based on this review, the limitations of existing studies are identified and the research gap addressed by this thesis is established.

2.1 UAV-Based Victim Detection and Aerial Perception

Effective UAV-assisted emergency medical response begins with accurate understanding of the disaster scene. Hence, the UAV dispatch or resource allocation decisions rely on first observing the affected environment, identifying victims, and extracting relevant situational information from aerial sensor data. For this reason, UAV-based victim detection and aerial perception play a central role in enabling perception-driven disaster response and supporting downstream emergency decision-making.

2.1.1 UAVs in Disaster Reconnaissance

UAVs have become important tools in disaster management operations because of their rapid deployability, mobility, and ability to access hazardous or infrastructure-compromised environments that may be difficult or unsafe for human responders to enter/reach. In comparison to the conventional ground-based reconnaissance and manned aerial surveying, the UAV platforms provide faster situational awareness acquisition while reducing operational risk. Specifically, the UAVs offer capability to capture aerial imagery and real-time video, which makes them valuable for search and rescue, damage assessment, and emergency monitoring. In this context, the authors of [7] have identified UAV-assisted disaster management as a promising paradigm for enhancing emergency response effectiveness through aerial sensing, monitoring, communication support, and autonomous mission execution. Their work highlighted the broad applicability of UAVs across disaster response tasks and highlighted their potential to improve the speed and coverage of early situational assessment.

Subsequent studies have further expanded UAV deployment toward integrated disaster response frameworks. The authors of [15] [16] and [28] have proposed UAV-enabled disaster management systems that combine aerial sensing with automated analysis and real-time operational support. These works demonstrated the growing transition from passive UAV observation toward intelligent UAV-assisted disaster response systems capable of contributing directly to emergency assessment workflows. Similarly, the authors of [31] investigated intelligent decision-support architectures for UAV-assisted search and rescue coordination. They demonstrated growing interest in integrating UAV perception and mission-support logic within broader emergency response systems. In this context, UAV-based reconnaissance in the early post-disaster stage is particularly valuable because responders require rapid understanding of affected areas before deploying limited rescue resources. UAV imagery can reveal blocked access routes, damaged infrastructure, debris concentration, and possible victim locations across large disaster zones. This capability significantly improves the speed and spatial coverage of initial situational assessment compared with manual surveying approaches.

Despite these advantages, effective UAV-based disaster reconnaissance remains challenging in practice. Disaster environments are visually complex and unstructured, often involving smoke, rubble, occlusion, unstable illumination, weather interference, and cluttered backgrounds. In addition, UAV platforms face operational limitations such as battery constraints, payload restrictions, communication bandwidth limitations, and restricted onboard computational resources. Overall, the literature establishes UAVs as highly promising platforms for rapid disaster reconnaissance and situational awareness generation. However, most broad UAV disaster management studies focus primarily on monitoring and general response support, without addressing how UAV-derived observations can be transformed into structured victim-centered decision signals for medical prioritization and emergency response optimization.

2.1.2 Human and Victim Detection From Aerial Imagery

Automated human detection from UAV imagery has become a central research direction in disaster response perception because manual inspection of large aerial image streams is slow, labor-intensive, and difficult to sustain under time-critical emergency conditions. Existing studies consistently show that UAV platforms can improve victim search efficiency by providing broad visual coverage and enabling computer-vision-based detection of humans in affected areas. The authors of [32] have reviewed this research direction comprehensively and showed that UAV-based human detection has evolved into an active search and rescue research area driven largely by deep learning (DL)-based object detection methods. Similarly, the authors of [15, 16] have proposed UAV-enabled disaster response frameworks in which automated human detection plays a key operational role in supporting rescue activities. Furthermore, several studies have explored concrete human-detection and localization pipelines for aerial imagery that goes beyond high-level disaster response frameworks. In this context, the authors of [33] have investigated DL-based human detection in aerial images, while the authors of [34] have examined UAV-based human detection and localization in indoor disaster environments. Collectively, these studies indicate that aerial human detection is no longer treated merely as a theoretical possibility, but as a technically feasible perception component for disaster response systems.

However, the existing literature remains strongly detection-centric. Most prior works focus on whether a human can be detected or localized from UAV imagery, rather than how such detections should be interpreted for downstream emergency response decision making. In practice, identifying the presence of a victim is only the first step in an effective medical response pipeline. Rescue teams also require information that helps distinguish which observations are more urgent, which locations should be prioritized, and how limited emergency resources should be allocated under operational constraints. This reveals a key limitation in the current aerial human-detection literature. Although prior studies demonstrate the usefulness of UAVs for victim search and localization, they rarely extend perception outputs into structured triage-aware or severity-aware representations. As a result, perception is often treated as an isolated vision task rather than as an integrated component of a larger emergency medical decision framework. This limitation directly motivates this thesis, which extends beyond victim detection to transform UAV-derived visual observations into decision-relevant severity and priority metrics for optimization-driven emergency response.

2.1.3 Small-Object Detection Challenges in UAV Imagery

A fundamental challenge in UAV-based victim detection arises from the small-object nature of aerial human targets. Because UAV platforms typically operate at elevated altitudes to maximize area coverage, individual victims often occupy only a limited number of pixels within captured imagery. This significantly reduces discriminative visual detail and makes aerial victim detection substantially more difficult than conventional ground-level object detection [32], [35]. In disaster scenarios, the problem is further intensified by debris occlusion, complex backgrounds, unstable illumination, motion blur, and extreme viewpoint variation [36].

Recognizing these challenges, recent research has increasingly focused on adapting object-detection architectures to better preserve fine-grained spatial information in aerial imagery. These efforts are motivated by the observation that repeated down sampling in conventional deep detection backbones can suppress shallow low-level features and degrade localization performance for very small targets [37]. To mitigate this issue, multiple studies have introduced enhanced multi-scale feature-fusion strategies, attention mechanisms, and modified prediction heads specifically tailored for UAV imagery [17, 18, 35–38]. Collectively, these approaches demonstrate that preserving shallow spatial detail and improving cross-scale representation are critical for robust aerial small-object detection.

Another notable trend is the continued refinement of you only look once (YOLO)-based one-stage detectors for UAV-specific applications. Rather than abandoning established detection frameworks, many studies build upon mature architectures such as YOLO by introducing targeted modifications to backbone, neck, or head components [17], [35], and [37]. This suggests that the research community increasingly views UAV perception challenges as requiring domain-specific architectural adaptation rather than entirely new detection paradigms.

Although these architectural enhancements have improved detection accuracy on aerial benchmarks, they introduce practical trade-offs that are often under emphasized in the literature. More sophisticated feature-fusion modules and attention mechanisms can increase model complexity, inference latency, and computational cost, which may reduce suitability for real-time deployment on resource-constrained UAV or edge computing platforms [35–37]. Consequently, improvements in benchmark accuracy do not necessarily translate directly into operationally deployable emergency response systems, particularly when rapid onboard or near-edge inference is required.

Furthermore, most existing small-object detection studies remain benchmark-oriented and detector-centric in their evaluation. Performance is typically assessed using conventional metrics such as precision, recall, and mean average precision [17, 18, 35–37], with limited discussion of how improved detections contribute to downstream emergency response effectiveness. In addition, many of these studies focus on generic aerial targets including vehicles, pedestrians, or surveillance objects rather than disaster-specific victim analysis. As a result, the literature has primarily advanced the technical capability of aerial detection without addressing how these detections should be interpreted for rescue prioritization or medical response planning.

Therefore, while specialized aerial detection architectures provide an important technical foundation for robust UAV perception, current small-object detection research remains largely disconnected from operational emergency response decision making. Improved detection performance alone does not ensure that UAV perception outputs are actionable in medical or disaster response settings. Bridging this gap requires additional modeling layers capable of transforming aerial detections into structured severity-aware and priority-aware decision signals. This motivates the next stream of literature, which examines AI-based approaches for severity estimation and triage-oriented assessment in emergency response applications.

2.2 AI-Based Severity and Triage Assessment

While victim detection enables the identification and localization of affected individuals in disaster scenes, effective emergency response requires additional assessment of victim condition and relative urgency. In practical emergency medical operations, limited rescue resources must often be allocated selectively, making it necessary to prioritize victims or affected regions based on severity, urgency, or estimated medical criticality. As a result, increasing research attention has been directed toward AI-based severity assessment and triage-support systems that assist in estimating patient condition and guiding emergency prioritization decisions [39, 40]. This section reviews existing literature on AI-driven triage and severity estimation methods, followed by UAV-assisted remote health and wellness assessment approaches, in order to examine how current systems support emergency prioritization and where important limitations remain.

2.2.1 AI-Driven Triage and Medical Severity Prediction

AI has increasingly been explored as a decision-support tool for medical triage and severity assessment in emergency response settings. Traditional triage processes often rely heavily on manual assessment by medical personnel, which can be time-consuming, subjective, and difficult to scale during large-scale emergencies or mass-casualty incidents [39, 41]. To address these limitations, recent research has investigated AI-based approaches for assisting pre-hospital triage and automated severity estimation.

A major direction in this literature involves learning-based models that infer patient urgency or triage category from structured medical or telemedicine data. For example, the authors of [39] have proposed a graph neural network-based triage framework for pre-hospital emergency telemedicine systems, demonstrating the use of AI to support automated triage decision making from patient-related information. Such approaches highlight the growing role of machine learning in assisting early-stage emergency assessment and reducing reliance on purely manual triage procedures.

Beyond direct triage prediction, several studies have incorporated criticality-aware prioritization mechanisms into broader emergency response systems. The authors of [40] have proposed a criticality-driven scheduling strategy for UAV-assisted remote health monitoring, where patient urgency is explicitly considered when prioritizing health-data transmission. Similarly, the authors of [42] have presented an AI-enabled UAV emergency response framework that integrates incident detection and triage-oriented response coordination for rapid roadside medical emergencies. Collectively, these studies demonstrate increasing recognition that emergency response systems should incorporate patient criticality or urgency as a core decision factor rather than relying solely on first-come-first-served or cost-driven dispatch logic.

Despite these advances, existing AI-based triage and severity assessment methods are predominantly designed for structured clinical, telemedicine, or sensor-based environments in which explicit patient health measurements such as physiological signals, medical telemetry, structured patient records, or direct clinical observations; are available during assessment [39, 40, 42]. While effective in healthcare and telemedicine environments, such assumptions may not hold in post-disaster UAV reconnaissance scenarios, where direct physiological measurements are often unavailable and only visual aerial observations may be accessible. As a result, current AI-driven triage literature provides limited support for UAV-based disaster response settings in which severity must be inferred indirectly from aerial scene observations rather than structured medical data. This limitation motivates the need for alternative severity-assessment formulations capable of estimating emergency criticality from visual and contextual UAV-derived indicators, which is more consistent with the operational constraints of aerial disaster reconnaissance.

2.2.1.1 UAV-Assisted Remote Health and Wellness Assessment

In parallel with AI-driven triage research, a growing body of work has explored the use of UAV platforms for remote health monitoring and wellness assessment in situations where direct medical access is limited [43, 44]. These approaches aim to extend healthcare assessment capabilities beyond conventional clinical settings by leveraging UAV-mounted sensing, communication, and monitoring technologies to support remote patient observation.

Several studies have investigated UAV-assisted remote healthcare architectures in which UAVs serve as airborne monitoring or communication nodes for collecting and transmitting patient-related information. For example, the authors of [43] have proposed an Internet-of-UAV-based e-health system that utilizes UAVs to collect patient health data and provide agile edge computing services for remote medical monitoring. Such systems demonstrate the feasibility of UAV-enabled remote healthcare support in dispersed or infrastructure-limited environments. Related conceptual and forward-looking studies have further examined the broader role of AI-powered healthcare assistants within drone-enabled medical monitoring ecosystems, highlighting the growing interest in intelligent UAV-assisted healthcare support frameworks beyond conventional telemedicine settings [45].

More recently, UAV-enabled wellness-monitoring frameworks have explored extracting physiological or condition-related indicators directly from remote sensing modalities. The authors of [44] have presented a drone-assisted remote wellness monitoring pipeline that estimates physiological indicators from red green blue (RGB) video for telehealth-oriented triage support. Their work demonstrates the potential of UAV-based visual sensing to support remote assessment beyond simple victim localization, moving toward richer condition-aware emergency evaluation.

Although the above-mentioned studies expand the role of UAVs from transportation and observation platforms toward health-aware sensing systems, their applicability to disaster response scenarios remains limited. Most existing UAV-based wellness-monitoring approaches assume relatively controlled sensing conditions, stable subject visibility, and close-range observation sufficient for extracting physiological or biometric signals. Such assumptions are difficult to satisfy in post-disaster aerial reconnaissance, where UAVs often operate at higher altitudes over cluttered, dynamic, and visually degraded environments.

Furthermore, many current UAV-assisted health-monitoring systems focus on physiological telemetry, remote patient monitoring, or telehealth support rather than large-scale victim prioritization in disaster scenes [43, 44]. Consequently, while these approaches demonstrate the broader feasibility of UAV-assisted remote health assessment, they do not directly address the problem of estimating victim severity from coarse aerial visual observations under disaster response constraints. This highlights an important gap between UAV-enabled health monitoring and UAV-based disaster triage. Existing remote wellness-assessment frameworks show that UAVs can support health-aware sensing, yet they remain insufficient for large-area disaster reconnaissance scenarios in which rapid severity estimation must be inferred from limited aerial visual evidence. Bridging this gap requires severity-assessment approaches specifically designed for aerial disaster imagery and scalable multi-victim prioritization contexts.

2.3 UAV Dispatch, Resource Allocation, and Facility Location Optimization

UAV-assisted emergency response requires optimization frameworks for dispatch decision-making, resource allocation, and facility location planning under operational constraints. Existing literature includes UAV dispatch models, emergency logistics optimization, and facility location allocation formulations for UAV-enabled medical and disaster response applications. While these works provide important foundations for UAV operational planning, most do not explicitly incorporate perception-derived victim severity or priority information into dispatch and allocation decisions. This limitation motivates the severity-aware optimization framework proposed in this thesis.

2.3.1 Resource Allocation and UAV-Assisted Emergency Medical Logistics

UAVs are increasingly studied for healthcare resource allocation and emergency medical logistics in disaster response settings. The authors of [46] have proposed a UAV-assisted healthcare resource distribution framework that allocates medical resources across disaster regions while accounting for demand distribution and vehicle constraints. Similarly, the authors of [47] have investigated UAV transportation of emergency medical supplies under payload and operational feasibility constraints, demonstrating the practical potential of UAV-supported medical delivery. Related studies have further explored autonomous UAV-based payload delivery systems for medical emergency response [48], while broader reviews of UAV-enabled healthcare delivery have highlighted the growing role of UAV-assisted logistics in critical and time-sensitive medical operations [49], [50].

Collectively, the above-mentioned studies show that UAVs can improve emergency logistics efficiency and complement conventional ground-based response mechanisms. Existing research consistently demonstrates the feasibility of UAV-assisted healthcare delivery for accelerating medical supply transportation, expanding emergency reach, and supporting rapid disaster response logistics. However, their optimization formulations primarily model demand in an abstract or homogeneous manner, emphasizing delivery cost, travel time, operational feasibility, or logistical efficiency without explicitly incorporating heterogeneous victim severity, medical urgency, or perception-derived triage information. Consequently, while operationally valuable, these approaches do not directly address perception-driven, severity-aware dispatch and prioritization in

UAV EMS scenarios, where limited UAV resources must be allocated according to dynamically inferred victim criticality rather than homogeneous delivery demand.

2.3.2 Facility Location Allocation Models

Facility location allocation models constitute one of the most structured and mathematically rigorous categories within UAV-assisted EMS optimization literature. More broadly, strategic disaster-relief and humanitarian logistics research has long employed facility-network design models to determine the placement of shelters, medical centers, and emergency supply facilities for efficient disaster preparedness and response [51]. These studies highlight the importance of strategic facility placement in enabling timely and scalable emergency operations.

Within UAV-assisted EMS specifically, the UAVs for emergency management system (UEMS) location allocation problem (ULAP) [52] serves as a foundational baseline for strategic UAV-assisted emergency medical planning. ULAP formulates the placement of UAV facilities and assignment of UAV resources to demand nodes as an integer programming problem that minimizes total deployment and operational cost, subject to facility activation, UAV availability, assignment, and service coverage constraints. ULAP has been influential due to its well-defined mathematical formulation, scalability, and direct applicability to strategic UAV EMS system design. In particular, it provides a tractable framework for cost-oriented facility planning and UAV resource deployment across service regions.

However, ULAP is formulated primarily for strategic cost-focused planning under static predefined demand assumptions. The model does not incorporate perception-derived severity or priority information, nor does it explicitly differentiate demand nodes based on heterogeneous victim urgency or medical criticality. Furthermore, its single-objective cost-minimization formulation does not consider triage-aware trade-offs such as prioritizing high-severity demand points or balancing cost against severity and priority coverage. Accordingly, while ULAP provides a strong strategic planning baseline for UAV EMS deployment, it is not designed for perception-driven, severity-aware dispatch and allocation scenarios in which demand importance varies according to

UAV-derived triage information. These limitations motivate the extension of cost-oriented facility location frameworks toward triage-aware multi-objective UAV EMS optimization.

2.3.3 Broader Multi-Objective UAV Optimization Literature

Beyond EMS planning, a growing body of UAV optimization research has explored multi-objective formulations in broader mission-planning and emergency response applications. These studies recognize that practical UAV deployment often requires balancing multiple competing operational objectives rather than optimizing a single performance metric.

The authors of [53] have proposed a multi-objective UAV optimization framework for wildfire monitoring that jointly considered surveillance coverage and energy efficiency. Similarly, the authors of [54] have optimized UAV deployment for post-disaster communication restoration while balancing communication coverage and network resilience. Similarly, the authors of [55] also investigated sensing-communication trade-offs in UAV platform selection and aerial integrated sensing networks. Additional disaster-oriented UAV deployment studies have examined reactive drone positioning and coverage-maximization strategies for autonomous disaster area monitoring and service provisioning [38]. Collectively, these studies demonstrate increasing recognition that UAV operational planning often involves multidimensional trade-offs across competing objectives.

However, despite adopting multi-objective formulations, existing UAV optimization studies remain focused primarily on non-medical operational criteria such as coverage quality, communication reliability, sensing performance, and energy efficiency. They do not incorporate victim severity, medical urgency, perception-derived triage information, or healthcare-oriented dispatch objectives into the optimization process. Furthermore, these formulations typically assume predefined mission priorities and do not integrate perception outputs from upstream AI-based UAV sensing systems. Consequently, while broader multi-objective UAV optimization literature provides important methodological foundations for multidimensional UAV decision making, existing formulations remain insufficient for UAV-assisted EMS scenarios requiring integrated severity-aware dispatch, triage-driven prioritization, and perception-informed optimization. This limitation further motivates the development of the unified perception to optimization UAV EMS framework proposed in this thesis.

2.4 Supporting Communication and Edge Infrastructure for UAV Emergency Response

Reliable communication and computing infrastructure are important enabling components of UAV-assisted emergency response systems, particularly in disaster environments where terrestrial networks may be degraded or unavailable. In practical deployments, perception outputs generated by UAV platforms must often be transmitted to remote command centers, optimization engines, or emergency coordinators for further analysis and decision making. Consequently, communication and edge computing frameworks play a supporting role in facilitating timely situational awareness and coordinated UAV-assisted response.

Prior research has investigated several communication approaches for UAV-enabled disaster management and emergency response systems. A major research direction focuses on connectivity restoration, where airborne communication through UAVs is considered as temporary communication infrastructure to restore network access over affected disaster regions [56], [54]. More recently, advanced UAV base station architectures operating over emerging communication paradigms, such as terahertz-band links, have also been investigated to improve post-disaster wireless coverage and high-capacity emergency communication support [57]. Complementary studies have explored fifth-generation (5G)-enabled and heterogeneous multi-radio access architectures to improve communication reliability, reduce latency, and enhance service robustness in UAV-assisted emergency deployments [58–60]. Additional public-safety communication research has further investigated localization and victim estimation-oriented emergency communication frameworks to support situational awareness in next-generation emergency response systems [61].

Beyond connectivity provision, recent work has increasingly emphasized distributed intelligence and computational offloading within UAV-assisted disaster systems. Edge computing and distributed processing frameworks have been introduced to reduce centralized processing dependence and enable near-real-time UAV data analysis. The authors of [62] proposed a UAV-assisted mobile edge computing (MEC) architecture for post-disaster emergency medical rescue. Furthermore, the authors of [28, 30] explored cloud and IoT-based disaster data acquisition frameworks for real-time disaster management support. Emerging digital twin-enabled UAV frameworks have further extended this direction by employing virtual replicas of UAV systems and disaster environments to support intelligent task offloading, adaptive resource management, and UAV coordination in dynamic disaster scenarios [4]. More recent work has also investigated semantic communication and bandwidth-efficient UAV networking strategies to improve transmission efficiency in communication-constrained disaster environments [63].

Collectively, the above-mentioned studies demonstrate that communication and infrastructure research has progressed from basic connectivity restoration toward increasingly intelligent, distributed, and adaptive UAV-enabled disaster support architectures. However, despite these advancements, existing communication and edge computing frameworks primarily optimize networking performance, computational efficiency, or system coordination rather than medical severity assessment or UAV dispatch optimization. Accordingly, communication layer design, bandwidth allocation, edge scheduling, and digital twin-assisted infrastructure management are considered outside the scope of this thesis. Nevertheless, these enabling technologies provide important deployment context for future real-world implementations of perception-driven UAV EMS systems.

2.5 Identified Research Gaps and Thesis Positioning

The reviewed literature demonstrates substantial progress in UAV-based aerial perception, AI-assisted triage, and UAV-enabled emergency medical service optimization. However, these research directions remain largely disconnected and insufficient for supporting integrated severity-aware UAV-assisted emergency response in disaster environments. In particular, several important limitations persist when the literature is examined collectively. The following subsections summarize the principal research gaps identified from the reviewed studies and position the proposed thesis relative to these unresolved challenges.

2.5.1 Absence of Severity-Aware Optimization in UAV-Based EMS

Although UAV-assisted EMS optimization has received growing research attention, existing frameworks remain predominantly logistics-oriented, with primary emphasis on cost and travel time. Strategic facility location allocation models such as ULAP [52] and related UAV EMS deployment formulations [46], [64] optimize resource placement and dispatch primarily with respect to operational cost, coverage, or response efficiency. While these objectives are important for practical deployment planning, they do not explicitly account for heterogeneous victim severity, medical urgency, or triage priority during UAV allocation and dispatch.

In disaster response settings, however, affected individuals and regions often exhibit significantly differing levels of medical urgency, which makes uniform treatment of all demand nodes operationally unrealistic. Despite this, current UAV EMS optimization models generally assume homogeneous demand importance or rely on static predefined demand weights rather than dynamically inferred victim-criticality information. Consequently, existing formulations remain insufficient for scenarios in which emergency resources must be selectively prioritized toward the most critical victims under constrained UAV availability. This reveals a fundamental limitation in the current UAV EMS optimization literature. Although optimization models effectively support cost-efficient UAV deployment and dispatch, they do not incorporate severity-aware or triage-aware prioritization mechanisms as core optimization objectives. To address this limitation, this thesis extends conventional UAV EMS optimization by embedding severity and priority directly into the multi-objective dispatch formulation, enabling UAV allocation decisions to account for both logistical efficiency and victim criticality.

2.5.2 Disconnection Between Perception and Optimization Pipelines

A second major limitation in the current literature is the disconnect between UAV perception systems and UAV optimization frameworks. Recent advances in UAV-based aerial perception have significantly improved the ability to detect and localize victims in disaster environments through DL-based detection and aerial vision models [15, 16, 32]. Related small-object detection and aerial feature-enhancement approaches have further improved perception robustness for challenging UAV imagery conditions [17, 18, 36, 37, 65]. In parallel, AI-driven triage and remote health-assessment studies have demonstrated the feasibility of estimating patient condition or medical urgency using AI-based decision-support frameworks [39, 40, 44].

Despite these advances, perception and optimization components remain largely isolated in existing UAV-assisted emergency response systems. Current optimization models typically assume that demand locations, severity levels, or service priorities are predefined and externally available, without explicitly incorporating real-time UAV-derived perception outputs into the optimization pipeline. Conversely, most UAV perception studies terminate at detection, localization, or severity estimation without integrating their outputs into downstream resource allocation or dispatch decision-making frameworks. As a result, the literature lacks an integrated perception to optimization pipeline capable of transforming real-time UAV perception outputs into actionable emergency medical dispatch decisions. This thesis addresses this gap by proposing a unified UAV-assisted EMS framework that integrates aerial victim detection, AI-based severity estimation, priority modeling, and multi-objective optimization within a single end-to-end decision-making pipeline.

2.5.3 Lack of Aerial Visual Severity Estimation for Disaster Triage

Although AI-based triage and severity assessment methods have shown promising results in healthcare and telemedicine applications, existing approaches remain predominantly designed for structured clinical or sensor-based environments [39, 40, 44]. In such cases, the explicit patient health information is available during assessments. Hence, these methods mainly rely on physiological measurements, medical telemetry, structured patient records, or direct clinical observations to estimate patient urgency or severity. Similarly, there has been substantial progress in UAV-based aerial perception literature, which relates to victim detection and localization [15, 16, 32, 36, 37]. However, in most cases, the perception systems terminate at identifying the presence or location of victims without attempting to infer medically relevant severity or urgency from visual evidence.

Overall, the insights reveal an important methodological gap between UAV perception and medical triage research. Specifically, the existing UAV perception methods provide localization without clinical prioritization, while existing AI-based triage methods assume structured health data that is usually unavailable in aerial disaster settings. To address this limitation, this thesis proposes an aerial visual severity estimation framework that infers victim criticality directly from UAV-derived visual observations and transforms these estimates into triage-aware decision signals for downstream optimization.

Table 2.1: Comparison of selected UAV/EMS works against the thesis objectives.

Paper	Publication year	Travel Time	Cost	Severity Cov.	Perception	Optimization
[16]	2024	×	×	×	✓	×
[17]	2025	×	×	×	✓	×
[18]	2025	×	×	×	✓	×
[27]	2025	✓	×	×	×	✓
[36]	2025	×	×	×	✓	×
[39]	2024	×	×	✓	×	×
[41]	2025	✓	×	✓	×	✓
[46]	2022	✓	✓	×	×	✓
[52]	2023	×	✓	×	×	✓
[66]	2025	✓	✓	×	×	✓
[67]	2024	✓	×	×	×	✓
Proposed	**–**	✓	✓	✓	✓	✓

Table 2.1 summarizes the comparative positioning of representative prior works relative to the core objectives addressed in this thesis. As shown, existing studies typically focus on isolated subsets of the overall UAV-assisted emergency response pipeline. Perception-oriented works primarily address victim detection and aerial image understanding without optimization capabilities, while optimization-focused studies emphasize travel time, cost, or resource allocation without incorporating perception-derived severity information. Similarly, AI-based triage methods consider

severity assessment but generally lacks integration with UAV-based aerial perception or downstream dispatch optimization. In contrast, the proposed framework uniquely combines perception, severity-aware prioritization, and multi-objective optimization within a unified UAV-assisted EMS decision-making pipeline.

2.6 Summary

This chapter has reviewed the existing literature relevant to UAV-assisted emergency medical response systems across four supporting research areas: UAV-based victim detection and aerial perception, AI-based severity and triage assessment, UAV EMS optimization and resource allocation, and enabling communication and edge computing infrastructure. The literature review has showed that substantial progress has been made in aerial victim detection, UAV-assisted healthcare monitoring, and UAV-enabled emergency medical logistics. However, existing studies largely treat these components in isolation. Current UAV perception literature primarily focuses on victim detection and localization without translating perception outputs into structured triage-aware decision signals [16, 17, 32]. Existing AI-based triage and severity assessment approaches are generally designed for structured clinical or physiological data settings and remain insufficient for aerial disaster reconnaissance scenarios relying on coarse visual observations [39,40,44]. Similarly, prevailing UAV EMS optimization models, including facility location allocation frameworks such as ULAP, are predominantly cost-driven and assume static homogeneous demand, without incorporating dynamically inferred severity or priority information from UAV perception systems [46, 52, 64]. Therefore, a significant research gap remains in the development of integrated UAV-assisted emergency medical response frameworks that unifies aerial perception, severity-aware victim prioritization, and optimization-based UAV dispatch within a single end-to-end decision-making pipeline. To address this gap, the next chapter presents the proposed UAV-based severity-aware EMS optimization framework developed in this thesis.

Chapter 3

Methodology

This chapter presents the methodological foundation of the proposed severity-aware uncrewed aerial vehicle (UAV)-based emergency medical service (EMS) framework. It is organized into two main parts. The first part introduces the system model, detailing the key components of the perception, severity inference, and decision-making pipeline. The second part formulates the multi-objective mixed-integer linear programming (MILP) problem and describes solution approaches used for implementation, training, and evaluation of the overall system.

3.1 System Model and Problem Formulation

The proposed system model provides a structured representation of a severity-aware UAV-based EMS framework for disaster response. It explicitly characterizes four key components. First, the physical elements, including UAVs, candidate facility locations, and affected individuals within the disaster environment. Second, the perception elements process UAV-acquired imagery using a vision module to detect individuals and estimate their criticality. Third, the spatial modeling elements, which aggregate detected individuals into geographically defined demand points (tile) and assign severity and priority scores. Lastly, the optimization elements determines facility activation and UAV dispatch decisions under operational constraints, including travel time, capacity, and budget. The framework assumes that severity and priority values are computed prior to optimization and remain fixed during decision-making process. Collectively, these components define a unified end-to-end pipeline that transforms raw UAV observations into structured, optimization-ready inputs for emergency response planning.

3.1.1 System Entities

UAV Fleet: The UAV fleet consists of autonomous aerial vehicles deployed from EMS facilities to support disaster response operations. In practical deployments, UAVs acquire real-time imagery using onboard sensors. However, in this work, the perception pipeline is evaluated using pre-collected aerial images from a UAV-based disaster dataset [68]. Each UAV operates under practical constraints, including limited battery capacity and a maximum operational range. During each dispatch cycle, a UAV is assigned to serve exactly one demand location (i.e., tile) and is required to return to its originating facility upon completion of the mission.

EMS Facilities: EMS facilities act as UAV launch and recovery bases. Each facility i may be activated ($y_i = 1$) or remain inactive ($y_i = 0$). Activating a facility incurs a fixed cost f_i and enables UAV deployment from that location, subject to capacity constraints. Each active facility can host a limited number of UAVs, and allocates u_i UAVs to serve demand locations (tiles). In addition, a per-UAV operational cost p_i is incurred for each deployed UAV.

Control Center / Decision-Making Unit: The control center represents the computational unit responsible for processing UAV imagery, running the perception and classification models, computing tile-level severity and priority scores, and solving the optimization model. The resulting decisions include facility activation, UAV allocation, and tile-service assignments. In this thesis, the control center is modeled as an offline decision-making module using pre-collected UAV images from the C2A dataset, rather than as a real-time communication or edge-computing system.

Disaster Region and Tiles: The disaster region is spatially partitioned into a finite set of tiles $J = \{1, 2, \ldots, 16\}$, representing a 4×4 grid the affected area. Each tile corresponds to a localized region that may contain an unknown number of affected individuals and varying levels of damage, environmental complexity, and operational risk. Each tile j is associated with a travel time T_{ij} from facility i, which influences UAV dispatch decisions. A tile is considered served if at least one UAV is assigned to it, as represented by the binary variable z_j. The terms "tile" and "demand point" are used interchangeably in this thesis.

Spatial Representation and Travel Time Modeling: The disaster region is represented using a normalized grid-based coordinate system derived from image tiling. The input image is partitioned into a fixed grid (e.g., 4×4), where each tile corresponds to a spatial location. The center of each tile defines the coordinates of a demand point. Candidate facility locations are defined within the same coordinate system to ensure consistency in spatial modeling.

The spatial distance between facility i and demand point j is computed using Euclidean distance:

$$d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}, \tag{3.1}$$

where (x_i, y_i) and (x_j, y_j) denote the coordinates of facility i and tile j, respectively.

The travel time between facility i and tile j is defined as:

$$T_{ij} = \frac{d_{ij}}{v}, \tag{3.2}$$

where v denotes the UAV speed.

The dataset lacks geo-referenced spatial information, so the coordinate system has been normalized. As a result, distances are represented in relative units instead of actual physical units. Travel time is also measured in relative time units. Here, one unit corresponds to the distance between the centers of adjacent tiles in the grid.

The overall workflow of the proposed severity-aware UAV-based EMS framework is illustrated in Fig. 3.1. The framework consists of three major modules: the vision layer, spatial triage layer, and optimization layer. In the vision layer, UAV-captured aerial imagery is processed to detect affected individuals and estimate their criticality levels using computer vision models. The spatial triage layer aggregates individual-level detections into tile-level severity, urgency, and priority metrics representing localized emergency demand. These aggregated triage scores are then passed to the optimization layer, which determines UAV facility activation and dispatch decisions to allocate limited emergency resources efficiently while prioritizing high-risk regions.

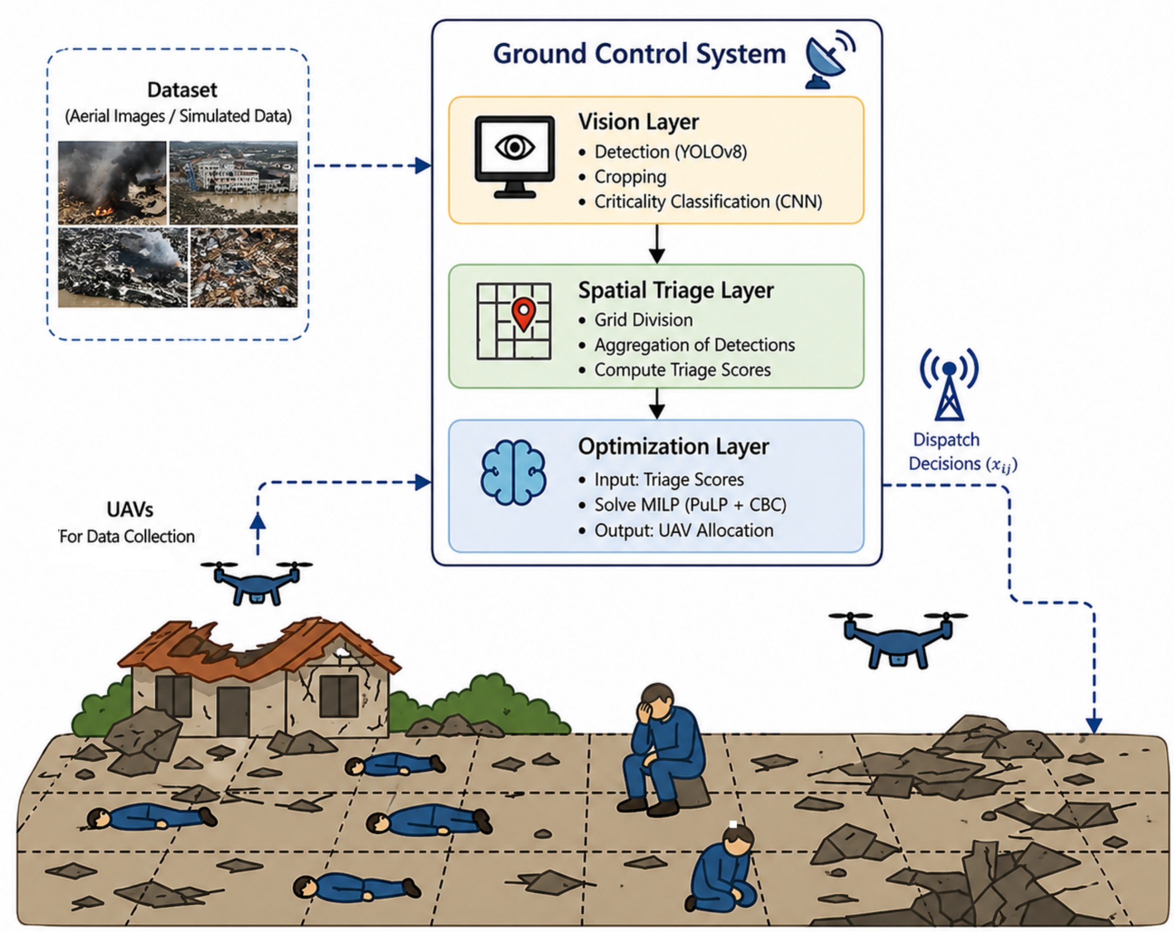

Figure 3.1: Architecture of the proposed severity-aware UAV-based EMS system.

Affected Individuals: Affected individuals are detected from UAV-acquired imagery and represented as discrete entities within the disaster environment. Detection is performed using a vision-based model that produces bounding boxes corresponding to individual locations. The spatial position of each individual is approximated using the center of the bounding-box, which is then used to map the individual to a corresponding tile. These detected individuals are subsequently aggregated at the tile-level, forming demand points that serve as the fundamental units for severity estimation, priority computation, and optimization-driven dispatch decisions. This aggregation enables the transition from individual-level perception outputs to structured inputs suitable for system-level decision-making.

Disaster Environment: The disaster environment is inherently heterogeneous, characterized by structural damage, occlusions, shadows, debris, and varying illumination conditions. Affected individuals may be distributed sparsely across the region or concentrated in localized clusters, leading to non-uniform demand patterns. Such environmental complexity introduces uncertainty in perception, as visual clutter and degraded conditions can impact detection and classification accuracy. Furthermore, during real-world deployment, ground-truth labels are not available, and the system must rely entirely on model-based inference to estimate severity and priority. This reflects realistic operational conditions in post-disaster scenarios.

System Assumptions: The proposed framework operates under the following assumptions:

- Affected individuals are assumed to remain stationary during observation and decision-making process [32, 35].
- Each UAV serves exactly one tile per dispatch event.
- Spatial coordinates are derived from UAV image grids and expressed in pixel units. As a result, distances and travel times are treated as relative values rather than real-world physical quantities.
- Intra-tile navigation and fine-grained movement within a tile are not explicitly modeled.
- Travel time T_{ij} between facility i and tile j is computed based on Euclidean distance [52].
- UAV battery limitations are enforced through a maximum allowable travel time $T_{\max}$ constraint [51].
- Severity and priority values are computed prior to optimization and remain fixed during the decision-making process.
- Tiles are treated as independent demand units within the optimization model.

- Communication delays and data transmission constraints are not explicitly modeled [34].
- The proposed framework is evaluated using a pre-collected dataset of UAV imagery rather than a real-time deployment. As a result, this thesis does not explicitly model system-level latency components such as image acquisition time, onboard processing delays, and communication time to medical responders.

3.2 Perception and Spatial Triage Modeling

The perception component detects affected individuals and estimates their criticality, while the spatial triage component aggregates individual-level predictions into tile-level severity and priority scores. These aggregated representations from the key inputs to the optimization model presented in the subsequent section.

The proposed framework integrates multiple computational components across perception, learning, spatial aggregation, and optimization stages. These modules operate sequentially, forming an end-to-end pipeline that converts raw aerial imagery into structured decision inputs for UAV dispatch. The framework combines deep learning–based perception models with classical machine learning baselines and an optimization layer to enable severity-aware decision-making.

- **You only look once version 8 (YOLOv8)-based detection module:** A deep learning-based object detection model is used to identify affected individuals in aerial imagery and estimate their spatial locations through bounding boxes [69].
- **EfficientNet-Lite0 criticality classifier:** A convolutional neural network is employed to estimate the probability of critical condition for each detected individual based on visual features extracted from the corresponding image regions [70].
- **Random forest (RF) and support vector machine (SVM) baselines:** Classical machine learning models are trained on feature embeddings extracted from the neural network to provide alternative criticality estimation approaches for comparative analysis [71, 72].

- **Spatial triage module:** This module aggregates individual-level detections and predictions into tile-level representations and computes severity and priority scores that captures both medical relevance and spatial demand characteristics.
- **MILP Optimization Module:** A mixed-integer linear programming model determines facility activation, UAV allocation, and tile-level service decisions under operational constraints, including travel time, capacity, and budget.

3.2.1 Vision-Based Perception

The perception component processes aerial imagery to identify affected individuals and estimate their level of criticality. This process is structured in two stages: (i) a supervised learning stage, where a criticality model is trained using labeled data, and (ii) a deployment stage, where the trained model is applied to unseen disaster imagery. This design enables generalization to real-world scenarios while operating without access to ground-truth labels during deployment.

Supervised learning (SC1): In the supervised learning stage, a binary criticality classifier is trained using labeled person-level data. The input consists of image crops of detected individuals along with posture-based labels. Multi-class posture labels are mapped into a binary representation:

$$y_k^{(\text{bin})} = \begin{cases} 0, & \text{if posture is upright,} \\ 1, & \text{if posture is lying, bent, sitting, or kneeling,} \end{cases} \tag{3.3}$$

where $y_k^{\text{bin}} \in \{0, 1\}$ denotes the binary criticality label of the k-th detected individual, and k represents the index of an individual detected in the image.

A convolutional neural network (EfficientNet-Lite0) is trained to estimate the probability of critical condition:

$$p_k = P(\text{critical} \mid \text{crop}_k), \tag{3.4}$$

where p_k denotes the predicted probability that the k-th detected individual is in a critical posture, given the corresponding image crop_k, and the output of this stage is a trained model that produces a criticality probability p_k for each individual. The dataset used for training (SC1) is constructed exclusively from training and validation images, ensuring no information leakage from the test set. To enable comparative analysis, feature embeddings extracted from the neural network are also used to train classical machine learning models, including random forest (RF) and support vector machine (SVM), providing alternative criticality estimation baselines.

Deployment and Inference (SC2): In the deployment stage, the trained perception model is applied to unseen aerial imagery representing disaster environments. The input images are partitioned into a grid structure to enable localized processing and improve small-object detectability. Detection is performed independently on each spatial tile using YOLOv8-based model. The resulting bounding boxes, initially expressed in tile-level coordinates, are transformed into the global image coordinates using spatial offsets. Since tile-wise processing may produce duplicate detections, particularly near tile boundaries, global non-maximum suppression (NMS) is applied to remove redundant bounding boxes and retain a single detection per individual. The final detections are used to extract person-level image crops, which are processed by the trained classifier to estimate the criticality probability p_k for each individual. The output of this stage consists of detected individuals, their spatial locations, and associated criticality probability p_k which serve as inputs to the spatial triage module. Overall, the vision layer operates as a two-stage pipeline combining supervised learning (SC1) and deployment inference (SC2), integrating YOLOv8-based detection with convolution neural network (CNN)-based criticality estimation.

The perception layer relies on a pre-collected dataset of UAV disaster images rather than real-time image acquisition. Specifically, we use the C2A dataset [68], which offers aerial views of disaster scenarios with annotated human poses. Each image is treated as a snapshot of a disaster scene, from which we extract spatial demand information. Consequently, image capture and wireless transmission processes are not explicitly modeled within the system. The computational latency associated with the perception and decision-making pipeline, including detection, classification, and optimization, is not explicitly modeled in this study. The primary focus is on evaluating the effectiveness of the proposed decision-making framework based on perception-derived inputs.

Accordingly, no explicit constraints on processing time are imposed in the system formulation. This modeling choice aligns with simulation-based studies on UAV-assisted emergency response, which emphasize allocation efficiency and prioritization performance rather than end-to-end system latency. The computed severity, urgency, demand, and priority metrics are therefore directly used as inputs to the optimization model, without incorporating delays associated with image acquisition, communication, or processing.

3.2.2 Spatial Triage Modeling

The spatial triage component converts individual-level perception outputs into structured, tile-level representations suitable for decision-making. The disaster region is partitioned into a predefined grid, where each tile represents a localized region of the affected area.

Detected individuals are assigned to tiles based on their spatial locations. This mapping establishes a direct link between perception outputs and tile-level aggregation units. By aggregating individual-level information, each tile becomes a demand point characterized by severity and priority scores, which serve as inputs to the optimization model.

Tile-level statistics: To quantify the distribution and severity of affected individuals within each tile, the following statistics are computed:

$$L_j = |k_j|, \tag{3.5}$$

$$C_j = \sum_{k \in K_j} \mathbf{1}(p_k > 0.5), \tag{3.6}$$

$$\rho_j = \max_{k \in K_j} p_k, \tag{3.7}$$

$$\bar{p}_{\text{crit},j} = \frac{1}{L_j} \sum_{k \in K_j} p_k, \tag{3.8}$$

where L_j denotes the total number of individuals in tile j (local demand),K_j denotes the set of individuals detected in tile j, and k indexes each detected individual within the tile. C_j denotes the number of individuals classified as critical, ρ_j represents the maximum criticality probability among individuals in the tile, and $\bar{p}_{\text{crit},j}$ denotes the average criticality probability within the tile, p_k denotes the predicted probability that the k-th detected individual is in a critical posture, as obtained from the trained classifier given the corresponding image crop_k.

These statistics capture both the quantity and severity distribution of affected individuals and form the basis for triage-aware scoring.

Severity and priority modeling: To quantify the operational importance of each tile, severity, urgency, demand, and priority scores are defined based on the computed statistics. These scores capture complementary aspects of the disaster scenario, including injury intensity (severity), time sensitivity (urgency), operational burden (demand), and decision-level importance (priority).

The formulation is inspired by prior work on triage-aware emergency response. For example, the authors of [41] incorporated severity-based prioritization into robotic planning, while [39] showed the use of computational models for estimating patient urgency. However, such approaches typically rely on discrete triage categories or physiological data not available in UAV-based settings. Therefore, this work proposes a continuous, perception-driven scoring framework. The scores are defined as:

$$S_j = \alpha_1 L_j + \alpha_2 C_j + \alpha_3 \rho_j + \alpha_4 \bar{p}_{\text{crit},j}, \tag{3.9}$$

$$U_j = \theta_1 T_j + \theta_2 E_j + \theta_3 S_j, \tag{3.10}$$

$$D_j = \beta_1 C_j + \beta_2 S_j, \tag{3.11}$$

$$P_j = \lambda D_j + (1 - \lambda) U_j, \tag{3.12}$$

where $\alpha_1, \alpha_2, \alpha_3, \alpha_4$, β_1, β_2, $\theta_1, \theta_2, \theta_3$, and λ are non-negative weights that determine the relative influence of each component based on operational priorities. The weight parameters α, β, θ, and λ were assigned fixed values to illustrate an engineered triage-prioritization strategy rather than clinically validated medical scores. Higher weights were assigned to features associated with more severe critical conditions, because these components are more likely to indicate immobility, injury, or reduced ability to self-evacuate. The purpose of assigning these weights is to guide the optimization model toward identifying high-risk regions, especially when UAV resources are limited.

The severity score S_j captures injury intensity within the j-th tile by combining victim density L_j, critical casualty count C_j, maximum criticality likelihood ρ_j, and average criticality confidence $\bar{p}_{\text{crit},j}$. The urgency score U_j incorporates factors like the time since report (T_j), and the distance from the epicentre (E_j), and the severity score (S_j), reflecting the time-sensitive nature of emergency intervention.The demand score D_j captures the operational burden associated with the j-th tile by combining the critical victim count (C_j) and the severity level (S_j). The priority score P_j determines the importance of the j-th tile for UAV dispatch and resource allocation decisions by integrating demand (D_j) and urgency (U_j).

In deployment, the urgency score is simplified as:

$$U_j = \theta_3 S_j. \tag{3.13}$$

These tile-level scores provide a quantitative representation of disaster conditions and serve as key inputs to the optimization model presented in the next section.

Posture-to-criticality mapping: The posture-to-criticality mapping adopted in this work is grounded in established triage principles, where prioritization is based on observable condition and functional status. In mass-casualty incidents, triage commonly relies on mobility, level of consciousness, and visible signs of life to assign priority levels.

In UAV-based scenarios, direct physiological measurements (e.g., respiration and pulse) are not available; therefore, assessment must rely on visual cues. Recent work on drone-assisted triage proposes remote algorithms that evaluate observable indicators such as walking ability, consciousness, and signs of life for victim classification [73]. Recent studies have also explored AI-driven triage approaches that utilize patient data and machine learning models to prioritize, highlighting the role of data-driven decision-making in emergency response [39]. In particular, mobility (i.e., the ability to walk) is used as a key decision factor in prioritization.

Accordingly, posture inferred from aerial imagery is employed as a practical proxy for mobility and observable condition. The proposed posture-to-criticality mapping is thus a heuristic derived from established triage practice and adapted to the constraints of vision-based, remote assessment. However, this mapping does not represent a clinically validated assessment and is considered a heuristic approximation.

3.3 From Spatial Triage to UAV Allocation Problem

The spatial triage process assigns each demand tile$j \in J$ a severity score S_j and a priority score P_j, derived from the perception pipeline. These scores quantify the medical urgency and relative importance of each tile within the disaster environment. Given these triage-aware scores, the UAV dispatch problem involves determining how limited fleet of UAVs should be allocated across multiple demand locations. Each allocation decision must satisfy practical constraints, including UAV availability, travel time limits, and deployment cost. At the same time, the allocation should prioritize tiles with higher severity and priority, rather than relying solely on operational efficiency.

This leads to the central optimization problem of this thesis: to allocate UAVs across demand locations that balances resource constraints with triage-driven priorities.

A commonly used optimization framework for UAV facility location and allocation is the uncrewed aerial vehicles for emergency management system (UEMS) location-allocation problem (ULAP) [52], which is designed for cost-efficient deployment of UAV resources. However, ULAP does not account for victim severity or priority, making it insufficient for real-time, triage driven disaster response. In its classical robust standard formulation (RSF), ULAP incorporates demand uncertainty by considering worst-case demand scenarios. This is achieved through additional variables and constraints that ensure feasibility under variations in demand. A key feature of ULAP is its coverage requirement, which ensures that every demand location must be served by at least one facility:

$$\sum_{i \in I} x_{ij} \geq 1, \quad \forall j \in J, \tag{3.14}$$

where x_{ij} is a binary decision variable if demand point j is served by facility i and I denotes the set of candidate UAV facility locations, and J denotes the set of demand point. This constraint enforces a full coverage requirement, ensuring that every demand tile is served by at least one facility. While appropriate for planning applications, this assumption is often unrealistic in disaster-response settings with limited resources.

Several limitations arise when applying ULAP to severity-aware UAV dispatch. First, ULAP is fundamentally cost-driven and does not incorporate severity and priority. As a result, all demand locations are treated uniformly, ignoring differences in medical urgency. Second, in real disaster scenarios, limited UAV availability and operational constraints make it impractical to serve all demand locations. Therefore, prioritization is necessary. Third, ULAP relies on predefined demand estimates and bounded uncertainty, which are better suited for pre-disaster planning rather than dynamic, perception-driven environments. These limitations highlight that the classical ULAP formulation does not adequately capture the requirements of triage-aware UAV dispatch. This motivates the development of an optimization framework that explicitly incorporates severity and priority into allocation decisions.

3.3.1 Cost-Focused Allocation Baseline (ULAP-Inspired)

To enable a meaningful comparison, a modified cost-focused allocation baseline is constructed based on the objective structure of ULAP. This baseline isolates cost-driven decision making and serves as a reference for evaluating the benefits of triage-aware optimization.

The baseline objective minimizes the total cost associated with facility activation and UAV deployment:

$$\min \sum_{i \in I} f_i y_i + \sum_{i \in I} p_i u_i, \tag{3.15}$$

where I denotes the set of candidate UAV facility locations and J denotes the set of demand tiles. The parameter f_i represents the fixed cost of activating facility i, and p_i denotes the cost associated with deploying a single UAV from facility i.The binary decision variable $y_i \in \{0, 1\}$ indicates whether facility i is activated or not, and u_i represents the number of UAVs deployed from facility i.

To reflect the limited-resource, selective-dispatch setting considered in this thesis, the classical ULAP full-coverage requirement is replaced with a selective-service dispatch constraint:

$$\sum_{i \in I} \sum_{j \in J} x_{ij} = K, \tag{3.16}$$

where the binary decision variable x_{ij} describes if a UAV from facility i is assigned to serve demand point j, and k denotes the total number of UAV dispatches available within the operational horizon.

To ensure that each selected demand point receives at most one UAV dispatch, the following constraint is imposed:

$$\sum_{i \in I} x_{ij} \leq 1, \quad \forall j \in J. \tag{3.17}$$

These modifications replace the classical ULAP full-coverage requirement, which is incompatible with resource-constrained disaster scenarios. As a result, UAV allocation decisions in this baseline are driven solely by operational cost considerations and do not incorporate severity or priority information. The resulting formulation serves as a cost-focused comparator, rather than a full reproduction of the classical ULAP RSF model. It enables a clear and controlled evaluation of the impact of incorporating triage-aware information into UAV dispatch decisions.

3.4 Optimization Problem Formulation

This section formulates the UAV-based emergency response problem which is a MILP model. The objective is to determine optimal facility activation and UAV dispatch decisions for serving demand points derived from the spatial triage stage. The formulation incorporates tile-level severity and priority scores to ensure that high-risk regions are prioritized under limited resources.

3.4.1 Decision Variables

We define the following decision variables:

$$x_{ij} = \begin{cases} 1, & \text{facility } i \text{ serves tile } j, \\ 0, & \text{otherwise,} \end{cases}$$

$$y_i = \begin{cases} 1, & \text{facility } i \text{ activated,} \\ 0, & \text{otherwise,} \end{cases}$$

$$u_i \in \mathbb{Z}_+ : \text{number of UAVs at facility } i,$$

$$z_j = \begin{cases} 1, & \text{tile } j \text{ is served,} \\ 0, & \text{otherwise,} \end{cases}$$

where x_{ij} represents the assignment of UAVs from facility i to demand tile j, y_i indicates whether a facility is activated or not, u_i specifies the number of UAVs available at each facility, and z_j denotes whether a demand tile is served. These variables collectively determine facility activation, UAV allocation, and demand coverage decisions. Each UAV is assumed to serve exactly one tile per dispatch cycle. Therefore, the total number of served tiles equals the total number of available UAVs N, establishing a one-to-one mapping between the UAVs and served tiles.

3.4.2 Parameters

The model uses the following parameters.

- T_{ij}: travel time from facility i to demand point/tile j
- S_j: severity score of demand point/tile j
- P_j: priority score of demand point/tile j
- f_i: fixed cost of activating facility i

- p_i: cost per UAV allocated from facility i
- N_i^{UAV}: maximum UAV capacity at facility i
- $T_{\max}$: maximum allowable round-trip travel time
- $B_{\max}$: total available budget
- $N_{\max}$: maximum number of facilities that can be activated
- N: total number of UAVs available

3.4.3 Objective Function

The objective balances operational efficiency with triage-driven priorities:

$$\min\left(w_1\sum_{i,j}T_{ij}x_{ij}+w_2\sum_i(f_iy_i+p_iu_i)-w_3\sum_j S_jz_j-w_4\sum_j P_jz_j\right), \tag{3.18}$$

where the first term minimizes total travel and the second term minimizes operation cost (facility activation and UAV deployment). The third and fourth terms promote servicing tiles with higher severity and priority. The weights $w_1, w_2, w_3, w_4 \geq 0$ control the trade-off between efficiency and triage-driven prioritization.

3.4.4 Constraints

The following constraints ensure feasibility and enforce operational requirements of the system.

$$\begin{aligned}
&\textbf{C1:} \quad \sum_i y_i \leq N_{\max} \\
&\textbf{C2:} \quad \sum_i x_{ij} = z_j,\ \forall j \\
&\textbf{C3:} \quad x_{ij} \leq y_i,\ \forall i, j \\
&\textbf{C4:} \quad \sum_j x_{ij} \leq u_i,\ \forall i \\
&\textbf{C5:} \quad u_i \leq N_i^{UAV}\, y_i,\ \forall i \\
&\textbf{C6:} \quad 2T_{ij}x_{ij} \leq T_{\max},\ \forall i, j \\
&\textbf{C7:} \quad \sum_i (f_i y_i + p_i u_i) \leq B_{\max} \\
&\textbf{C8:} \quad \sum_i u_i = K \\
&\textbf{C9:} \quad \sum_j z_j = K
\end{aligned}$$

C1 limits the number of activated facilities to the maximum allowable value N_{max}. C2 links assignment and service decisions; a tile is served if and only if it is assigned. C3 enforces that demand point j can only be assigned to facility i if facility i is activated. C4 limits the number of demand points assigned to facility i by its allocated UAVs. C5 ensures that UAV allocation at facility i should not exceed its maximum capacity and is permitted only when the facility is activated. C6 enforces that the round-trip travel time for UAV operations does not exceed the maximum allowable limit T_{max}. C7 enforces the total budget constraint, i.e., facility activation and UAV allocation should not exceed the available budget B_{max}. C8 fixes the total number of UAVs deployed in the system to k. C9 ensures that exactly k demand points are served (one UAV per demand point/tile).

The decision variables x_{ij}, y_i, and z_j are binary, and u_i is a non-negative integer. All constraints are linear, and the objective function is linear in the decision variables. Therefore, the problem MILP problem. By incorporating severity and priority directly into the objective function, the formulation enables triage-aware UAV allocation, ensuring that limited resources are directed toward high-risk regions while satisfying operational constraints.

3.5 Solution Approaches

This section presents the implementation of the proposed severity-aware UAV-based EMS framework, describing the computational pipeline used to operationalize the system model and solve the formulated optimization problem. The pipeline integrates perception, spatial triage, and optimization components, transforming raw aerial imagery into actionable UAV deployment decisions.

3.5.1 Vision Layer Implementation

The vision layer serves as the perceptual backbone of the system, transforming raw aerial imagery into structured, person-level criticality information for downstream processing. It is implemented as a two-stage pipeline consisting of a supervised training phase (SC1) and a deployment inference phase (SC2), enabling both learning from labeled data and generalization to unseen disaster environments.

CNN-Based Criticality Inference: The vision layer identifies individuals in UAV imagery and extracts person-level regions using bounding boxes. Each cropped region is resized and passed to a CNN (EfficientNet-Lite0), which outputs a probability score indicating whether the individual is in a critical condition. Based on this score, each person is classified as critical or non-critical, enabling automated victim assessment for downstream prioritization.

The process of CNN-based criticality inference is illustrated in Fig. 3.2. The figure demonstrates how the vision layer operates on UAV imagery to assess the condition of detected individuals. At first, we obtain images from the UAV dataset, where individuals are localized using YOLOv8 detection. YOLOv8 detection provides us with bounding boxes, and then we get the crops. These cropped images are then passed to a convolutional neural network, specifically the EfficientNet-Lite0 model, which classifies the individual as critical or non-critical based on visual features. This step enables automated and scalable assessment of victim conditions from aerial imagery, forming the basis for downstream severity estimation and resource allocation.

Fig. 3.3 illustrates the overall computational pipeline of the proposed severity-aware UAV-

(a) UAV Image

(b) zoomed view of a crop

(c) CNN Inference

Figure 3.2: CNN-based criticality inference.

based EMS framework. The framework is organized into three main layers: the vision layer, the spatial triage layer, and the optimization layer, which together enable the transformation of raw UAV-acquired imagery into structured inputs for emergency response.

In the vision layer, the system operates in two stages: supervised training (SC1) and deployment inference (SC2). In SC1, a criticality classifier based on EfficientNet-Lite0 is trained using labeled person-level image crops, where posture-based annotations are mapped into a binary criticality representation. Feature embeddings extracted from the trained network are further used to train RF and SVM models for comparative evaluation. In SC2, unseen UAV images are processed using a YOLOv8-based detector to identify individuals, followed by EfficientNet-Lite0 inference to estimate criticality probabilities for each detected individual. The trained models from SC1 are transferred to SC2, as indicated by the dashed connections in Fig. 3.3, ensuring that inference is performed using models learned exclusively from training data.

The outputs of the vision layer are aggregated in the spatial triage layer, where detected individuals are mapped to a predefined 4×4 grid of tiles based on their spatial locations. This spatial partitioning enables localized analysis of disaster conditions across the affected area. For each tile, statistical measures such as the total number of detected individuals (L_j), number of critical individuals (C_j), maximum criticality probability (ρ_j), and average criticality probability ($\bar{p}_{\text{crit},j}$)

are computed. These statistics capture both the quantity and condition of individuals within each region. Using these measures, severity, urgency, demand, and priority scores are derived, which collectively represent the medical importance and operational relevance of each spatial region.

Finally, these triage-aware scores are provided as inputs to the optimization layer, where a MILP model determines UAV allocation and dispatch decisions under operational constraints. This includes decisions related to resource allocation and prioritization across different regions. Overall, this layered structure enables the transformation of raw UAV imagery into structured, decision-ready inputs for prioritized emergency response.

Stage SC1: Supervised Training Environment: In the SC1 stage, a binary criticality classifier is trained using cropped images of individuals obtained from UAV datasets with pose annotations. Each individual is assigned a posture-based label, which is mapped into a binary criticality representation, where upright postures are treated as non-critical and all other postures are considered critical. This mapping enables the model to focus on identifying potentially injured or vulnerable individuals.

A convolutional neural network (CNN)-based on EfficientNet-Lite0 is trained to estimate the criticality probability p_k for each individual. The model takes a person-level image crop as input and outputs the likelihood of a critical condition.

To enable comparative evaluation, feature embeddings extracted from the trained network are used to train classical machine learning models, including RF and SVM. These models serve as baseline approaches for criticality estimation. The SC1 dataset is constructed exclusively from training and validation samples, ensuring that no information from the test set is used during model training. The output of this stage is trained model capable of producing probabilistic criticality estimates.

Stage SC2: Deployment and Inference: The SC2 stage represents the deployment environment, where the trained models are applied to full UAV images without access to ground-truth annotations. SC2 processes full-resolution aerial imagery and performs detection, classification, aggregation, and decision-making in an integrated pipeline. The input to SC2 consists exclusively

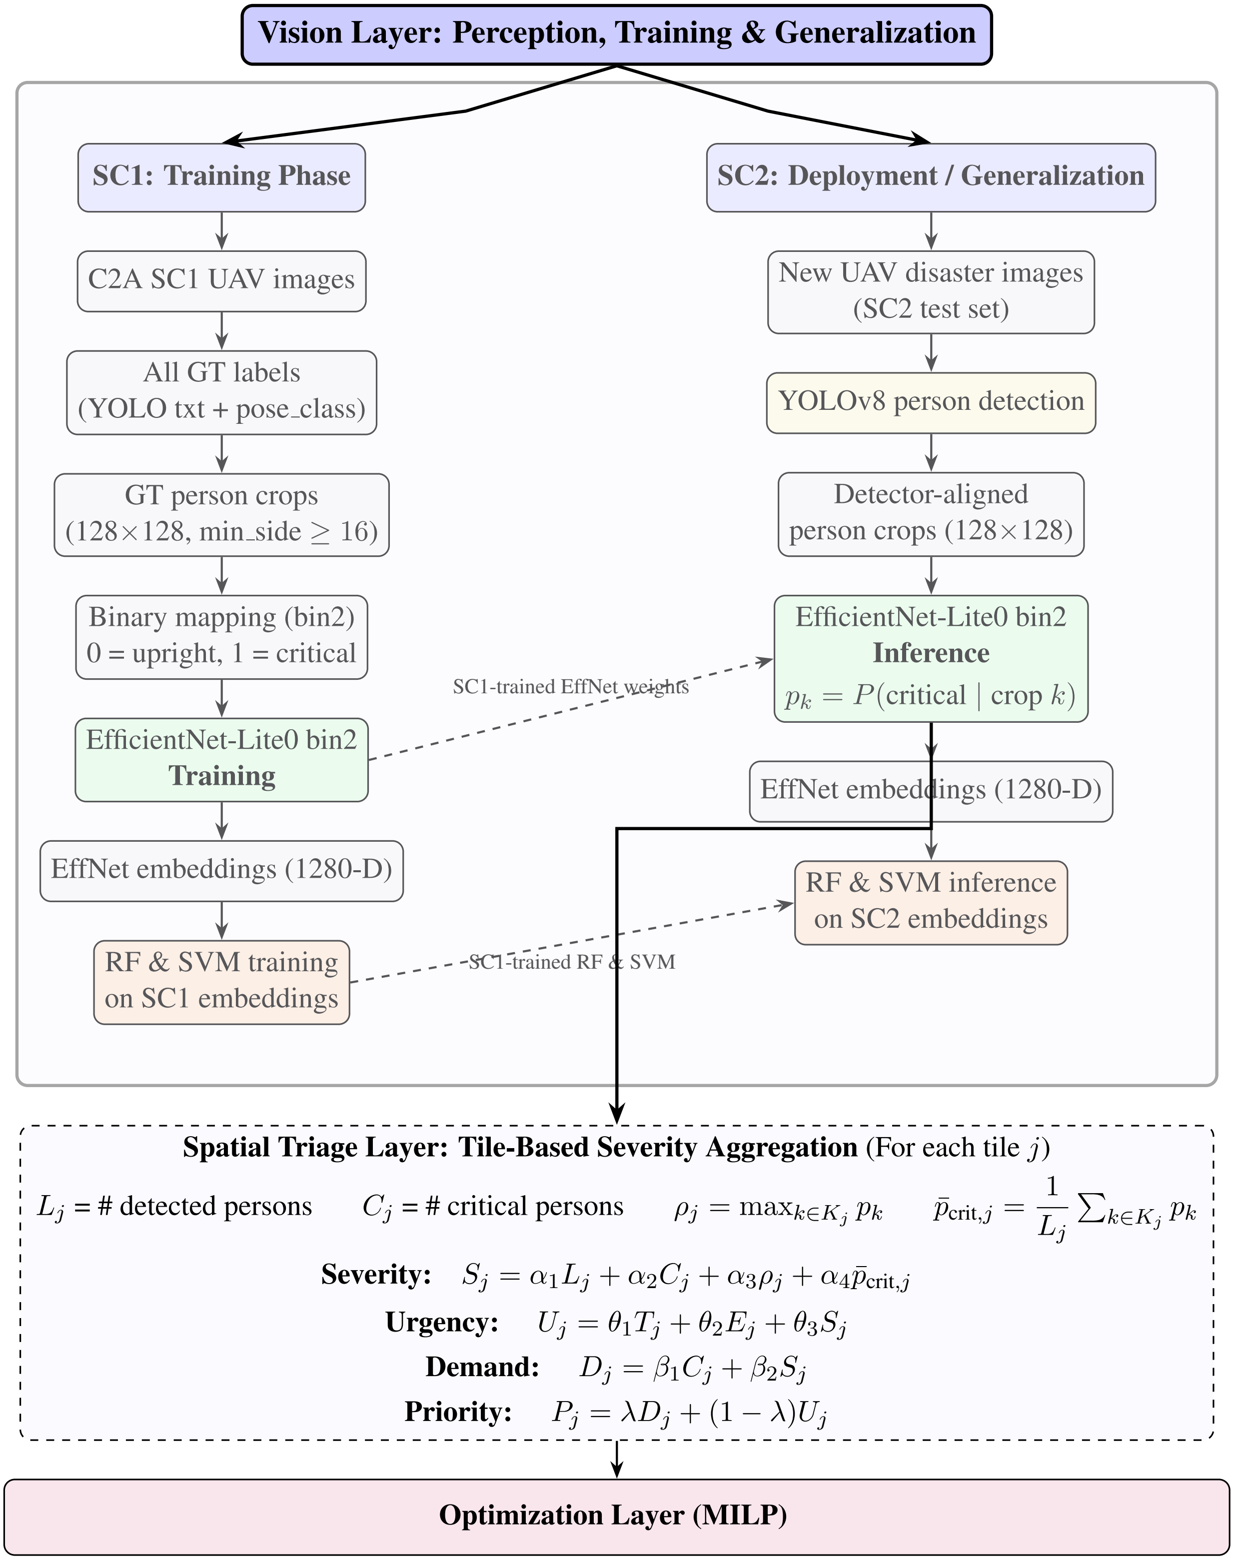

Figure 3.3: Architecture of the proposed severity-aware UAV-based EMS framework, illustrating the vision layer, spatial triage layer, and optimization layer.

of full-resolution UAV images from the held-out test split. These images are strictly excluded from all training processes, ensuring that evaluation reflects true generalization performance without any data leakage. To address the small scale of human instances in aerial imagery, each image is partitioned into a grid of non-overlapping tiles. This tiling strategy increases the relative size of objects within each tile, thereby improving detection performance for small and distant individuals.

Detection is performed independently on each tile using a YOLOv8-based model, improving small-object detectability. The detector is trained on the original training split and validated on the validation split using person bounding-box annotations. The model is trained solely for person detection and does not utilize pose or criticality labels, ensuring that no test data is used during training. Since detections are performed in tile-local coordinate systems, each detected bounding box is transformed into the global coordinate system of the original image using spatial offsets corresponding to the tile position. Tile-wise detection may produce overlapping bounding boxes for the same individual, particularly near tile boundaries. To address this, all detections are aggregated in the global coordinate space, and global non-maximum suppression (NMS) is applied to remove redundant detections based on an intersection-over-union (IoU) threshold. This ensures a single bounding box represents each individual.

The final set of bounding boxes is used to extract person-level crops directly from the original image. This preserves spatial detail and avoids distortions introduced by tile resizing. Each extracted crop is processed using the SC1-trained EfficientNet-Lite0 classifier to estimate the criticality probability p_k for each detected individual. Each detected individual is assigned to a tile based on its spatial location. For each tile, aggregated statistics are computed, including the total number of detected individuals, the number of critical individuals, and the average criticality probability. Using the aggregated statistics, severity, urgency, demand, and priority scores are computed according to Equations 3.9–3.12. These values capture both the quantity and condition of individuals within each tile. The computed tile-level scores are used as inputs to the optimization model. The optimization stage determines facility activation, UAV allocation, and assignment decisions while satisfying operational constraints, including travel time, UAV capacity, and budget limitations.

Together, these actions ensure that raw UAV imagery is transformed into structured tile-level decision inputs for triage-aware optimization.

3.5.2 Spatial Triage Layer Implementation

The outputs of the vision layer consist of individual detections and their associated criticality probabilities. However, these outputs cannot be directly used in the optimization model for several reasons: (i) the optimization model operates on spatial demand points rather than individual detections, (ii) UAVs are deployed to regions rather than individual bounding boxes, (iii) disaster scenes exhibit heterogeneous spatial distributions of affected individuals, and (iv) decision-making must consider collective risk rather than isolated cases. Therefore, a spatial aggregation mechanism is applied to transform individual-level detection outputs into structured demand points. The spatial triage layer performs this transformation by aggregating individual detections into tile-level severity and priority scores.

Spatial Partitioning and Assignment: Each UAV image is partitioned into a fixed grid structure (e.g., 4×4), where each grid cell represents a candidate demand point. Each detected individual is assigned to a tile based on the center of its bounding box, resulting in a grouping of individuals within each tile. This grouping enables the computation of tile-level statistics required for triage modeling.

Computation of Tile-Level Severity: For each tile, aggregated statistics are computed from the assigned individuals, including the number of detected individuals, the number of critical individuals, and summary measures of the predicted criticality probabilities. These statistics are used to compute severity score S_j for each tile using Equation 3.9. This severity score captures the overall intensity of risk within a tile, incorporating both the concentration of individuals and likelihood of critical conditions. By integrating these components, the severity score provides a robust representation of the urgency associated with each spatial region.

The relative contribution of each component is controlled by weighting parameters defined in the system model. These weights allow the model to adjust the influence of density-based and probability-based factors, enabling flexible tuning of the severity representation according to operational priorities.

Computation of Urgency, Demand, and Priority: Based on the computed severity values, additional triage quantities are derived to support decision-making. Urgency is computed by combining severity with contextual factors such as temporal or environmental conditions using Equation 3.10, reflecting how critical and time-sensitive a region is. Demand represents the estimated service requirement within each tile and is derived from both the number of affected individuals and their associated severity levels using Equation 3.11. Priority integrates both demand and urgency to determine the relative importance of each region for UAV allocation using Equation 3.12.

The relative influence of these components is controlled by weighting parameters defined in the system model. These weights regulate the contribution of severity, contextual factors, and demand in determining urgency and priority, allowing the model to balance different operational objectives. These computed values are then used as direct inputs to the optimization model.

3.5.3 Optimization Layer Implementation

This section describes the implementation of the optimization problem formulated in Section 3.4. The objective of this stage is to determine optimal UAV deployment and assignment decisions based on the triage information generated by the spatial aggregation layer. The optimization operates on tile-level inputs, including severity, urgency, demand, and priority scores, and produces allocation decisions under operational constraints.

The formulation follows Section 3.4 and determines: (i) which facilities to activate, (ii) how many UAVs to allocate to each facility, (iii) which tiles to serve, and (iv) how UAVs are assigned to selected tiles. The decision variables include assignment variables x_{ij}, facility activation variables y_i, UAV allocation variables u_i, and service indicators z_j. The relationship between assignment and service decisions is enforced through the constraint:

$$\sum_{i} x_{ij} = z_j \quad \forall j, \tag{3.19}$$

which ensures that a demand point is considered served only if it is assigned to a facility.

The objective function jointly minimizes operational costs and response times while maximizing coverage of high-severity and high-priority regions. The severity (S_j), urgency (U_j), demand (D_j), and priority (p_j) scores computed in the spatial triage layer (Equations 3.7–3.10) are directly incorporated into the objective, enabling triage-aware decision-making. All variables are binary or integer, and all constraints are linear, ensuring compatibility with standard MILP solvers.

The model is implemented using a Python-based optimization framework such as PuLP and solved using a standard MILP solver, named as Coin-or Branch and Cut (CBC) solver [74]. The solver processes facility-level parameters (activation cost, UAV capacity), tile-level inputs (severity, priority, travel time), and global constraints (budget, fleet size, and travel limits). All operational constraints defined in Section 3.4 are enforced, ensuring feasibility under realistic operational conditions.

The optimization outputs include selected facility locations (y_i), UAV allocation per facility (u_i), assignment decisions (x_{ij}), and the set of served demand points (z_j). These outputs define the final UAV deployment plan and are used to evaluate system performance in terms of coverage, cost, and response efficiency. The stage completes the pipeline by translating perception-driven triage information into actionable deployment decisions, forming a unified framework that integrates perception and decision-making for UAV-based emergency response.

3.5.4 Baseline Methods for Comparison

To evaluate the effectiveness of the proposed optimization model, several baseline methods are implemented for comparison.

First, a cost-focused optimization baseline inspired by the ULAP framework is considered.

This model minimizes operational cost without incorporating severity or priority information. This baseline reflects traditional facility location-allocation approaches used in emergency service planning. It serves as a reference for assessing the impact of triage-aware decision-making.

Second, machine learning-based baseline approaches are implemented using RF and SVM models. In these approaches, tiles are ranked based on predicted criticality, and UAVs are assigned to the highest-ranked locations in a greedy manner until available resources are exhausted.

Unlike the proposed MILP formulation, these baseline methods do not explicitly enforce global constraints such as UAV capacity, travel time limits, or budget restrictions. Instead, they provide computationally efficient heuristic allocation strategies based on local ranking. These baseline methods are evaluated alongside the proposed model to compare the performance in terms of cost efficiency, coverage of critical regions, and overall response effectiveness.

3.6 Summary

This chapter presented the methodological foundation of a severity-aware UAV-based EMS framework for disaster response. The proposed system integrates computer vision-based perception, spatial triage modeling, and MILP optimization to transform UAV-acquired aerial imagery into prioritized resource allocation decisions. The system model consists of four core components: (i) physical entities, including UAVs, EMS facilities, affected individuals, and the disaster environment; (ii) perception elements for detecting individuals and estimating their criticality; (iii) spatial modeling for aggregating detections into tile-level demand representations; and (iv) an optimization layer for facility activation and UAV dispatch under operational constraints. The disaster region is partitioned into a 4×4 grid of tiles, each representing a localized demand zone, while UAV operations are subject to constraints such as battery capacity, travel time, and facility limits. The perception pipeline operates in two stages. In the SC1, a binary criticality classifier based on EfficientNet-Lite0 is trained using labeled person-level image crops, where posture labels are mapped into critical and non-critical categories. Classical machine learning baselines, including

RF and SVM, are also trained for comparative evaluation. In the deployment stage SC2, YOLOv8 is used to detect individuals in unseen UAV imagery, global non-maximum suppression removes duplicate detections, and the trained classifier estimates the criticality probability for each detected individual.

The spatial triage module aggregates individual-level detections into tile-level representations. For each tile, statistics such as total victim count, critical victim count (C_j), maximum criticality probability (ρ_j), and average criticality probability ($\bar{p}_{\text{crit},j}$) are computed. These statistics are used to derive composite scores, including severity (S_j), urgency (U_j), demand (D_j), and priority (P_j), which capture the medical relevance and operational importance of each region. The UAV allocation problem is formulated as a MILP model that extends the classical . While ULAP focuses on cost minimization under full-coverage assumptions, it does not incorporate severity or priority information. To address this limitation, a cost-focused baseline with selective dispatch constraints is introduced for comparison. The proposed formulation balances multiple objectives, including minimizing travel time and operational cost, while maximizing coverage of high-severity and high-priority regions. Overall, the proposed framework establishes an end-to-end pipeline that connects perception-driven criticality estimation with optimization-based UAV dispatch, enabling informed and prioritized emergency response under realistic operational constraints.

Chapter 4

Simulation Results

This chapter presents the comprehensive evaluation of the proposed severity-aware uncrewed aerial vehicles (UAV)-based emergency management system (EMS) optimization framework. The evaluation examines both the perception capability of the system, specifically binary criticality estimation and derived severity scoring, and the operational effectiveness of UAV resource allocation decisions under different optimization settings. All simulations follow the two-stage pipeline described in Chapter 3. Stage 1 (SC1) focuses on supervised learning for binary criticality classification and severity estimation, while Stage 2 (SC2) evaluates end-to-end UAV allocation behaviour on unseen real UAV imagery using the optimization model and baseline methods.

4.1 Simulation Setup and Parameters

In Stage SC1, a binary criticality classifier is trained on ground-truth labeled individual-level image. The objective is to classify each detected individual as either critical or non-critical based on posture-derived labels. The model takes cropped person images of size 128×128 as RGB (three-channel) inputs and outputs a binary prediction indicating the likelihood of a critical condition. Multi-class posture labels are mapped into a binary representation, where upright postures are considered non-critical and all other postures are treated as critical.

The training configuration for the SC1 classifier is summarized inTable 4.1.

Table 4.1: SC1 binary criticality classification training configuration (EfficientNet-Lite0).

Parameter	**Value**
Task	Binary classification (critical vs non-critical)
Original classes	5 pose classes (0–4)
Binary mapping	Upright → non-critical; all other poses → critical
Random seed	42
Training split	85% train / 15% validation
Training subset size	8000 samples
Validation subset size	2000 samples
Model backbone	EfficientNet-Lite0
Pretrained initialization	Yes (ImageNet)
Number of output classes	2
Input image size	128×128
Training transforms	Resize, horizontal flip, rotation (10°), color jitter (0.2)
Batch size	64
Number of epochs	10
Optimizer	Adam
Learning rate	1×10^{-3}
Loss function	Cross-entropy loss
Framework	PyTorch, torchvision, timm
Computation device	CUDA (if available), otherwise CPU
Dataset	C2A dataset [68]

In Stage SC2, the trained perception model is applied to unseen UAV imagery. Detected individuals are aggregated into tile-level demand points, from which severity, urgency, demand, and priority scores are computed using Equations 3.9–3.12. These scores serve as inputs to the UAV allocation models.

Unlike SC1, SC2 does not include ground-truth criticality labels and therefore represents the deployment setting of the proposed system. The evaluation in this stage focuses on system-level behavior, specifically UAV allocation decisions under different policies.

The evaluation considers four UAV resource-allocation strategies:

- **Proposed Triage-Aware MILP:** The proposed model is a multi-objective MILP framework that incorporates severity, priority, travel time, and operational cost into UAV allocation decision.

- **Cost-Focused Baseline (ULAP-inspired):** A cost-focused objective adapted from the uncrewed aerial vehicles for emergency management system (UEMS) location-allocation problem (ULAP) formulation [52], which does not incorporate severity or priority information and serves as a non-triage benchmark.

- **Random Forest (RF) Heuristic:** A greedy UAV allocation strategy in which demand locations are ranked based on severity estimates derived from the RF model, and UAVs are assigned sequentially [71].

- **Support Vector Machine (SVM) Heuristic:** A greedy allocation strategy similar to the RF baseline, using severity estimates derived from the SVM model [72].

Performance is evaluated using four metrics across multiple UAV fleet sizes K. Severity coverage and priority coverage, which quantify how well high-risk regions are served, and total travel time and operational cost, which measure efficiency. All metrics are post-normalized to the range $[0, 1]$ to enable consistent comparison range across methods and fleet sizes.

Table 4.2 summarizes the simulation parameters and optimization settings used in SC2 ~~experiments~~. All parameters are kept fixed across simulations to ensure fair and consistent comparison between the proposed model and baseline methods.

The travel-time-related quantities in Table 4.2 including $T_{\max}$ and the UAV speed reference, are defined under the normalized grid-based representation described in Section 3.1 and are interpreted as relative simulation values rather than real-world physical units.

4.2 Binary Criticality Classification Evaluation (SC1)

This section evaluates the performance of the binary criticality classification models that constitute the perception front-end of the proposed UAV-based EMS framework. Since the SC2 deployment dataset does not contain ground-truth criticality labels, all classification metrics are reported on the SC1 dataset, where ground-truth pose annotations are available and mapped to binary criticality labels. The original five C2A pose categories (upright, sitting, bending, lying, upside-down) are

Table 4.2: Simulation and optimization parameters used in SC2

Parameter	Value
Number of demand points/ tiles	696
UAV fleet sizes (K)	$\{2, 4, 6, 8, 10\}$
Maximum travel time ($T_{\max}$)	1200
UAV Speed Reference	15
N_{max}	3
B_{max}	35000
Assignment constraint	$\sum_i x_{ij} = z_j$
Solver	CBC (via PuLP)
Severity weights	$a_1 = 1,\ a_2 = 1,\ a_3 = 1,\ a_4 = 1$
Demand and urgency weights	$\beta_1 = 1,\ \beta_2 = 1,\ \theta_1 = 1,\ \theta_2 = 1$
Priority balance parameter	$\lambda = 0.5$
Balanced objective weights	$w_1 = w_2 = w_3 = w_4 = 1$
Single-objective setting	One weight = 1; others = 0
Metric normalization	Post-normalized to $[0, 1]$

mapped into a binary representation aligned with triage requirements. Specifically, the upright posture is treated as non-critical (class 0), while all other postures are considered (class 1). This mapping reflects a conservative triage strategy aimed at minimizing missed critical cases.

Three models are evaluated: EfficientNet-Lite0 (convolutional neural network (CNN) baseline), RF, and SVM, where the latter two are trained on feature embeddings extracted from the CNN.

Table 4.3 summarizes the classification performance of all models on the SC1 validation dataset. The results show that EfficientNet-Lite0 achieves the strongest overall performance, with the highest accuracy (79.49%) and F1-score (87.98%) among the evaluated models. While the overall accuracy may appear moderate, this is expected given the challenging nature of the task because the model operates on small cropped person regions extracted from aerial imagery, where individuals often occupy only a few pixels. These low-resolution crops limit the amount of discriminative visual information available for classification. It achieves a very high recall for the critical class (99.49%), indicating that nearly all critical individuals are correctly identified. In the context of emergency response, this is the most important criterion, as missing critical victims is far more costly than false positives. Therefore, the observed performance represents a strong and practi-

Table 4.3: SC1 Binary Criticality Classification Performance

Model	Accuracy	Precision (crit)	Recall (crit)	F1 (crit)
EfficientNet-Lite0	79.49%	78.86%	99.49%	87.98%
Random Forest	60.50%	58.14%	75.00%	65.50%
SVM	58.00%	54.65%	94.00%	69.12%

cally effective result for this application. The SVM model also achieves high recall (94.00%), but at the cost of significant lower precision, indicating a tendency to over-classify non-critical cases as critical. The RF model achieves comparatively lower performance across all metrics, indicating limited effectiveness in capturing the visual patterns associated with criticality. From a triage perspective, recall for the critical class is a key performance indicator, as failing to identify critical individuals may lead to severe consequences. In this context, both EfficientNet-Lite0 and SVM exhibit a high-recall behavior. However, EfficientNet-Lite0 provides a more balanced trade-off between recall and precision, resulting in superior overall performance. These results establish the reliability of the perception module for criticality estimation. In particular, the high recall achieved by the CNN model supports its use in downstream severity and priority computation, where underestimation of critical cases is more detrimental than overestimation.

4.3 SC2 Deployment Dataset and Setup

The SC2 deployment stage evaluates the proposed UAV-based EMS framework under realistic disaster response conditions. Unlike SC1, which focuses on supervised training and validation of the classification models, SC2 represents a deployment scenario where the trained perception module is applied to unseen UAV images for decision-making. The SC2 deployment dataset is derived from the C2A UAV-based disaster imagery dataset, which contains aerial images of collapsed building scenarios and flood scenarios with human pose annotations.

In this stage, the UAV images are processed using a YOLOv8-based person detection model to identify human instances in aerial scenes. Each detected individual is extracted as a cropped image of resolution 128×128 and classified into critical and non-critical categories using the trained EfficientNet-Lite0 binary classifier. The classification outputs are subsequently used to compute severity, urgency, demand, and priority scores, which serve as inputs to the optimization model.

A total of 238 UAV images are processed, resulting in 1000 detected person instances. Each UAV image contains multiple individuals, resulting in a larger number of cropped instances than input images. Due to the disaster-specific nature of the dataset, the majority of detected individuals are classified as critical. Specifically, 953 instances (95.3%) are classified as critical, while 47 instances (4.7%) are classified as non-critical. This imbalance is expected in collapsed building disaster scenarios, where individuals are more likely to be in non-upright or high-risk conditions.

To enable spatial decision-making, a 4×4 grid-based tiling strategy is applied to partition each UAV image into non-overlapping regions. Detection outputs are aggregated within each tile to construct spatial demand points, which are then used as input to the optimization model. Due to computational constraints, a representative subset of the SC2 dataset is used for optimization experiments while preserving its statistical characteristics.

Table 4.4: SC2 deployment dataset characteristics.

Property	Value
Number of UAV images	238
Detected person instances	1000
Critical instances	953 (95.3%)
Non-critical instances	47 (4.7%)
Average persons per image	4
Crop resolution	128×128
Tiling strategy	4×4 grid
Scene type	Collapsed building and flood disaster

4.4 Balanced Multi-Objective Optimization

In this simulation, all objective weights are set to one ($w_1 = w_2 = w_3 = w_4 = 1$), resulting in a balanced optimization framework that simultaneously considers severity coverage, priority coverage, travel time, and operational cost. This configuration reflects realistic emergency response conditions, where multiple competing objectives must be addressed jointly.

4.4.1 Performance Under Balanced Objectives

Fig. 4.1 illustrates the performance of the evaluated methods as the UAV fleet size K increases. Since all methods rely on the same CNN-derived demand representation, the observed differences are attributable solely to the optimization strategies. Fig. 4.1(a) shows the severity coverage achieved by each method across increasing UAV fleet sizes K. The proposed CNN-based perception with mixed integer linear programming (MILP)-based optimization model consistently achieves the highest severity coverage, indicating its ability to effectively prioritize high-severity demand points. CNN-based heuristic method achieves moderate performance, while the ULAP-inspired cost baseline results in lower severity coverage. This behavior is expected, as the ULAP formulation does not incorporate severity information into its objective function. Fig. 4.1(b) presents the priority coverage across different fleet sizes. The proposed CNN-based perception with MILP-based optimization model again achieves the highest priority coverage, demonstrating its effectiveness in jointly considering severity and urgency through the priority formulation. The CNN-based heuristic method achieves moderate performance, while the ULAP-inspired cost baseline exhibits lower priority coverage, as it does not account for priority information during allocation.Fig. 4.1(c) illustrates the total travel time for each method. The CNN + heuristic approach achieves the lowest travel time due to its simple assignment strategy, which favors locally efficient decisions. The proposed CNN-based perception with MILP-based optimization model results in slightly higher travel time, reflecting the trade-off required to prioritize high-severity demand points. The ULAP-inspired cost baseline exhibits the highest travel time, indicating less efficient spatial allocation under a purely cost-driven objective. Fig. 4.1(d) shows the operational cost associated with each method. The ULAP-inspired cost baseline and the proposed CNN-based perception with MILP-based optimization model achieve nearly identical cost values across all fleet sizes. This occurs because both methods operate under the same resource and feasibility constraints, resulting in similar UAV and facility usage. The proposed model improves severity and

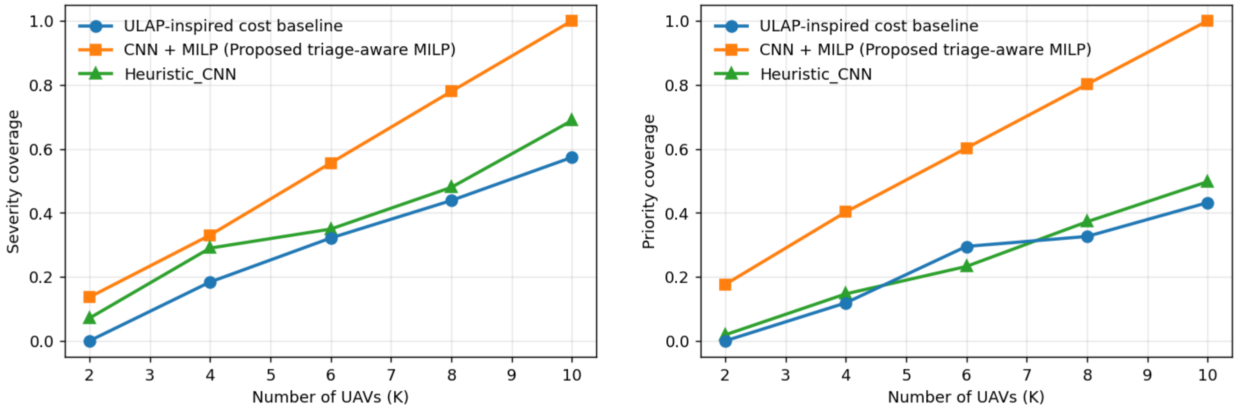

(a) Severity coverage

(b) Priority coverage

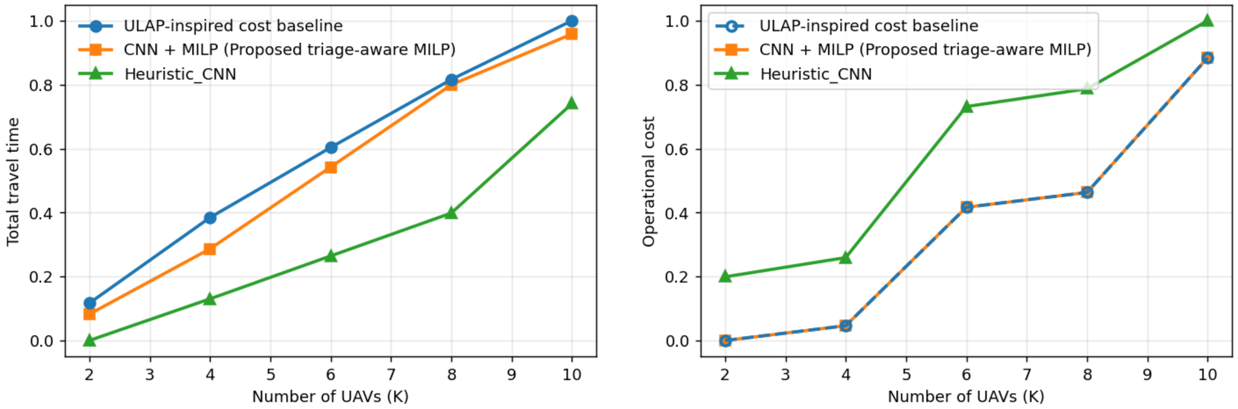

(c) Total travel time

(d) Operational cost

Figure 4.1: Performance under balanced multi-objective optimization (w_1–$w_4 = 1$) across different UAV fleet size K.

priority coverage primarily through more effective assignment decisions rather than increased resource utilization. In contrast, the CNN-based heuristic method results in significantly higher cost due to the lack of coordinated resource allocation.

4.4.2 Evaluation Using the Scalar Objective Function

Fig. 4.2 presents the normalized scalar objective value under the balanced multi-objective configuration. The horizontal axis shows the number of UAVs deployed (K), while the vertical axis shows the normalized scalar objective value, with lower values indicating better performance. The scalar objective aggregates multiple components, including severity coverage, priority coverage,

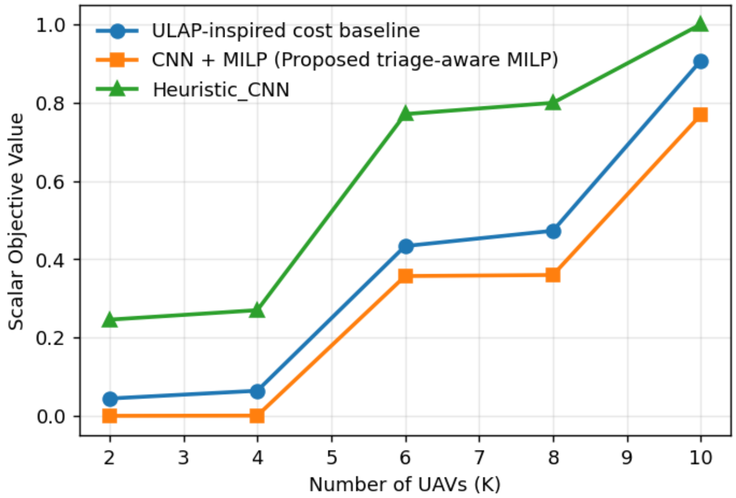

Figure 4.2: Normalized scalar objective value under balanced multi-objective optimization (w_1–$w_5 = 1$). Lower values indicate better performance.

travel time, and operational cost, into a single measure consistent with the optimization formulation. It is important to note that although the compared methods are based on different optimization principles, all models are evaluated using a common scalar objective function to ensure consistent comparison across efficiency and triage-related metrics. The proposed CNN-based perception with MILP-based optimization model achieves the lowest scalar objective values across all fleet sizes, indicating the best performance according to the defined objective. The ULAP-inspired baseline produces higher values because it focuses on cost minimization without incorporating prioritization metrics. The heuristic method results in the highest scalar objective values, reflecting its lack of coordinated multi-objective optimization. These results demonstrate that the proposed model achieves improved performance in jointly optimizing efficiency and triage-driven priorities within the defined optimization framework.

4.4.3 Average Performance Summary

Table 4.5 summarizes the average performance across all UAV fleet sizes under balanced objective weights. The proposed CNN-based perception with MILP-based optimization model achieves the highest severity and priority coverage, indicating that the triage-aware optimization effectively

Table 4.5: Average performance across all UAV fleet sizes (K) under balanced weights. Higher values indicate better performance for severity and priority coverage, while lower values indicate better performance for travel time and operational cost.

Model	SevCov ↑	PriCov ↑	Travel ↓	Cost ↓
ULAP-Inspired (Cost-Only Baseline)	0.007999	0.008179	25.821349	**15820.0**
Proposed (CNN + MILP)	**0.013705**	**0.017448**	24.003423	**15820.0**
CNN + Heuristic	0.009602	0.008669	**15.793329**	20860.0

prioritizes high-severity demand regions. The CNN-based heuristic method achieves lower coverage values compared to the proposed model due to its lack of global coordination. In terms of travel time, the CNN-based heuristic method achieves the lowest values due to its localized assignment strategy. In contrast, the proposed model exhibits higher travel time as it prioritizes more critical demand points. Notably, the ULAP-inspired cost-only baseline and the proposed model yield identical average operational costs, indicating that both approaches use comparable resources. However, the proposed model achieves significantly higher severity and priority coverage without increasing cost, demonstrating improved performance in severity-aware allocation.

4.4.4 Overall Comparison Using Scalar Objective

Table 4.6 summarizes the average scalar objective values across all UAV fleet sizes. Lower values indicate better overall multi-objective performance. The proposed CNN-based perception with MILP-based optimization model achieves the lowest scalar objective value among the evaluated methods, with an average value of 0.1136. The ULAP-inspired cost baseline achieves the second-lowest value of 0.1289, while the CNN-based heuristic method obtains the highest scalar objective value of 0.1698. These results indicate that the proposed model achieves the best overall performance across severity coverage, priority coverage, travel time, and operational cost, while maintaining a balance between efficiency and triage-driven priorities. Although the ULAP-inspired baseline maintains low operational cost, it achieves lower severity and priority coverage because these factors are not included in its objective function. The CNN-based heuristic method achieves lower travel time; however, its higher cost and lower coverage reduce its overall scalar

Table 4.6: Average scalar objective across all UAV fleet sizes (K) under balanced weights. Lower values indicate better overall multi-objective performance.

Model	Scalar Objective (↓)
Proposed (CNN + MILP)	**0.113646**
ULAP-Inspired (Cost-Only Baseline)	0.128947
CNN + Heuristic	0.169816

performance. Overall, the scalar objective results show that the proposed CNN-based perception with MILP-based optimization model utilizes available UAV resources more effectively, improving severity and priority coverage while maintaining comparable operational cost.

4.5 Single-Objective SC2 Simulation

To analyze the influence of individual decision objectives, a single-objective severity-based simulation is performed. In this setup, the severity weight is set to 1 while all other objective weights are set to 0. This configuration isolates the effect of severity prioritization on UAV allocation behavior. The performance is evaluated using four metrics, including severity coverage, priority coverage, total travel time, and operational cost. The comparison includes the proposed triage-aware MILP model, a cost-focused baseline (ULAP-inspired), and two heuristic approaches based on RF and SVM predictions. All results are reported across increasing UAV fleet sizes K. Metrics are normalized to the range $[0, 1]$ for consistent comparison.

4.5.1 Severity-Only Optimization

In this simulation, only the severity objective is activated (w_3 = 1) , forcing the allocation strategy to prioritize demand points with the highest severity scores. Fig. 4.3 illustrates the performance of all methods across increasing UAV fleet size K. As shown in 4.3(a), the proposed CNN-based perception with MILP-based optimization model consistently achieves the highest severity coverage across all fleet sizes. The RF and SVM heuristics achieve very similar coverage values, closely matching the proposed model. This behavior is consistent with the severity-based objective, as all three approaches prioritize demand points based on severity scores. However, the RF and SVM methods rely on greedy selection, assigning UAVs sequentially without considering global feasibility constraints. In contrast, the proposed CNN-based perception with MILP-based optimization model jointly optimizes all assignments, ensuring that the selected set of demand points is globally feasible under the given constraints. The small difference between the methods indicates that severity-based prioritization is effective, while the MILP provides a principled and reliable framework for handling constraints and guaranteeing feasible allocation. The ULAP-inspired cost baseline achieves significantly lower severity coverage, as it does not incorporate severity information into its objective function. A similar trend is observed in Fig. 4.3(b), where the proposed model also achieves the highest priority coverage, despite priority not being explicitly optimized. This occurs because priority is partially derived from severity, leading to indirect improvement. The RF and SVM heuristics again perform closely but slightly below the proposed model, while the ULAP baseline remains the weakest.

Fig. 4.3(c) shows that the RF and SVM heuristic methods achieve the lowest travel time, as their greedy assignment tends to select nearby demand points. The proposed CNN-based perception with MILP-based optimization model results in moderately higher travel time, reflecting the trade-off required to prioritize high-severity demand points that may be spatially dispersed. The ULAP baseline exhibits the highest travel time, indicating less efficient spatial allocation under cost-only optimization. Fig. 4.3(d) shows the operational cost associated with each method under the severity-only setting. The ULAP-inspired cost baseline achieves the lowest cost across all fleet sizes, as expected, because its objective is cost minimization and does not include severity information. The proposed CNN-based perception with MILP-based optimization model incurs a higher cost than the ULAP-inspired baseline because it prioritizes high-severity demand points rather than selecting only the lowest-cost assignments. The RF and SVM heuristics exhibit the highest operational cost, especially as K increases. This occurs because the severity-ranked greedy heuristics select high-severity demand points sequentially without globally coordinating facility activation and UAV usage. Overall, the cost trend highlights the trade-off between severity prioritization and

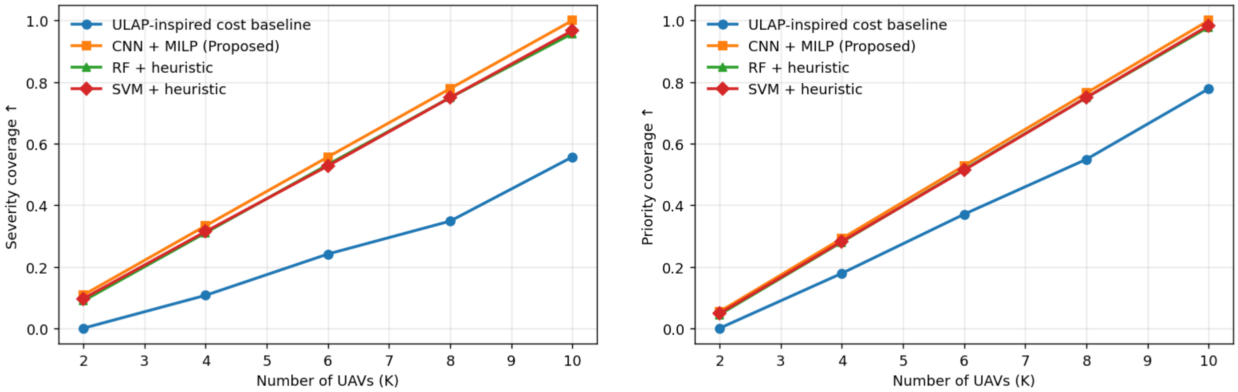

(a) Severity coverage

(b) Priority coverage

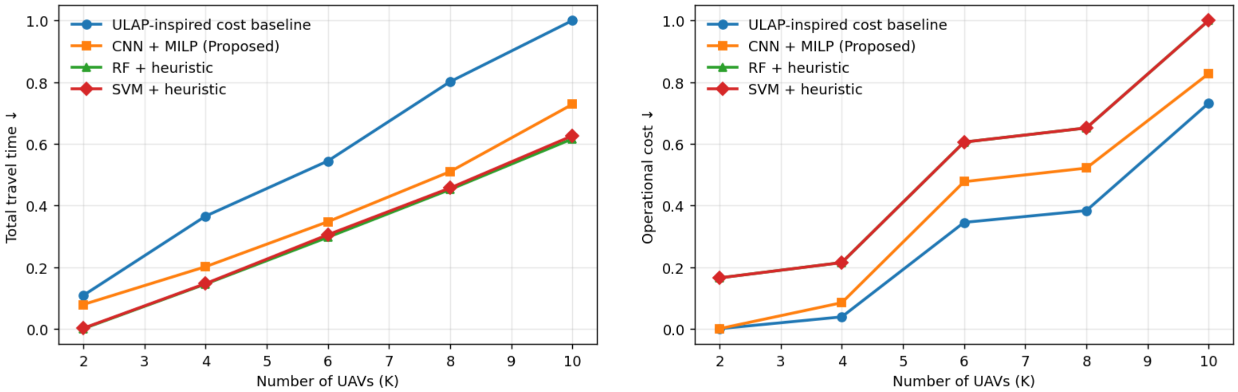

(c) Total travel time

(d) Operational cost

Figure 4.3: Performance under severity-only optimization (w_3=1) across different UAV fleet sizes K.

operational expense. The operational cost in this context consists of the fixed facility activation cost and the UAV deployment cost, as defined in the objective function. These results validate the behavior of the proposed model under severity-only optimization, while also highlighting the associated trade-offs in other performance metrics.

Top-Severity Victim Analysis: To complement the aggregate performance metrics, a victim-level analysis is conducted under the severity-only setting with K=10. This analysis examines which high-severity demand points are selected by each method. Table 4.7 presents the top 20 demand points ranked in descending order based on the CNN-derived severity scores. All models operate on an identical set of demand points with consistent spatial coordinates and ordering, ensuring a fair and direct comparison of allocation decisions across different severity estimation

Table 4.7: Top-20 demand points ranked by CNN-derived severity score and their selection by different methods under severity-only optimization (w_3=1). A: ULAP-inspired cost baseline, B: CNN-based perception with MILP-based optimization (Proposed), C: RF heuristic, and D: SVM heuristic. A checkmark (✓) indicates selection, while (×) indicates non-selection.

Rank	**Demand ID**	**Severity Score**	**A**	**B**	**C**	**D**
1	308	0.921050	×	✓	×	×
2	497	0.910734	×	✓	×	×
3	441	0.905975	×	✓	×	✓
4	560	0.897476	×	✓	✓	✓
5	556	0.897287	×	✓	×	×
6	581	0.892560	×	✓	×	✓
7	523	0.892071	×	✓	×	×
8	137	0.888438	×	✓	✓	×
9	239	0.888321	×	✓	×	×
10	388	0.886078	×	✓	✓	✓
11	203	0.879412	×	×	×	×
12	433	0.878804	×	×	✓	×
13	195	0.878566	×	×	×	×
14	428	0.876574	×	×	×	✓
15	380	0.876537	×	×	×	×
16	98	0.875780	×	×	✓	×
17	543	0.871955	×	×	×	×
18	70	0.871863	×	×	×	×
19	165	0.870835	×	×	×	✓
20	667	0.870781	×	×	×	×

methods. For each demand point, a check mark (✓) indicates selection by the corresponding method, while a cross (×) indicates non-selection. The proposed CNN-based perception with MILP-based optimization model selects nearly all of the highest-severity demand points within the top-K selection, which is consistent with maximizing severity. This indicates that the model effectively prioritizes critical regions while ensuring feasibility under system constraints. In contrast, the RF and SVM heuristic methods achieve similar aggregate severity coverage, as shown in Fig. 4.3(a). However, their individual selections differ from the CNN-based ranking presented in Table 4.7. In particular, these methods do not consistently select the highest-ranked demand points and may instead include lower-ranked points. This behavior arises because their severity estimates are derived from different classifiers, leading to variations in ranking. This highlights that although RF and SVM heuristics can achieve similar aggregate severity coverage, they do not consistently select the highest-severity demand points, resulting in less reliable prioritization.

The ULAP-inspired cost baseline fails to select high-severity demand points, as it does not incorporate severity information into its objective function. This highlights the limitations of cost-driven allocation strategies in scenarios that require prioritizing critical demand points. Overall, the results demonstrate that while multiple methods can achieve similar aggregate performance under severity-only optimization, the proposed CNN-based perception with MILP-based optimization framework more reliably aligns with the highest-severity demand points, providing a principled and constraint-aware prioritization mechanism. This analysis provides a finer-grained validation of allocation decisions, demonstrating that the proposed model not only achieves high aggregate performance but also correctly prioritizes the most critical demand points.

4.6 Computational Performance Analysis

The computational performance of the proposed mixed-integer linear programming (MILP) model and the ULAP cost-based baseline are assessed in terms of solver runtime and optimality gap across various UAV fleet sizes. All models are implemented in Python using the PuLP framework and solved using the CBC solver. The runtime represents the total time required by the solver to obtain a solution. At the same time, the optimality gap measures the relative difference between the best feasible solution and the solver's bound. A zero optimality gap indicates that the global optimal solution has been found. Table 4.8 presents the runtime and optimality gap for both models over varying fleet sizes.

The results in Table 4.8 show that both models achieve optimal solutions for all tested instances, with zero optimality gap (0%), indicating that the global optimum is reached in every case. This behavior is attributed to the moderate problem size and the structured linear formulation, which allow the solver to explore the solution space and prove optimality efficiently. This indicates that the problem instances considered in this study are computationally tractable for exact optimization using standard MILP solvers. The runtime remains low for both models across all UAV fleet sizes.

Table 4.8: Solver runtime and optimality gap for ULAP and the proposed triage-aware MILP model.

Model	K	Status	Runtime (s)	Gap (%)
ULAP Cost Baseline K	2	Optimal	2.565759	0.0
Proposed Triage	2	Optimal	7.195815	0.0
ULAP Cost Baseline K	4	Optimal	0.593082	0.0
Proposed Triage	4	Optimal	5.022970	0.0
ULAP Cost Baseline K	6	Optimal	2.512523	0.0
Proposed Triage	6	Optimal	7.034943	0.0
ULAP Cost Baseline K	8	Optimal	0.639016	0.0
Proposed Triage	8	Optimal	6.258020	0.0
ULAP Cost Baseline K	10	Optimal	1.303173	0.0
Proposed Triage	10	Optimal	8.007467	0.0

The ULAP cost-based baseline is solved within approximately 0.5–2.5 seconds, while the proposed triage-aware model requires approximately 5–8 seconds. The increased runtime of the proposed model is expected due to its multi-objective formulation and additional constraints that incorporate severity and priority into the decision-making process. Despite this increase, the computation time remains within a few seconds, demonstrating that the proposed model is computationally efficient and suitable for practical UAV-based emergency response scenarios. Overall, these results confirm that the proposed model achieves a favorable balance between solution quality and computational efficiency. In summary, the proposed triage-aware MILP model achieves globally optimal solutions with negligible computational overhead for all tested problem sizes. While the inclusion of additional decision variables and constraints increases the runtime compared to the ULAP baseline, the overall computation time remains within practical limits. This demonstrates that the proposed model not only improves decision quality but also maintains computational tractability, making it suitable for real-time or near-real-time UAV-based emergency response applications.

4.7 Summary

This chapter presented a comprehensive evaluation of the proposed severity-aware UAV-based EMS framework across both perception (SC1) and deployment (SC2) stages. In SC1, CNN, RF, and SVM classifiers were evaluated for binary criticality estimation. The CNN model demonstrated the most reliable and consistent performance, particularly in identifying critical cases, thereby justifying its use for generating severity scores in subsequent experiments. While RF and SVM achieved reasonable results, their comparatively lower accuracy limited their suitability for downstream decision-making. In SC2, the analysis examined allocation strategies under both single-objective and multi-objective settings. In the severity-focused scenario, the proposed CNN-based perception with MILP-based optimization model achieved the highest severity coverage, benefiting from its ability to produce globally optimal and constraint-aware assignments. Heuristic approaches based on RF and SVM yielded competitive but slightly inferior results due to their reliance on greedy selection, which does not consider global feasibility. The ULAP-inspired baseline performed poorly in this setting, as it does not incorporate severity into its objective. The balanced multi-objective evaluation, which jointly considered severity coverage, priority coverage, travel time, and operational cost, provided a more realistic assessment. Using a common CNN-derived demand representation, the comparison isolated the impact of allocation strategies. The proposed CNN-based perception with MILP-based optimization model achieved the most effective trade-off, maintaining high coverage while keeping travel time and cost under control. In contrast, the heuristic approach favored lower travel time at the expense of coordinated decision-making, and the ULAP baseline remained cost-efficient but failed to prioritize critical demand points. Finally, the Top-20 victim analysis offered fine-grained validation at the decision level. The proposed CNN-based perception with MILP-based optimization model consistently selected the highest-priority demand points, whereas heuristic methods captured only a subset due to their greedy nature, and the ULAP baseline did not reflect severity-driven prioritization. Overall, the results confirm that integrating CNN-based perception with MILP-based optimization enables effective prioritization of critical demand while ensuring feasible and efficient resource allocation, making the framework well-suited for UAV-assisted emergency response.

Chapter 5

Conclusion and Future Work

5.1 Conclusions

This thesis explored the integration of artificial intelligence (AI)-driven perception and optimization for uncrewed aerial vehicles (UAV)-based emergency medical service (EMS) systems, with an emphasis on severity-aware decision-making. The findings demonstrate that coupling perception outputs with optimization models enables more informed, context-aware allocation compared to traditional cost-driven approaches. At the perception level (SC1), convolutional neural network (CNN)-based models were shown to provide the most reliable and consistent criticality estimation, outperforming classical methods such as random forest (RF) and support vector machine (SVM). Although RF and SVM offered reasonable approximations, their lower accuracy and variability limit their effectiveness for downstream decision-making, supporting the use of CNN-derived severity scores as the primary input to the allocation stage. At the system level (SC2), the balanced multi-objective evaluation incorporates severity coverage, priority coverage, travel time, operational cost, and offered a realistic assessment of deployment performance. The proposed CNN-based perception with multi integer linear programming (MILP)-based optimization framework achieved the most effective trade-off across these competing objectives, maintaining high coverage while controlling resource usage. In contrast, the CNN-based heuristic approach reduced travel time through localized decisions but incurred higher operational cost and lacked global coordination. The uncrewed aerial vehicles for emergency management system (UEMS) location-allocation problem (ULAP)-inspired baseline remained cost-efficient but consistently failed to prioritize high-severity demand, underscoring the limitations of cost-centric strategies in emergency response contexts. Further analysis under a severity-only objective highlighted that while both MILP and heuristic approaches can achieve comparable coverage when prioritization is explicit, their underlying behaviors differ significantly. Heuristic methods rely on greedy selection and do not guarantee feasibility under complex constraints, potentially leading to suboptimal or inconsistent allocations. The MILP formulation, however, ensures globally coordinated and constraint-compliant decisions. This distinction was further reinforced by the Top-20 victim analysis, where the CNN-based perception with MILP-based optimization model consistently selected the highest-priority demand points. In contrast, heuristic methods captured only a subset, and the ULAP baseline failed to reflect severity-driven prioritization. Overall, the results confirm that incorporating severity-aware objectives into UAV allocation leads to more meaningful and reliable prioritization of critical demand points without compromising operational feasibility. While improvements in aggregate metrics may appear modest under certain conditions, the decision-level analysis reveals clear advantages in consistency and robustness. This work demonstrates that integrating perception and optimization is both practical and effective for enhancing UAV-based EMS systems and highlights the importance of balancing prioritization with efficiency in real-world deployments.

5.2 Challenges and Future Research Directions

The development of a unified, real-time, severity-aware UAV-based framework introduces several interdisciplinary challenges that must be addressed before deployment at an operational scale. While the proposed system demonstrates the feasibility of integrating UAV perception, AI-driven severity modelling, and multi-objective optimization, several open research problems remain. These challenges are consistent with limitations identified in prior work on UAV-based disaster response, communication systems, perception, and optimization [7,32,33,56,66,67].

Perception robustness and dataset limitations: A key limitation of the proposed framework is its reliance on vision-based perception and the representativeness of the training dataset. The detection and posture classification models may degrade under adverse conditions such as occlusion, smoke, debris, poor illumination, or extreme viewing angles, as also noted in prior perception studies [32,35–37]. Since severity estimation is directly derived from posture driven classification, such perception errors propagate to downstream decision-making. Future research should focus on improving robustness through domain adaptation, uncertainty-aware learning, and temporal consistency across video frames. Additionally, integrating multi-modal sensing such as thermal imaging, LiDAR, and acoustic signals can significantly enhance reliability under degraded environmental conditions.

Deterministic severity modeling and lack of uncertainty representation: The current severity and priority scoring mechanisms are deterministic, relying on fixed weighting schemes and static mappings from posture to severity. This formulation does not capture uncertainty in perception outputs (e.g., low-confidence classifications) or variability in real-world victim conditions. Future work should incorporate probabilistic and uncertainty-aware severity modeling, such as Bayesian neural networks or Monte Carlo dropout techniques, enabling uncertainty propagation from perception to optimization. This would allow the system to make risk-aware decisions under incomplete or noisy information.

Static optimization and limited real-time adaptation: The proposed MILP-based optimization framework assumes static demand, deterministic travel times, and stable communication conditions. However, disaster environments are inherently dynamic, with evolving victim conditions, environmental hazards, and UAV states. Future research should extend the framework toward dynamic and stochastic optimization, including receding-horizon control, model predictive control, and reinforcement learning-based dispatch strategies. These approaches would enable continuous adaptation to real-time changes in both perception inputs and environmental conditions.

Limited clinical fidelity of severity modeling: While posture-based cues provide a useful proxy for victim condition, they do not fully capture the complexity of medical triage, which involves physiological indicators such as respiration, bleeding, and consciousness. This limits the clinical accuracy of severity estimation. Future systems should incorporate richer clinical information through multi-modal sensing and closer alignment with established medical triage protocols. Collaboration with healthcare professionals and integration of medically validated scoring systems will be essential to improve real-world applicability.

Scalability and computational constraints: Another limitation arises from the computational complexity of the MILP formulation. As the number of UAVs and demand points increases, the optimization problem grows significantly in size, potentially limiting real-time applicability in large-scale disaster scenarios. Future directions include developing scalable solver architectures, such as distributed optimization, hierarchical decomposition, and graphics processing unit (GPU)-accelerated computation. These approaches would support deployment in city-scale environments involving large UAV fleets and dense victim distributions.

Lack of temporal victim modeling: The current framework provides a snapshot-based assessment of severity and does not model the evolution of the victim's condition over time. In real disaster scenarios, victim states may change dynamically, influencing prioritization decisions. Future work should incorporate temporal modeling of victim states, such as survival probability functions, deterioration curves, or time-dependent urgency metrics. This would enable predictive prioritization rather than static scoring.

Communication constraints and information prioritization: The framework assumes reliable communication between UAVs and the optimization layer, which may not hold in post-disaster environments characterized by intermittent connectivity and limited bandwidth [56,58,63]. Delays or data loss can lead to outdated severity information and suboptimal dispatch decisions. Future research should explore communication-aware optimization and semantic data prioritization, where high-severity or high-uncertainty information is transmitted with priority. This aligns with emerging semantic communication paradigms that focus on transmitting task-relevant information rather than raw data.

Limited integration with multi-agent and human-centric systems: The proposed framework focuses primarily on UAV-based decision-making and does not fully integrate with broader EMS ecosystems, including ambulances, ground responders, and hospital networks. Additionally, it lacks mechanisms for human-in-the-loop interaction. Future work should extend the framework to multi-agent coordination, incorporating UAVs, ground vehicles, and human responders within a unified decision-making architecture. Human-in-the-loop systems, including explainable severity scoring and override mechanisms, will be essential for ensuring trust, safety, and operational acceptance.

Addressing these limitations will be essential for transitioning severity-aware UAV EMS systems from proof-of-concept implementations to real-world deployment. The framework proposed in this thesis provides a foundational step by demonstrating how perception, severity modeling, and optimization can be tightly integrated into a coherent decision-making pipeline. Building on this foundation, future research can advance scalable, robust, and clinically aligned UAV-assisted emergency response systems capable of operating effectively in complex, uncertain disaster environments.

Bibliography

[1] Z. Shah, U. Javed, M. Naeem, S. Zeadally, and W. Ejaz, "Mobile edge computing (MEC)-enabled UAV placement and computation efficiency maximization in disaster scenario," *IEEE Transactions on Vehicular Technology*, vol. PP, pp. 1–12, Oct. 2023.

[2] R. K. M, P. Chaudhary, M. R. Reddy, G. Madhumathi, S. Gangwar, and D. R. S. S, "Real-time flood monitoring and resource allocation using AI-driven UAVs in disaster-prone areas," in *2nd International Conference on Artificial Intelligence and Knowledge Discovery in Concurrent Engineering (ICECONF)*, pp. 1–8, Chennai, India, Oct. 2025.

[3] R. Masroor, M. Naeem, and W. Ejaz, "Efficient deployment of UAVs for disaster management: A multi-criterion optimization approach," *Computer Communications*, vol. 177, pp. 185–194, Sep. 2021.

[4] M. A. Qadir, S. Sharif, S. Zeadally, and W. Ejaz, "Digital twin-enabled unmanned aerial vehicles network for disaster management," *IEEE Internet Computing*, vol. 29, no. 3, pp. 48–55, Jun. 2025.

[5] W. Ejaz, A. Ahmed, A. Mushtaq, and M. Ibnkahla, "Energy-efficient task scheduling and physiological assessment in disaster management using UAV-assisted networks," *Computer Communications*, vol. 155, pp. 150–157, Apr. 2020.

[6] C. Qu, F. B. Sorbelli, R. Singh, P. Calyam, and S. K. Das, "Environmentally-aware and energy-efficient multi-drone coordination and networking for disaster response," *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1093–1109, Jun. 2023.

[7] M. Erdelj and E. Natalizio, “UAV-assisted disaster management: Applications and open issues,” in *International Conference on Computing, Networking and Communications (ICNC)*, pp. 1–5, 2016.

[8] T. Calamoneri, F. Corò, and S. Mancini, “A realistic model to support rescue operations after an earthquake via UAVs,” *IEEE Access*, vol. 10, pp. 6109–6125, 2022.

[9] D. Vikram, B. D. Devi, H. P, S. A. M, S. Chandra, and S. P. Santhoshkumar, “Designing intelligent drones and robots for medical and rescue operations in Indian terrains,” in *International Conference on Future Technologies (ICFT)*, pp. 1–6, Sangli, India, Nov. 2025.

[10] R. G. Ribeiro, L. P. Cota, T. A. M. Euzébio, J. A. Ramírez, and F. G. Guimarães, “Unmanned-aerial-vehicle routing problem with mobile charging stations for assisting search and rescue missions in postdisaster scenarios,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 11, pp. 6682–6696, 2022.

[11] Z. Yao, W. Cheng, W. Zhang, and H. Zhang, “Resource allocation for 5G-UAV-based emergency wireless communications,” *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 11, pp. 3395–3410, 2021.

[12] X. Zou, T. Peng, and Y. Zhou, “UAV-Based human detection with visible-thermal fused YOLOv5 network,” *IEEE Transactions on Industrial Informatics*, vol. 20, no. 3, pp. 3814–3823, 2024.

[13] M. Muzammul, M. Assam, and A. Qahmash, “Quantum-inspired multi-scale object detection in UAV imagery: advancing ultra-small object accuracy and efficiency for real-time applications,” *IEEE Access*, vol. 13, pp. 2173–2186, 2025.

[14] R. Khalid, Z. Shah, M. Naeem, A. Ali, A. Al-Fuqaha, and W. Ejaz, “Computational efficiency maximization for UAV-assisted MEC networks with energy harvesting in disaster scenarios,” *IEEE Internet of Things Journal*, vol. 11, no. 5, pp. 9004–9018, 2024.

[15] K. S. Kalyan, M. Rakesh, N. Rohith, and Y. B, “An efficient AI based ensemble model for human detection and rescue during disaster using UAV,” in *International Conference on Smart Technologies for Sustainable Development Goals (ICSTSDG)*, pp. 1–6, 2024.

[16] G. M. Upadhyay, M. Joshi, P. Rathi, P. Vats, M. Narula, and S. K. Gupta, “A comprehensive framework for unmanned aerial vehicle (UAV)-enabled real-time human detection system for disaster management,” in *International Conference on Electrical Electronics and Computing Technologies (ICEECT)*, vol. 1, pp. 1–6, Greater Noida, India, Aug. 2024.

[17] X. Lu, H. Chen, and H. Yang, “HFE-YOLO for small object detection in UAV aerial image,” in *4th International Conference on Artificial Intelligence, Internet of Things and Cloud Computing Technology (AIoTC)*, pp. 44–47, 2025.

[18] F. Gao, H. Xu, Y. Zhang, and L. Cheng, “ACL-YOLO: A novel real-time network for UAV-Based infrared object detection,” in *10th International Conference on Automation, Control and Robotics Engineering (CACRE)*, pp. 150–154, 2025.

[19] J. Chen, B. He, C. D. Singh, C. Fermüller, and Y. Aloimonos, “Active human pose estimation via an autonomous UAV agent,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7801–7808, IEEE, 2024.

[20] X.-W. Tang, Y. Huang, Y. Shi, and Q. Wu, “MUL-VR: Multi-UAV collaborative layered visual perception and transmission for virtual reality,” *IEEE Transactions on Wireless Communications*, vol. 24, no. 4, pp. 2734–2749, 2025.

[21] S. Du, M. Xu, Y. Li, W. Fang, Q. Liu, P. Yi, and Y. Hong, “Multiview landmark-assisted UAV swarm 6-DoF pose estimation using resonant beam and VIO,” *IEEE Internet of Things Journal*, vol. 12, no. 15, pp. 31511–31525, 2025.

[22] C. Li, Y. Wu, S. Zhuge, X. Yang, B. Lin, X. Xu, and X. Zhang, “Robust 6-D pose estimation of the UAV based on hybrid features,” *IEEE Transactions on Instrumentation and Measurement*, vol. 73, pp. 1–13, 2024.

[23] Z. Wang and W. Dong, “Toward aerial collaborative stereo: Real-time cross-camera feature association and relative pose estimation for UAVs,” *IEEE Sensors Journal*, vol. 25, no. 6, pp. 9861–9875, 2025.

[24] Z. Yu, B. Guan, S. Liang, Z. Li, S. Ye, and Q. Yu, “Globally optimal relative pose estimation using affine correspondences with known vertical direction,” *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–12, 2023.

[25] Y. Li, “IMU-aided geographic pose estimation method for UAVs using satellite imageries matching,” *IEEE Robotics and Automation Letters*, vol. 10, no. 3, pp. 2902–2909, 2025.

[26] P. Yang, Y. Fan, J. Zhou, S. Zheng, and L. Han, “Vision-based estimation and tracking on fixed-wing multi-UAV systems,” in *44th Chinese Control Conference (CCC)*, pp. 4366–4371, 2025.

[27] U. Adeel and A. Alnoman, “A multi-GPU edge server-based UAV trajectory optimization for emergencies and disasters,” in *17th International Conference on Computer and Automation Engineering (ICCAE)*, pp. 581–585, 2025.

[28] R. Hejazi, M. Iqbal, A. Ara, and D. Jaber, “An integrated UGV-UAV system for real-time disaster management: A multi-protocol communication framework,” in *8th International Conference on Enterprise Systems (ES)*, pp. 1–6, 2025.

[29] W. Ejaz, M. A. Azam, S. Saadat, F. Iqbal, and A. Hanan, “Unmanned aerial vehicles enabled IoT platform for disaster management,” *Energies*, vol. 12, no. 14, 2019.

[30] M. Choksi, M. A. Zaveri, J. S. Kumar, and S. K. Pandey, “Cloud-based real time data acquisition in IoT environment for post disaster management,” in *9th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pp. 1–6, 2018.

[31] Y. Yan, X. Chen, M. Shi, and R. Li, “A decision support system architecture for intelligent driven unmanned aerial vehicles maritime search and rescue,” in *10th International Symposium on System Security, Safety, and Reliability (ISSSR)*, pp. 424–428, 2024.

[32] A. A. Bany Abdelnabi and G. Rabadi, “Human detection from unmanned aerial vehicles’ images for search and rescue missions: A state-of-the-art review,” *IEEE Access*, vol. 12, pp. 152009–152035, Oct. 2024.

[33] S. Gundu, H. Syed, and J. Harikiran, “Human detection in aerial images using deep learning techniques,” in *2nd International Conference on Artificial Intelligence and Signal Processing (AISP)*, pp. 1–10, 2022.

[34] S. M. Ali Rizvi, R. Mobin Ahmed, K. G. Alamdar, M. Raza Khorasani, and J. A. Memon, “Human detection and localization in indoor disaster environments using UAVs,” in *4th International Conference on Robotics and Computer Vision (ICRCV)*, pp. 159–163, 2022.

[35] Y.-T. Shen, Y. Lee, H. Kwon, D. M. Conover, S. S. Bhattacharyya, N. Vale, J. D. Gray, G. J. Leong, K. Evensen, and F. Skirlo, “”Archangel: A hybrid UAV-based human detection benchmark with position and pose metadata”,” *IEEE Access*, vol. 11, pp. 80958–80972, 2023.

[36] J. Qiu, F. Cai, N. Fu, and Y. Yao, “YOLO-Air: An efficient deep learning network for small object detection in drone-based imagery,” *IEEE Access*, vol. 13, pp. 79718–79735, 2025.

[37] X. Xiong, M. He, T. Li, G. Zheng, W. Xu, X. Fan, and Y. Zhang, “Adaptive feature fusion and improved attention mechanism-based small object detection for UAV target tracking,” *IEEE Internet of Things Journal*, vol. 11, no. 12, pp. 21239–21249, 2024.

[38] A. V. Savkin and H. Huang, “Range-based reactive deployment of autonomous drones for optimal coverage in disaster areas,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 51, no. 7, pp. 4606–4610, 2021.

[39] A. A. Jasim, O. Ata, and O. H. Salman, “AI-driven triage: A graph neural network approach for prehospital emergency triage patients in IoMT-based telemedicine systems,” in *International Symposium on Electronics and Telecommunications (ISETC)*, pp. 1–7, 2024.

[40] M. B. Singh, H. Singh, and A. Pratap, “Criticality-driven scheduling for UAV-assisted remote health monitoring system,” in *IEEE ANTS*, pp. 159–164, 2023.

[41] G. Weaver, A. Manalang, P. Sherman, H. S. Mortveit, and N. Bezzo, “Toward robotic triage: A distributed task and motion planning framework for efficient human-robot emergency response,” in *Systems and Information Engineering Design Symposium (SIEDS)*, pp. 143–148, Greater Noida Gautam Budh Nagar, India, Aug. 2025.

[42] M. Bhende, N. Yamsani, A. Salati, A. Karthikayen, M. Suriya, and V. K. Shukla, “AI-enabled UAVs for rapid response in roadside medical emergencies: A real-time lifeline framework,” in *International Conference on Intelligent and Secure Engineering Solutions (CISES)*, pp. 835–840, 2025.

[43] P. Dong, X. Wang, S. Wang, Y. Wang, Z. Ning, and M. S. Obaidat, “Internet of UAVs based remote health monitoring: An online eHealth system,” *IEEE Wireless Communications*, vol. 28, no. 3, pp. 15–21, 2021.

[44] E. B. Tuncarslan and I. Yilmaz, “Drone assisted remote wellness monitoring using RGB camera,” in *14th International Symposium on Digital Forensics and Security (ISDFS)*, pp. 1–6, 2026.

[45] K. Hussain, P. Saradha, S. Govindasamy, P. Neeraja, G. Suresh, and R. V, “The future role of AI-powered healthcare assistants in medical sector drone monitoring systems,” in *7th International Conference on Intelligent Sustainable Systems (ICISS)*, pp. 618–623, 2025.

[46] L. Diao, Y. Liu, W. Liu, and L. Chiaraviglio, “Unmanned aerial vehicle assisted healthcare resource allocation in disasters,” in *32nd International Telecommunication Networks and Applications Conference (ITNAC)*, pp. 1–5, 2022.

[47] M. R. Sri Ratna Sai Ramya, V. K. J, Y. Moulika, and V. R. Bathina, “Unmanned aerial vehicles to transport healthcare equipment under emergency,” in *OITS International Conference on Information Technology (OCIT)*, pp. 737–742, 2024.

[48] D. R. B. Prasad, C. Madineni, I. T. Kumar, I. Basumatary, K. R. P. Kasyap, and A. Mohan, “Simulation of an autonomous website controlled payload delivery UAV for medical emergency response,” in *6th International Conference On Computing, Communication, Control And Automation (ICCUBEA*, pp. 1–5, 2022.

[49] I. Hussain, T. Sibirajeshwaran, and S. Kishore, “Rapid response drones: UAV-based healthcare delivery for critical situations,” in *2nd International Conference on Advances in Computation, Communication and Information Technology (ICAICCIT)*, pp. 1229–1234, 2024.

[50] D. S. Hewage, W. M. Fernando, and A. Thibbotuwawa, “Advancing UAV implementation for medical supply distribution: A systematic review of current practices and future directions,” in *2025 Moratuwa Engineering Research Conference (MERCon)*, pp. 818–823, 2025.

[51] L. Zhen, K. Wang, and H.-C. Liu, “Disaster relief facility network design in metropolises,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 5, pp. 751–761, 2015.

[52] Y. Park, S. Lee, I. Sung, P. Nielsen, and I. Moon, “Facility location-allocation problem for emergency medical service with unmanned aerial vehicle,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 2, pp. 1465–1479, 2023.

[53] A. Umunnakwe and K. Davis, “An optimization of UAV-based remote monitoring for improving wildfire response in power systems,” *IEEE Open Access Journal of Power and Energy*, vol. 10, pp. 678–688, 2023.

[54] M. Matracia, M. A. Kishk, and M.-S. Alouini, “UAV-aided post-disaster cellular networks: A novel stochastic geometry approach,” *IEEE Transactions on Vehicular Technology*, vol. 72, no. 7, pp. 9406–9418, 2023.

[55] P. S. Bithas, G. P. Efthymoglou, A. G. Kanatas, and K. Maliatsos, “UAV selection in aerial integrated sensing and communication networks,” in *IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp. 1–5, 2024.

[56] X. Tang, F. Chen, F. Wang, and Z. Jia, “Disaster-resilient emergency communication with intelligent air–ground cooperation,” *IEEE Internet of Things Journal*, vol. 11, no. 3, pp. 5331–5346, 2024.

[57] S. Barhoumi, M. Erdem, and M.-S. Alouini, “Terahertz band UAV base stations for post-disaster communication,” *IEEE Open Journal of the Communications Society*, vol. 6, pp. 4788–4803, 2025.

[58] D. R. Bolla, A. Jogur, M. S. Gagan, V. Sastry, G. B. Arjun Kumar, and H. Sanvik, “5G-enabled disaster management: enhancing emergency response and recovery,” in *International Conference on Advances in Modern Age Technologies for Health and Engineering Science (AMATHE)*, pp. 1–5, 2025.

[59] H. C. de Resende, J. Francisco Nunes Pinheiro, P. Reiter, C. B. Both, and J. M. Marquez-Barja, “4G/5G performance of a multi-RAT UAV for medical parcel delivery,” in *IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*, pp. 268–271, 2022.

[60] H. Shakhatreh, A. Khreishah, and B. Ji, “UAVs to the rescue: Prolonging the lifetime of wireless devices under disaster situations,” *IEEE Transactions on Green Communications and Networking*, vol. 3, no. 4, pp. 942–954, 2019.

[61] J. Yu and F. Ye, “User equipment localization and victim estimation with next-generation PSC in emergency response,” in *IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, 2019.

[62] J. Li, Q. He, X. Wang, A. Hawbani, K. Yu, Y. Bi, and L. Zhao, “UAV-assisted microservice mobile edge computing architecture: Addressing post-disaster emergency medical rescue,” *IEEE Transactions on Computers*, vol. 74, no. 8, pp. 2635–2648, 2025.

[63] J. Xu, K. Ota, and M. Dong, “Ideas in the air: Unmanned aerial semantic communication for post-disaster scenarios,” *IEEE Wireless Communications Letters*, vol. 14, no. 6, pp. 1598–1602, 2025.

[64] R. Mao, B. Du, D. Sun, and N. Kong, “Optimizing a UAV-based emergency medical service network for trauma injury patients,” in *Systems and Information Engineering Design Symposium (SIEDS)*, pp. 721–726, 2021.

[65] F. Wang, H. Wang, Z. Qin, and J. Tang, “UAV target detection algorithm based on improved YOLOv8,” *IEEE Access*, vol. 11, pp. 116534–116544, 2023.

[66] J. Chen, P. Wan, and G. Xu, “Cooperative learning-based joint UAV and human courier scheduling for emergency medical delivery service,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 1, pp. 935–949, 2025.

[67] M. Ashraf, A. Gaydamaka, B. Tan, D. Moltchanov, and Y. Koucheryavy, “Low complexity algorithms for mission completion time minimization in UAV-based emergency response,” *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 12, pp. 7945–7960, 2024.

[68] R. A. Nihal, B. Yen, K. Itoyama, and K. Nakadai, “UAV-enhanced combination to application: Comprehensive analysis and benchmarking of a human detection dataset for disaster scenarios,” 2024.

[69] N. Narayan, A. R. Devidas, and R. Prabhakaran, “Deep learning based human detection in flooded environments using multi-UAV systems,” in *5th Asian Conference on Innovation in Technology (ASIANCON)*, pp. 1–7, 2025.

[70] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 6105–6114, 2019.

[71] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[72] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[73] C. Alvarez-Garcia, S. Camara-Anguita, J. M. Lopez-Hens, N. Granero-Moya, M. D. Lopez-Franco, I. M. Comino-Sanz, S. Sanz-Martos, and P. L. Pancorbo-Hidalgo, “Development of the aerial remote triage system using drones in mass casualty scenarios: A survey of international experts,” *PLOS ONE*, vol. 16, no. 5, p. e0242947, 2021.

[74] COIN-OR Foundation, “CBC (Coin-or branch and cut) solver.” `https://github.com/coin-or/Cbc`, 2023. Accessed: 2026-04-17.