Navigating Through Very Low-Level Airspace

Designing Risk-Aware UAV Operations in Dutch Airspace via Infrastructure-Aligned Corridors

MSc Engineering and Policy Analysis

C.F. Martens

Delft University of Technology

Navigating Through Very Low-Level Airspace

Designing Risk-Aware UAV Operations in Dutch Airspace via Infrastructure-Aligned Corridors

by

C.F. Martens

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on Friday July 18, 2025 at 13:00.

Chairperson:	Prof.dr.ir. A. Verbraeck,
First Supervisor:	dr. P.S.A. Stokkink,
Second Supervisor:	Prof.dr.ir. A. Verbraeck,
External Supervisor:	ir. I. Heitlager.
Project Duration:	February, 2025 - July, 2025
Faculty:	Technology, Policy and Management, TU Delft

Cover: reisgraag.nl, *Typische Nederlands landschap met molen en weide*

TU Delft

Preface

This thesis has been written in partial fulfillment of the requirements for the Master's degree in Engineering and Policy Analysis at TU Delft. The research was conducted in collaboration with Schuberg Philis, whose innovative Boxslot Lane concept provided the foundation for this study.

The concept emerged from Lab271 at Schuberg Philis as a long-term research initiative. Their primary concern centered on providing reliable network coverage for mission-critical IT services in drone operations. It was during these conversations that the Boxslot Lane Concept took shape. Rather than pursuing comprehensive area coverage, they proposed a paradigm shift: establish designated corridors with guaranteed full coverage along these defined routes.

Through this thesis, I have translated Schuberg Philis's vision into a model that is both policy-relevant and aligned with strategic thinking in the field, validating it through geospatial data analysis. The framework developed here extends beyond drone applications and offers a new lens through which similar infrastructure challenges can be approached. My hope is that readers will find inspiration to reconsider traditional approaches to complex spatial, technological and regulatory problems. By making both my thesis and code publicly available, I encourage others to extend, examine, adapt, critique, or implement this framework in their own unique ways.

I would like to express my sincere gratitude to Ilja Heitlager for providing this opportunity and for the numerous developmental moments I experienced at Schuberg Philis, both professionally and personally. It has been an enriching and enjoyable journey. I also want to thank the entire Lab271 team for creating such a welcoming and collaborative environment.

I am also deeply grateful to my supervisors, Alexander Verbraeck and Patrick Stokkink, for their excellent guidance throughout this research. Their valuable insights, close engagement with my topic, and well-timed nudges in the right direction enabled me to make important discoveries independently while staying on track. Their enthusiasm for the subject matter was particularly motivating, and I greatly appreciated the approachable and open communication throughout our collaboration.

Special thanks go to Ids for digitizing my less-than-perfect drawings, and to my friends and family for their unwavering support throughout this journey.

C.F. Martens
Delft, July 2025

Executive Summary

The rapid growth of e-commerce over the past decade has significantly increased freight transport volumes, intensifying economic, social, and environmental challenges such as urban congestion, air and noise pollution, delivery delays, excessive energy consumption, and reduced quality of life [28, 66, 65]. These impacts further contribute to greenhouse gas emissions, depletion of non-renewable resources, and increased waste generation. Unmanned Aerial Vehicles (UAVs), or drones, have emerged as a promising alternative, offering reduced operational costs, faster and more precise deliveries, and congestion-independent operations [22, 46]. Studies show UAV logistics can substantially reduce greenhouse gas emissions and energy use compared to traditional freight transport [62, 18, 52], with particular promise in critical sectors such as healthcare for organ transport and medical deliveries to vulnerable populations [33, 20, 54, 55].

The UAV sector has matured into an innovative, internationally oriented industry, and while the Dutch government aims to facilitate UAS integration into civil airspace [37, 39], large-scale UAV delivery deployment remains limited due to complex regulatory frameworks, spatial constraints, and operational challenges that must still be systematically addressed.

The central challenge is one of spatial design and optimisation. UAV corridors must navigate the complex constraints of Dutch airspace while complying with European U-space regulations. In doing so, they must remain economically competitive with ground-based delivery, which requires a careful balance between risk mitigation and efficiency in addition to design requirements. Current approaches fail to integrate the full spectrum of these requirements into a single implementable model. While the existing literature addresses individual aspects such as conceptual airspace structures, risk-based routing or regulatory compliance, these are not combined into an integrated framework suitable for use in highly constrained environments such as the Netherlands, where dense airspace usage and extensive no-fly zones create unique design challenges.

This study addresses this gap by developing and validating a comprehensive framework for risk-aware UAV corridor planning that integrates regulatory requirements, spatial constraints, and operational considerations to support practical implementation within Dutch airspace.

Research Approach

To address these barriers, this study develops a comprehensive framework for designing risk-aware UAV cargo corridors in Dutch airspace, addressing the operational, regulatory and spatial constraints that currently limit implementation. The framework focuses on three critical requirements:

- Competitive feasibility: UAV operations must be cost-effective relative to conventional delivery, achievable through fast delivery requirements, low route density, short-to-medium distances, and fixed-point delivery structures.
- Spatial feasibility: Corridors must navigate extensive no-fly zones, including controlled airspace and environmentally protected areas, while maintaining network connectivity.
- Regulatory compliance: Designs must enable compliance with European U-space regulations, supporting formal risk assessments and mandatory services like geo-awareness and flight authorization.

The research extends beyond existing road-based design approaches by incorporating multiple linear infrastructures and natural corridors as potential flight paths. Railways, waterways, and natural corridors are integrated alongside traditional road networks, providing lower-risk routing alternatives that exploit the Netherlands' unique geographic features.

This multi-infrastructure approach directly addresses the three key requirements. First, it ensures com-

petitive feasibility through simplified operations and faster regulatory approval of fixed routes. Additionally, it enables strategic avoidance of no-fly zones through flexible infrastructure selection. Finally, it facilitates U-space compliance by allowing corridor-level risk assessments rather than flight-by-flight evaluation while supporting efficient deployment of mandatory services along predefined infrastructure paths. The design approach is operationalized through a graph-based network model, where fixed one-way corridors connect predefined origin-destination points using PostNL distribution centers and pickup locations as test cases.

A Semi-Quantitative Risk Assessment (SQRA) methodology evaluates third-party ground risks across five external factors and three consequence domains (fatality risk, property damage, and societal impact). The resulting composite risk scores enable systematic comparison of corridor alternatives while remaining transparent and interpretable for policy makers. The framework employs Dijkstra's algorithm with a weighted objective function controlled by parameter α, enabling flexible prioritization between risk minimization (α = 1) and energy efficiency (α = 0).

Key Findings

The experimental validation demonstrates that fixed UAV corridors can be feasibly designed within the constraints of Dutch airspace while effectively balancing safety and efficiency objectives. The results show that risk exposure can be reduced by 70–80% with only a 15–20% penalty in efficiency at specific parameter settings (α = 0.10–0.20). This confirms that practical implementation of UAV corridors is achievable without prohibitive efficiency losses, while still meeting regulatory, spatial and competitive requirements. Additionally, findings highlight that while corridors are feasible across all tested environments, the optimal safety-efficiency configuration is geographically dependent.

The experiments further revealed that risk-averse routing tends to prioritise waterways, natural corridors, and rural infrastructure to avoid populated areas, while efficiency-focused routing favours direct paths along roads and railways despite their higher associated risks. This differentiation confirms that the framework effectively supports policy-driven corridor selection, enabling authorities to align UAV network design with local safety priorities and operational needs. The successful integration of polygon-based areas through grid structures, as validated in rural environments, demonstrates the framework's capability to handle both linear infrastructure and area-based zones, ensuring comprehensive network coverage even where traditional corridors are sparse.

However, battery constraints emerged as a critical factor limiting rural viability. While urban areas maintained high service coverage even under risk-averse routing, rural deployments faced severe reductions. For example, at a battery capacity of 400 Wh, rural coverage dropped to only 17.4% under risk-averse planning compared to 68.8% in urban settings, underscoring the importance of operational range when selecting appropriate α values for different contexts.

Model sensitivity analyses validate the robustness of the assessment framework. Severity parameters, particularly for heavily used corridors, were found to drive routing decisions more strongly than probability estimates. The framework exhibited expected behaviour: energy parameters had no effect under pure risk optimisation (α = 1.0), while risk parameters had minimal impact under efficiency-focused routing (α = 0.0), demonstrating consistent internal logic.

Overall, the Semi-Quantitative Risk Assessment (SQRA) framework successfully operationalises regulatory requirements while maintaining transparency for stakeholder communication. By evaluating five external risk factors across three consequence domains, the methodology provides a systematic corridor-level assessment aligned with U-space mandates. Importantly, the SQRA approach enables risk classification of potential routes in advance rather than retrospectively, allowing proactive integration of safety considerations into the planning process. The inclusion of the α parameter mechanism enables flexible policy control, allowing authorities to adjust safety-efficiency priorities based on local contexts and strategic objectives.

Limitations and challenges

The main limitations in this study are the simplified height and energy models of the framework, which limit three-dimensional optimisation and realistic energy consumption estimates. In addition, grid-based routing enforces rectilinear paths and distorts routing results. Also, the SQRA risk assessment relies

on structured assumptions rather than empirical data and ignores dynamic hazards and intersection conflicts. Data quality knows uncertainties from OpenStreetMap and PostNL. Finally, operational prerequisites remain theoretical, including mature U-space services, continuous GNSS/5G coverage, and strategic deconfliction capabilities.

Further research should focus on developing refined altitude and energy models that incorporate continuous climb/descent profiles, turning penalties, and environmental factors. Next to this, research should explore advanced routing algorithms beyond grid-based discretisation to enable direct path generation. Additionally, risk parameters should be empirically validated through incident data and expert elicitation. Finally, standardised data collection protocols should be established to support evidence-based model refinement and regulatory acceptance as UAV infrastructure and supporting technologies evolve.

Recommendations

The framework's adaptability enables nuanced risk management by customising parameters to local contexts, allowing differentiated safety standards while maintaining transparency between modelling and policy. Systematic design should combine initial network-wide corridor generation with local refinement to incorporate unmapped sensitive sites and targeted mitigation at unavoidable crossings. A phased deployment strategy, progressing from rural to urban areas with gradually adjusted α values, ensures safety and public trust. Finally, positioning UAV corridors as multi-use platforms supports broader societal applications such as medical logistics, emergency response, and industrial monitoring, maximising infrastructure efficiency while aligning operational standards with diverse needs.

Strategic facility placement should prioritise locations at low-risk corridor intersections rather than highway-adjacent sites to reduce risk without compromising efficiency, complemented by repositioning pickup points and healthcare access sites. U-space services can focus on corridor-specific solutions—such as weather monitoring, optimised 5G, and pre-approved flight authorisations—ensuring cost-effective scalability through phased expansion. Collaborative standardised data collection with operators, insurers, and regulators will validate and refine models for regulatory acceptance. Comprehensive benchmarking against van-based logistics is needed to confirm where this corridor framework adds value. Finally, clear operational guidelines accounting for battery, payload, and communication constraints will enable feasible deployments while maintaining flexibility as technology advances.

Conclusion

This framework provides essential groundwork for integrating UAV freight operations into Dutch airspace. By enabling transparent trade-offs between safety and efficiency, it offers policymakers a practical tool for corridor design that adapts to local contexts while meeting regulatory requirements. The research presents the first comprehensive framework that simultaneously addresses regulatory compliance, spatial feasibility, and operational viability for UAV corridor planning, extending beyond traditional road-focused approaches to integrate diverse infrastructure types including waterways, railways, and natural corridors. The novel corridor-level SQRA methodology effectively bridges theoretical risk assessment and practical routing decisions.

The modular design supports iterative refinement as the UAV ecosystem matures, facilitating the transition from proof-of-concept to operational reality. For policymakers, the framework provides transparent, parametric optimization with clear mechanisms for balancing safety and efficiency based on local priorities, while directly supporting European U-space compliance and Member State implementation guidance. By demonstrating viable corridor designs across diverse geographic contexts, this research proves UAV freight transport can be practically implemented within existing constraints, contributing to sustainable urban development through low-emission delivery alternatives.

Success requires coordinated advancement across regulatory frameworks, technological capabilities, and operational experience. The framework's strength lies in bridging technical optimization with policy implementation, providing a foundation for sustainable transformation of urban logistics through UAV integration. The study's core contribution demonstrates that structured corridor design, combined with systematic risk assessment, can overcome the apparent incompatibility between safety requirements and operational efficiency in UAV freight transport, positioning UAVs as a viable component of future transportation systems.

Contents

List of Figures

List of Tables

List of Abbreviations

Abbreviation	Written in full
4G / 5G	Fourth / Fifth Generation Mobile Network
A*	A-Star Path Planning Algorithm
A1, A2	EU Drone Subcategory Classes under the Open Category Framework
ANSP	Air Navigation Service Provider
BVLOS	Beyond Visual Line of Sight
CBS	Centraal Bureau voor de Statistiek (Statistics Netherlands)
C0–C4	UAV Classes under European Drone Regulation
CONOPS	Concept of Operations
CTA	Control Area
CTR	Control Zone
DGP	Data Generating Process
EASA	European Union Aviation Safety Agency
EDA-FRA*	Enhanced Directed Acyclic–Forward Risk Aware A*
EU	European Union
FL	Flight Level
GLV	Gebied Laag Vliegen (Military Low-Flight Zone)
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
ICAO	International Civil Aviation Organization
IFR	Instrument Flight Rules
IoD	Internet of Drones
ISO	International Organization for Standardization
KM^2	Square Kilometer
LIDAR	Light Detection and Ranging
LUC	Light UAS Operator Certificate
MJ	Megajoule
OSM	OpenStreetMap
OSMnx	Python Library for OSM-based Network Extraction
PLr	Performance Level Required (ISO 13849)
PostNL	Dutch National Parcel and Postal Delivery Company
PSO	Particle Swarm Optimization
RF	Radio Frequency
SDG	Sustainable Development Goal
SESAR	Single European Sky ATM Research
SM	Synchronized Delivery Model
SORA	Specific Operations Risk Assessment
SQRA	Semi-Quantitative Risk Assessment
TMA	Terminal Manoeuvring Area
TPR	Third-Party Risk
UAV	Unmanned Aerial Vehicle
UAS	Unmanned Aerial System
UTA	Upper Control Area
VFR	Visual Flight Rules
VLL	Very Low Level (Airspace)
Wh	Watt-hour

1

Introduction

The global rise in e-commerce, particularly since the COVID-19 pandemic, has significantly increased the volume of goods being transported [28, 66]. This growth has brought considerable economic, social, and environmental challenges [65]. Most prominently, the current reliance on road freight transport contributes to congestion, air pollution, noise pollution and general urban nuisance. From an economic standpoint, such traffic causes delays, unreliable deliveries, and high resource and energy consumption. Socially, it reduces the quality of life in urban areas. Environmentally, it accelerates greenhouse gas emissions, depletes non-renewable resources, and increases product waste.

Unmanned Aerial Vehicles (UAVs), widely known as drones, have emerged as a promising technology to address these issues [22]. As core components of Unmanned Aerial Systems (UAS), UAVs are gaining traction in logistics for their ability to reduce costs, increase delivery speed and precision, and operate independently of ground-based congestion [46]. Several studies have demonstrated the potential of UAV-based logistics to reduce greenhouse gas emissions and overall energy consumption compared to traditional vehicles [62, 18, 52]. Beyond commercial logistics, UAVs have also shown potential in critical applications such as healthcare, including organ transport and medical deliveries to vulnerable populations [33, 20, 54, 55].

Over the past decade, the UAV sector has evolved into an innovative and internationally oriented industry. By 2050, the unmanned aviation sector in the Netherlands is projected to grow from 2,460 to over 11,000 operational drones, with an estimated economic value increasing from €103-176 million to €374-805 million [57]. The Dutch government recognizes the potential of UAS applications and are actively working to facilitate their integration into civil airspace [37, 39]. The technology has already demonstrated practical viability in specialized applications, with companies like ZipLine successfully delivering medical supplies in Africa and Amazon Prime Air conducting commercial trials in selected markets [2, 70]. However, despite these promising demonstrations, widespread implementation of UAV operations faces fundamental challenges.

The integration of UAVs into national airspace systems requires navigating a complex web of spatial, regulatory, and operational constraints that are particularly pronounced in densely populated countries like the Netherlands. The Dutch airspace represents one of Europe's most complex and congested aviation environments. With multiple international airports, extensive military training areas, dense urban development, and numerous protected natural areas, the available airspace for UAV operations is severely constrained. Current regulations prohibit UAV flights over highly populated areas, critical infrastructure, and numerous other sensitive locations, effectively eliminating direct point-to-point delivery for freight transport in most urban contexts.

The European regulatory framework provides essential safety standards through the U-space concept, which defines digital traffic management services to enable safe and scalable UAV operations, particularly in dense urban areas [13, 61]. However, while Member States must comply with mandatory risk assessments and minimum U-space service requirements [11], the actual spatial design of UAV

corridors is left to them. This creates a critical gap: regulations specify safety standards but not how to design networks that are simultaneously safe, operationally efficient, and spatially feasible.

The central challenge is thus one of spatial design and optimisation. UAV corridors must navigate the complex constraints of Dutch airspace while remaining economically competitive with ground-based delivery, requiring a careful balance between risk mitigation and efficiency. Avoiding populated areas increases safety but often results in longer routes with substantial altitude variation, reducing operational viability.

Current approaches fail to integrate the full spectrum of these requirements within a single implementable model. While existing literature addresses individual aspects such as conceptual airspace structures, risk-based routing, or regulatory compliance, it does not combine these into an integrated framework suitable for deployment in highly constrained environments like the Netherlands, where dense airspace usage and extensive no-fly zones create unique design challenges. This study addresses this gap by developing and validating a comprehensive framework for risk-aware UAV corridor planning that integrates regulatory requirements, spatial constraints, and operational considerations to support practical implementation.

Based on this aim, the following research question is formulated:

> **Main Research Question:** Can fixed UAV corridors be feasibly designed within the spatial, operational and regulatory constraints of Dutch airspace, while balancing ground risk and energy efficiency?

To answer this overarching question, six sub-questions are defined:

- **Sub-question 1:** What societal, technical, and regulatory constraints define the design space for UAV operations in the Netherlands?
- **Sub-question 2:** How can UAV corridors be spatially structured within this design space in a manner that is both safe and practically feasible?
- **Sub-question 3:** How can the resulting spatial corridor design be translated into a network routing model suitable for balancing route efficiency and risk exposure?
- **Sub-question 4:** How can third-party ground risk be systematically assessed and scalably integrated into the design of UAV corridors?
- **Sub-question 5:** How can fixed UAV corridors be optimally planned within a spatial network by jointly minimising third-party risk and energy consumption?
- **Sub-question 6:** How does the proposed corridor design model perform in generating feasible UAV routes across different geographic contexts, and what trade-offs emerge between safety and efficiency under varying parameter settings?

This study systematically addresses the gaps by developing a framework that integrates regulatory compliance, spatial feasibility, and operational viability for UAV freight operations in the Netherlands. It establishes a feasible operational space by explicitly incorporating legally designated no-fly zones and structuring corridors along various linear infrastructures and natural corridors, inherently supporting compliance with U-space regulatory requirements. Within this space, a semi-quantitative risk assessment (SQRA) framework is applied to evaluate third-party ground risks across spatial segments. Risk-based path planning is then conducted using a weighted objective function governed by a configurable parameter α, enabling flexible trade-offs between safety and operational efficiency.

Scientific contributions include presenting the first comprehensive framework that simultaneously addresses regulatory compliance, spatial feasibility, and operational viability for UAV corridor planning in highly constrained airspace environments. Although existing studies have explored corridor conceptualisation or risk-based routing individually, none have provided an integrated approach that operationalises these concepts within regulatory and spatial constraints for practical implementation. Furthermore, while most previous research has focused primarily on road networks, this study demonstrates how diverse infrastructure types can be systematically integrated into UAV routing networks. The development of a corridor-level SQRA methodology also provides a novel approach for incorporat-

ing spatially varying ground risks into network design, effectively bridging theoretical risk assessment and practical routing decisions.

Societal contributions include offering policymakers a transparent, parametric optimisation approach with clear mechanisms for balancing safety and efficiency based on local priorities and constraints. Furthermore, the framework directly supports compliance with European U-space requirements while providing practical guidance for Member State implementation. By demonstrating viable corridor designs across diverse geographic contexts, this research shows that UAV freight transport can be practically implemented within existing regulatory and spatial constraints. The framework also contributes to sustainable urban development by enabling low-emission delivery alternatives that reduce traffic congestion and environmental impact.

This study is structured as follows. Chapter 2 provides the theoretical foundation and identifies key design requirements by establishing three main constraints for UAV operations, reviewing existing corridor design principles, and identifying infrastructure-based routing as the primary approach. Chapter 3 translates this theoretical framework into a spatial network model, developing the core model components including graph-based network construction, no-fly zone integration, and fundamental design assumptions. Chapter 4 develops the Semi-Quantitative Risk Assessment (SQRA) methodology, encompassing risk identification, risk analysis, and risk calculation procedures. Chapter 5 integrates both risk and efficiency considerations into the routing algorithm through a weighted objective function. Chapter 6 tests the complete framework across three case studies representing urban, rural, and suburban environments. Chapter 7 synthesizes the findings, evaluates the overall framework performance, and provides actionable guidance for future research, policy implementation, and technical enhancements. Figure 1.1 visualizes the research process and framework structure.

Figure 1.1: Research process and framework structure

2

Literature Synthesis and Design Context for UAV Corridor Planning

The emergence of unmanned aerial vehicles (UAVs) offers promising opportunities for freight transport, but also brings significant complexity in terms of safety, airspace integration and regulatory compliance. This chapter summarizes the main academic and policy literature to identify the societal relevance, technical constraints, and design considerations for UAV corridor planning.

First, the societal and operational potential of UAVs is discussed, with a special focus on their role in urban logistics, healthcare and sustainability goals. This is followed by an analysis of the regulatory and infrastructural constraints that determine UAV deployment, focusing on the unique characteristics of Dutch airspace and U-space implementation. It then explores structuring principles for airspace design. Finally, the chapter introduces the concept of using linear infrastructures, such as roads, canals and railways, as the spatial basis for fixed-route UAV corridors.

Together, these sections lay the foundation for the methodological approach developed in the following chapters and justify key design choices for risk-aware, infrastructure-aligned UAV corridor planning.

2.1. Societal and operational relevance of UAVs

The global rise in e-commerce, particularly since the COVID-19 pandemic has significantly increased the volume of goods being transported [28, 66], bringing numerous economic, social and environmental challenges. Studies indicate that 10 to 20 percent of motorized vehicles in cities are linked to goods transportation or service provision. However, these vehicles account for over 50 percent of the negative effects of road usage, including traffic congestion, pollution, and noise [65]. In addition, freight transport contributes in economic terms to lost time, unreliable deliveries due to congestion, high resource consumption and the costs associated with government regulations and urban planning. Socially, the impacts include health risks from traffic accidents as well as local air and noise pollution. Furthermore, quality of life is diminished through damage to buildings and infrastructure. Finally, the sector compounds polluting emissions, exhaustion of non-renewable resources and more product waste, which has negative environmental impacts [65]. To address these challenges, e-commerce and logistics companies need to turn to innovative solutions. Unmanned Aerial Vehicles (UAVs), widely known as drones, have emerged as a promising technology to tackle these problems [22]. UAVs are key components of UAS (Unmanned Aerial Systems) and are becoming increasingly popular in operations because they reduce costs, facilitate operations and can increase the granularity of delivery [46].

By providing an alternative to traditional delivery methods, UAVs have the potential to reduce road delays [3], and when carefully deployed, UAVs can significantly reduce greenhouse gas emissions and energy consumption in logistics [62, 18, 52]. For example, Stolaroff et al. [62] highlights the potential for carbon reductions in the cargo sector. Similarly, Goodchild and Toy [18] identifies conditions under which UAS outperform trucks in terms of carbon dioxide emissions, particularly in delivery zones close

to depots with fewer stops. Rodrigues et al. [52] further strengthens this evidence by demonstrating that UAS can consume up to 94% less energy per package than traditional modes of transportation.

The economic benefits are equally compelling; according to estimates by Lyon-Hill et al. [33], businesses using drones can generate more revenue, in addition to significant time savings and reduced city congestion. The revenue is estimated to increases for businesses adopting drones, ranging from $25,300 to $284,300 annually within five years, depending on the sector. Drones can also offer significant time savings for consumers, with annual savings of 23–45 hours per person in less densely populated areas. Communities benefit from reduced congestion and fewer traffic accidents, resulting in both financial and societal gains. Operationally, drones alleviate urban traffic congestion, enhance road safety, and improve delivery efficiency, particularly in densely populated areas Lyon-Hill et al. [33].

UAS technology also holds considerable promise for applications in healthcare logistics. The research by Lyon-Hill et al. [33] proves that drones can improve access for vulnerable populations, such as people who are housebound, and improve health outcomes through reliable delivery of essential goods such as medicines. Additionally, Güngören et al. [20] explored the feasibility of drones for medical logistics, focusing on the transport of monoclonal antibodies. Their findings confirmed that the stability of these substances remained unaffected during drone transport. Furthermore, the study by Scalea et al. [55] evaluated the use of drones for organ transportation, demonstrating the suitability of drones for organ delivery over short distances. Also Sage et al. [54] conducted test flights using a drone specifically designed for healthcare transport, in this case donor lungs. The study validated the capability of drones to safely and reliably maintain the stability of donor organs during transit, and highlighted the efficiency of drones in bypassing urban congestion, offering significant potential for urban logistics.

In recent years, the unmanned aviation sector has rapidly matured into an innovative and internationally oriented industry. Opportunities with UAVs are increasingly being exploited, while future innovations are quickly taking shape [37]. This rapid growth is reflected in the development of UAVs in the Netherlands. By 2023, there are already 2,460 drones deployed in sectors such as agriculture, energy and infrastructure, security, delivery and mobility, and the estimated value of unmanned aviation in these sectors is 103 to 176 million euros. These numbers are expected to grow to 11,230 drones and 374 to 805 million euros by 2050 [57]. The Dutch government acknowledges the opportunities presented by Unmanned Aerial Systems (UAS) and aims to facilitate innovative UAS services, such as air-based delivery systems [37, 39]. As noted by SDU UAS Center [56], implementation of UAS can contribute to six EU Sustainable Development Goals, including improving healthcare access (SDG 3), fostering economic growth and quality jobs (SDG 8), and delivering sustainable solutions for infrastructure and transport (SDG 9).

Despite these promising developments, the integration of UAVs into national airspace systems, especially in densely populated regions such as the Netherlands, remains a complex and unresolved challenge. A key aspect of this challenge is to ensure that UAV-based cargo transport can function as a competitive alternative to traditional delivery systems. The following section explores the conditions under which such competitive feasibility can be achieved.

2.2. Strategic Considerations for UAV Delivery Feasibility

When evaluating the market potential for UAV-based freight transport, the benefits should outweigh the costs, and this balance must be more favourable than that of existing alternatives, particularly vans, which currently represent the main competing modality [27]. According to Kennisinstituut voor Mobiliteitsbeleid [27], UAV operations can be a viable alternative when speed is a key requirement and shipment volumes are limited. Similarly, Sesar [58] emphasise that UAVs offer particular added value for premium, same-day deliveries of lightweight packages in urban environments.

In addition to delivery speed, the literature identifies route density, transport distance, and investment costs as critical factors in comparing UAV-based delivery networks to traditional van operations [27, 17].

1. Route density: High route density lowers the relative cost of van-based delivery, as multiple stops can be combined efficiently within a compact area.

2. Distance: Longer distances improve van efficiency through increased route density, whereas drones must perform each delivery individually, making extended distances less cost-effective unless driven by urgency.
3. Investment costs: Setting up and maintaining drone networks requires significant upfront investment, which can only be recouped through high occupancy rates.

To structure the possible applications of UAV-based parcel delivery, Kennisinstituut voor Mobiliteitsbeleid [27] has proposed a categorization into three variants:

- Variant A: parcel delivery from a fixed location to a fixed location
- Variant B: parcel delivery from a fixed location to a variable location
- Variant C: parcel delivery from a variable location to a variable location

Variant A is generally considered easier to implement than Variants B and C.This is primarily because, in Variant A, the aviation authority only needs to authorise specific routes, whereas Variants B and C require permits for entire operational areas, which involves greater regulatory complexity. In addition, autonomous flight operations are technically more straightforward in Variant A. For example, in industrial zones, the establishment of dedicated no-fly zones for other drones can facilitate safe and predictable flight paths. Autonomous operations in such controlled environments also enable labour cost reductions [27]. While vans remain more cost-effective in scenarios with high route density and large customer volumes, UAVs may offer competitive advantages in terms of delivery speed, operational costs, and environmental impact when delivery points are geographically dispersed [17]. However, when evaluating Variant B, parcel delivery from fixed distribution centres to variable consumer locations, Kennisinstituut voor Mobiliteitsbeleid [27] conclude that vans remain preferable due to their route density advantage.

Garg et al. [17] also highlight potential cost savings and time reductions through the concept of Synchronized Delivery (SM), which corresponds to Variant C, where UAVs collaborate with trucks or other vehicles. In this configuration, the delivery van serves as a mobile distribution centre, enabling UAVs to perform short-range deliveries [27]. Despite its conceptual advantages, variant C faces significant regulatory and operational challenges.

The literature presents different perspectives on the optimal geographical context for UAV-based delivery. On the one hand, the [27] suggests that drones can offer cost advantages in rural areas, where low route density makes conventional delivery methods less efficient. On the other hand, [58] argues that UAVs are likely to gain their greatest competitive advantage in densely populated (sub-)urban areas, where faster delivery times and congestion avoidance create additional value. These contrasting findings underline the importance of exploring both rural and (sub-)urban scenarios within the Dutch context.

In summary, the reviewed literature indicates that the competitive feasibility of UAV parcel delivery is highly dependent on the specific delivery scenario. UAVs are not expected to replace vans in general freight transport, but may offer meaningful advantages in niche markets. The most promising use cases for UAV delivery involve:

1. Fast delivery requirements,
2. Low route density,
3. Short to medium delivery distances,
4. High utilisation rates to ensure cost-effectiveness.

This study focuses on Variant A, which offers the highest practical feasibility. Operating between fixed origin and destination points simplifies regulatory approval, reduces operational complexity, and allows for cost savings through automation. To comprehensively assess the competitive feasibility of UAV corridors for freight transport, both rural and urban contexts are explored within the Dutch airspace environment.

While competitive feasibility is a necessary condition for UAV-based freight transport, it is not sufficient

on its own. UAV deployment must also comply with safety requirements and legal frameworks within the existing airspace structure. The following section examines the spatial and legal constraints that determine where and how UAV operations can be realised within Dutch airspace.

2.3. Airspace constraints for UAV corridor integration

The Dutch airspace is small, densely utilized, and highly complex [45, 38]. These characteristics pose a significant challenge for integrating a fixed-location UAV network into the existing airspace structure, as conventional aviation already occupies a substantial portion of the available airspace. Beyond commercial aviation, Dutch airspace is also used by numerous other stakeholders, including private and business aviation, recreational flying, military operations, and emergency services [45].

2.3.1. Structural complexity of Dutch airspace

As mentioned earlier, Dutch airspace is densely and very complexly organized [45, 38]. Figure 2.1 presents a snippet of the Dutch airspace map. Each segment of airspace is classified according to the seven ICAO air traffic service classes (A to G), based on the intended use of that specific region. Traffic within these classes is regulated under Instrument Flight Rules (IFR), Visual Flight Rules (VFR), or a combination of both [45].

Figure 2.1: Aeroneutical chart ICAO of the Netherlands [32]

The dominant airspace structures visible on the map are Terminal Manoeuvring Areas (TMAs), Control Areas (CTAs), Upper Control Areas (UTAs) and Control Zones (CTR). TMAs are established to safely and efficiently route arriving and departing aircraft around major airports. CTAs define the operating area of Luchtverkeersleiding Nederland's (LVNL) general air traffic control service (Area Control Service), which is responsible for en-route traffic along designated airways. In addition to airway management, CTAs include arrival routes for several Dutch airports. UTAs, on the other hand, are responsible for handling en-route traffic in upper airspace. Finally, the Control Zone (CTR), a circular area around airports where tower control is responsible for air traffic services also provided by LVNL [45].

At first glance, this level of airspace utilisation may suggest limited room for UAV operations. However, this is not necessarily the case. Under the U-space airspace review framework of Ministerie van Infrastructuur en Waterstaat [38], U-space services are primarily intended for Very Low Level (VLL) airspace, defined as extending up to 500 feet above ground level. This region largely falls within Class G uncontrolled airspace. As a result, UAV corridors operating within U-space do not intersect with TMAs, CTAs, or UTAs, as illustrated in Figure 2.2.

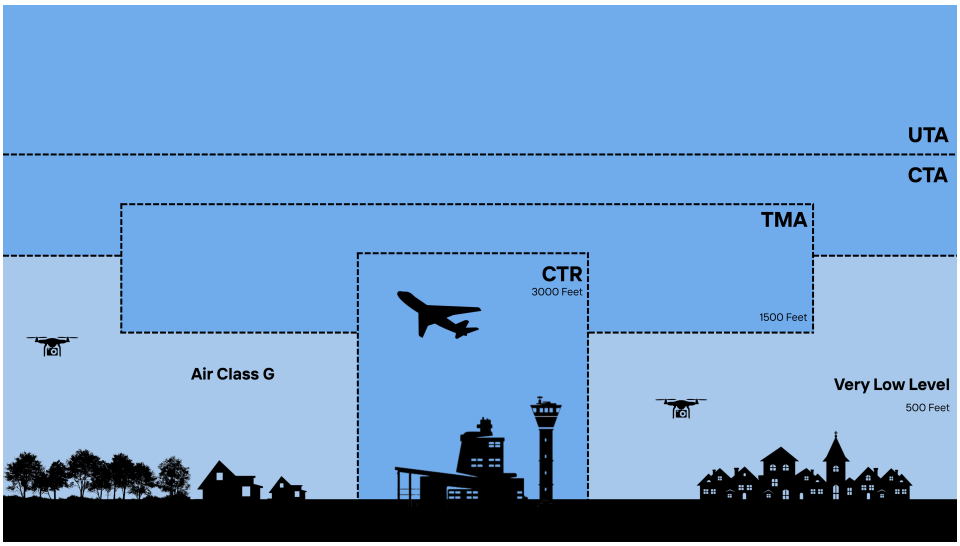

Figure 2.2: Representation of Dutch Airspace

2.4. Classification of Very Low Level (VLL) airspace

Based on the assessment framework by Ministerie van Infrastructuur en Waterstaat [38], the UAV operations will primarily target very low level (VLL) airspace, defined as extending up to 500 feet. However, the report also states that there is a direct interaction between traffic operating in VLL airspace and in adjacent airspace above. This interaction suggests that higher altitudes, up to 1500 feet, could potentially be included in UAV operations in the future. Therefore, drawing on the reports by Ministerie van Infrastructuur en Waterstaat [38] and Platform Nederlandse Luchtvaart [45], the following section provides an overview of the Dutch airspace classification up to 1500 feet, in order to identify the relevant considerations for the placement of fixed-to-fixed UAV networks.

2.4.1. Controlled airspace (CTRs)

The Dutch controlled airspace below 1,500 feet contains ‘Control Zones’ (CTRs) [45, 38]. Within these CTRs, air traffic control is provided by tower controllers to regulate air traffic near airports. Dutch civil airspace consists of five CTRs around Schiphol, Rotterdam, Beek, Eelde and Lelystad, all within airspace class C. Drone operators maintain contact with air traffic control by means of radiotelephony within this area, and both IFR and VFR traffic is allowed provided permission is obtained from air traffic control. Besides civil CTRs around airports, there are also seven CTRs in military controlled airspace in Leeuwarden, Volkel, Gilze-Rijen, Eindhoven, Woensdrecht, De Peel and De Kooy, classified in airspace class D [38].

2.4.2. Special use airspace

In Dutch airspace, specific regions are designated as special air traffic areas, which may be either permanently assigned or temporarily established within certain sections of the airspace. This designation takes place in situations and locations where potential hazards to air traffic or ground installations exist, resulting in temporary or permanent restrictions on access to airspace [38]. The different special air traffic areas distinguished in the Netherlands are described in Table 2.1 [45, 38].

A distinct category are the military low flight zones (GLV) and low flight routes (link). Lower altitudes may be flown in these areas, in deviation from the standard minimum safe flight altitude. Low-flight

Table 2.1: Special Airspace Categories in the Netherlands [45]

Name	Abbreviation	Description
Restricted Area	EHR	Area with temporarily restricted access for civil traffic. This is communicated via daily publications (NOTAM). Often, an EHR is related to military exercises.
Prohibited Area	EHP	Area permanently prohibited for civil air traffic. This area can often be found near Royal Palaces.
Danger Area	EHD	Area that cannot be classified as restricted or prohibited but should be avoided by civil air traffic during specific times. Often used for military shooting exercises. Pilots must assess whether or not to overfly these areas.
Temporary Reserved Airspace	TRA	Areas designated for specific activities during certain periods of the day, such as test flights. These activities are often incompatible with other air traffic.
Temporary Segregated Airspace	TSA	Areas reserved exclusively for specific activities during certain periods of the day for exclusive and special use.
Aerodrome Traffic Zone	ATZ	Areas around airports designated to protect local air traffic. Aircraft not operating to or from the designated airport should avoid these zones.
Special Rules Zone	SRZ	Area where special rules apply to ensure the protection of specific types of air traffic, such as gliders.

areas are not closed to other air traffic; however, they are marked on the map as an area of concern for other pilots [38].

2.4.3. Uncontrolled airspace (Class G)

In addition to civil and military controlled airspace and special air traffic areas, the remaining portions of Dutch airspace are uncontrolled for both IFR and VFR traffic and are classified as Class G airspace [38]. This airspace includes airspace over land up to a height of 1,500 feet, excluding civil and military CTRs and over the North Sea above sea level up to FL 55. In this airspace, ensuring separation between aircraft is one of the captain's responsibilities. Several VFR activities mainly take place in this uncontrolled airspace [45].

The activities conducted in this airspace include social flights (such as rescue services, aviation police, fire brigade, and medical services), business aviation, aerial work (including test and inspection flights, as well as aerial photography), education and training, and various forms of sport and recreational flying (such as gliding, private flying, and ballooning). In addition, this airspace is also used by military VFR operations, including helicopter flights and low-level training missions by fighter aircraft [38].

2.5. Infrastructure-based no-fly restrictions

Operating UAVs in densely populated urban areas introduces additional challenges related to ground-level risks such as noise, privacy concerns, and security threats [3]. To address these concerns, the Rijksoverheid [49] has established regulations specifying where UAV operations are restricted in the Netherlands for all UAV categories. These restricted areas include:

- Government buildings;
- Palaces and buildings of the royal family;
- Military (training) areas;
- Events or other temporarily restricted areas;
- Protected nature areas above the Wadden Sea;

- Certain Natura 2000 areas.

In addition, for the 'open' category, which will be discussed in more detail in Section 2.7, specific rules apply, including additional restricted zones for UAV operations [49].

- Roads with a speed limit of 80 kilometres per hour or higher;
- Railway tracks;
- Areas around airports;
- Hospital helipads;
- High-voltage power lines;
- Drinking water areas;
- Major seaports;
- Companies handling hazardous substances;
- Secured sites;
- Areas where emergency services are active;
- Low-flying routes and zones. However, drones are permitted up to a maximum altitude of 30 meters above the ground if they:
    1. Belong to category A1 and weigh a maximum of 250 grams;
    2. Belong to category A2 and weigh a maximum of 4 kilograms.

The regulations outlined in Sections 2.3, 2.4, and above reflect the current regulatory framework in the Netherlands. However, in addition to national rules, UAV operations are also governed by regulations at the European level, which will be discussed in the following section.

2.6. European regulatory foundations for UAV corridors

The design and construction of UAV networks must comply with an evolving set of European regulations that govern where and under what conditions UAVs may operate, and which airspace management services are required to ensure safe integration with manned aviation. This section introduces the key elements of the European regulatory framework, including the definition of geographic UAS zones, the U-space framework, mandatory U-space services, and operational categories. Together, these components establish the legal foundation for UAV corridor design within EU member states.

2.6.1. Geographic UAS zones

At the European level, the Implementing Regulation (EU) 2021/664 introduces the concept of UAS geographic zones, which are designated portions of airspace where UAV flights may be allowed, restricted, or prohibited to mitigate safety, security, privacy, or environmental risks [16]. In the Netherlands, the delineation and designation of these zones are regulated by national authorities [38]. UAS geographic zones may include restrictions related to proximity to low-level flight paths, airports, heliports, ports, industrial sites, critical infrastructure, or sensitive security areas [38].

Palmerius et al. [43] further distinguish between two types of spatial design elements within UAV networks:

- No-fly zones, which define three-dimensional regions that must be fully avoided by UAVs;
- Grid zones, which allow UAVs to operate along predefined grid-based routes that ensure spatial separation.

2.6.2. The U-space framework and member state responsibilities

To enable large-scale UAV integration into airspace while addressing safety and operational complexity, the European Union Aviation Safety Agency (EASA) has introduced the U-space framework [13]. U-space establishes a structured set of airspace management services and procedures to allow safe, efficient, and scalable UAV operations, including in densely populated urban areas [61].

U-space airspace is designated at the discretion of individual Member States, who retain full authority over its geographic scope, structure, access conditions, and restrictions [14]. In the Netherlands, U-space implementation remains at a preliminary stage, with exploratory vision documents and test cases, but without a fully defined spatial design for high-density urban areas [9, 36].

U-space airspace may be established in both controlled and uncontrolled airspace [13]:

- In controlled airspace, Air Navigation Service Providers (ANSPs) dynamically coordinate manned and unmanned operations to ensure separation.
- In uncontrolled airspace, U-space airspace becomes a restricted area where UAS operators must rely on U-space services, and manned aircraft are required to share positional data with U-space service providers.

All UAV operations conducted within U-space airspace must utilise a defined set of core services, as mandated by the Implementing Regulation (EU) 2021/664 [15]. These services are mandatory for all U-space participants and collectively enable real-time coordination between manned and unmanned aircraft, thereby reducing collision risks and enhancing overall airspace safety [13]. The required U-space services include:

- Network identification service;
- Geo-awareness service;
- UAS flight authorisation service;
- Traffic information service.

In addition, the designation of U-space airspace must be based on a formal airspace risk assessment that ensures the safe integration of UAV operations while maintaining the safety of manned aviation [11]. This assessment primarily addresses environmental, privacy, safety, and security considerations, with a particular focus on both air risk and the associated ground risk resulting from potential mid-air collisions. The airspace risk assessment must at a minimum include:

1. Hazard identification, including safety, security, privacy, and environmental hazards;
2. Risk analysis, evaluating the likelihood and severity of harmful effects resulting from the identified hazards;
3. The definition of mitigation measures required to maintain an acceptable level of risk.

The outcomes of this assessment serve as the basis for U-space airspace design, defining performance requirements and operational constraints. Reassessments must be conducted to accommodate major changes or to refine the airspace configuration based on operational experience or evolving environmental conditions, such as urban expansion or the development of new critical infrastructure. The assessment process must also incorporate the coordination mechanisms established under Regulation (EU) 2021/664 [11].

In addition to airspace design and risk assessment, the regulatory framework classifies UAV operations into different operational categories based on risk levels and safety requirements. These categories have direct implications for corridor design, as discussed in the next section.

2.7. Operational categories and implications for corridor use

EASA [10] distinguishes three categories of UAS operations based on safety requirements and risk levels:

1. 'Open' category: UAS operations where prior activity does not require the approval of competent authority or declaration of the UAS operator given the risks;
2. 'Specific' category: UAS operations where a permit from competent authority is required before the activity takes place given the risks. This takes into account risk mitigation measures identified in an operational risk assessment, except for certain standard scenarios where a declaration by the operator is sufficient or when the operator holds a light UAS operator certificate (LUC) with

the appropriate privileges;

3. 'Certified' category: UAS operations where, given the risks, an appropriate level of safety must be ensured. This requires certification of the UAS, a competent remote pilot and an operator approved by the competent authority.

EASA [10] indicates that, based on market priorities, regulatory development has primarily focused on the 'open' and 'specific' categories. The 'open' category serves as the regulatory standard for most recreational UAV operations and low-risk commercial applications. The 'specific' category applies when UAV operations exceed the limitations of the 'open' category and therefore require additional authorisation. Typical examples include Beyond Visual Line of Sight (BVLOS) operations and flights exceeding 120 meters in altitude [10].

BVLOS operations involve UAV flights that extend beyond the direct visual range of the remote pilot [59]. In such cases, navigation relies on instrumentation such as on-board cameras and detect-and-avoid technologies [46]. BVLOS operations are particularly relevant in environments where natural or built obstacles (such as mountains, dense forests, or urban structures) obstruct visual line of sight [6, 46]. The autonomous nature of BVLOS flights, combined with their ability to access otherwise hard-to-reach locations, makes them particularly well-suited for freight transport applications [6]. The relationship between the different operational categories and their typical operating environments is also illustrated in Figure 2.3 [30].

The growing number of UAS in the airspace, combined with the complexity of BVLOS flights in Very Low-Level (VVL) airspace, will pose significant risks to safety, security, privacy, and the environment [14]. However, the implementation of U-space services substantially facilitates the organisation of BVLOS operations, making such flights considerably more feasible than was previously the case [59].

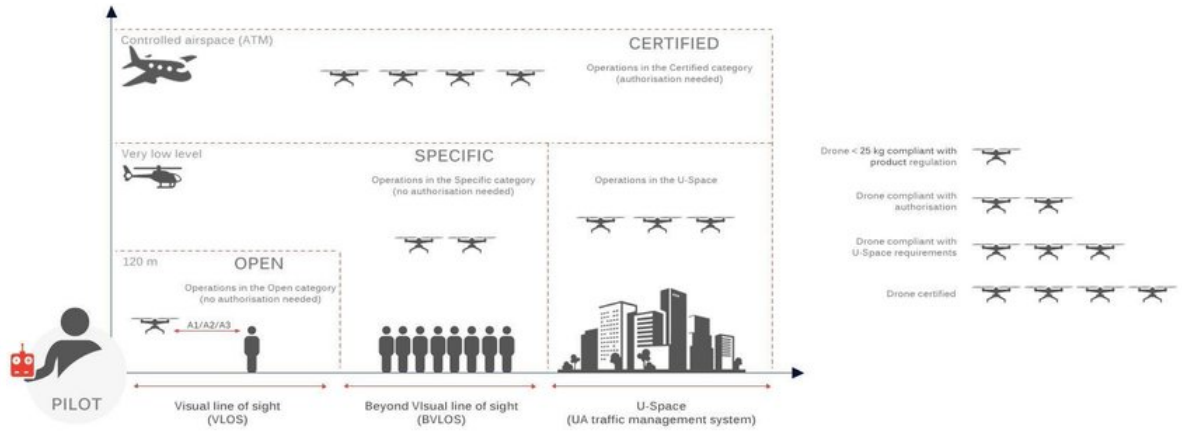

Figure 2.3: A visual representation of the adopted three categories of UAV operations [30]

While regulatory frameworks such as U-space define the legal and procedural boundaries for UAV operations and make freight transport with UAVs more feasible, they do not specify how airspace should be spatially structured. Member States remain responsible for determining this. Translating the regulations into practical and scalable solutions requires specific design principles for UAV corridors. The following section discusses these principles.

2.8. Design principles for UAV corridors

According to studies like Doole [8], Badea et al. [3], and Lyon-Hill et al. [33], autonomous UAV delivery has the potential to mitigate several societal challenges, and the U-space framework provides the necessary regulatory foundation. However, it has not yet been widely implemented in urban areas. This slow adoption is partly due to complex regulatory barriers, but the biggest challenge remains safety concerns [13], due to numerous unpredictable risks [1]. Replacing traditional parcel delivery systems with UAV-based delivery is expected to significantly increase traffic in the lower airspace [24],

particularly in urban areas. One of the key indicators of safety in airspace operations is the number of separation violations, which occur in situations in which aircraft or drones come too close to each other. To mitigate these risks, airspace design acts as a critical foundational layer of protection [63].

2.8.1. Airspace structuring concepts

The Metropolis project [63] investigated how increasing levels of airspace structure affect capacity, safety, and efficiency in UAV networks under extremely high traffic densities. It introduced four conceptual designs for open (unconstrained) airspace:

- Full mix: An unstructured airspace design, where traffic is only subject to physical constraints such as weather, obstacles, and terrain. UAVs follow direct point-to-point routes, optimizing for fuel or battery efficiency.
- Layers: In this concept, airspace is segmented into vertically stacked bands, each assigned to a range of horizontal headings. While horizontal routing remains flexible, the vertical position is constrained by heading direction, promoting traffic alignment and segmentation.
- Zones: In this approach, traffic is segmented based on similarities in driving direction. However, unlike the layer concept, the zone concept also takes into account the layout of the city when designing the structure. Here, two main types of zones are used: circular and radial zones. Both types of zones segment the airspace in the horizontal plane.
- Tubes: The final concept of tubes proposes a maximally structured design for fixed routes in airspace. This design makes use of graph theory [50], with nodes as connection points for one or more routes, and edges connecting the two nodes as tubes. The tubes never intersect horizontally beyond the nodes and are sized to allow a UAV to fly.

The Metropolis findings suggest that some structuring is necessary to maintain safety at scale, but excessive structuring, as in the Tubes concept, can significantly reduce overall capacity. The Layers concept, which combines moderate structure with flexibility, was found to be optimal in unconstrained airspace.

2.8.2. Structuring corridors in VLL urban airspace

While the Metropolis study suggests that layered designs are optimal for unconstrained airspace [63], UAV delivery flights typically operate within very low level (VLL) urban airspace, generally defined as up to 500 feet above ground level [38]. In this more constrained domain, Zhang et al. [67] argue that unstructured airspace configurations such as Full Mix and Layers are insufficient to ensure safety and efficiency. Instead, they propose a publicly organized network of fixed flight paths for low-altitude UAVs as a more effective alternative, offering enhanced safety, operational predictability, and regulatory alignment. This network structure corresponds to the fixed-to-fixed delivery model (Variant A), previously identified as the most practically viable configuration in Section 2.2.

Doole [8] adapt the Metropolis [63] insights into the highly constrained VLL urban context by aligning UAV air corridors with existing street networks. Recognizing that urban airspace is as limited and conflict-prone as the ground-level street network, they argue for spatially organized airspace corridors as aerial analogues to road systems. In their framework, two corridor configurations are compared:

- Two-Way: UAVs can travel in both directions along the same street corridor, separated vertically by heading.
- One-Way: Each corridor permits travel in only one direction

Simulation results show that the one-way concept outperforms two-way configurations in terms of safety and stability, even though it may result in longer routes. This finding aligns with broader literature suggesting that implementing structured airspace designs often entails trade-offs between safety and operational efficiency [67]. The combination of horizontal structuring, which is used to enforce one-way flow along specific corridors, and vertical segmentation, which separates travel directions across altitude layers, has proven to be the most effective configuration for managing dense drone traffic safely and predictably [8]. To be practically viable, however, such corridors must also align with ground-based infrastructure, which is explored in the next section.

Simulation results indicate that the one-way design outperforms the two-way configuration in terms of safety and stability, even though it may increase route lengths. This trade-off aligns with broader literature suggesting that structured airspace designs often improve safety at the expense of operational efficiency [67, 4]. The combination of horizontal structuring (one-way flow) and vertical segmentation (altitude-based separation) provides the most effective framework for managing dense drone traffic in VLL airspace [8]. FTo be practically viable, however, such corridors must also align with ground-based infrastructure, which is explored in the next section.

2.9. Infrastructure as spatial foundation for corridors

Doole [8] adopts the urban street network as the spatial foundation for organizing UAV corridors. The design principle applied in the study is structured routing along spatially defined linear features (in this case road networks). This could be extended to other forms of linear infrastructure and natural corridors. The Netherlands is a densely populated country with many different types of infrastructure. This provides several other opportunities for the base of networks such as railroads, canals and drainage systems. These can similarly create structured, navigable paths through both urban and rural environments. In addition, natural corridors such as rivers, ditches and meadows could provide low-risk routing opportunities in less developed areas. This study therefore expands the spatial scope of UAV corridor design by including a broader set of linearly structured land features. This builds on the core insight of Doole [8] that organized routing along constrained corridors improves safety and operational feasibility in VLL airspace.

The linear infrastructures considered in this study are;

- Infrastructure: Motorways, major roads, regional roads, tracks and rural access roads, living and residential streets, pedestrian and cycling paths and railways;
- Ground use: Agricultural lands, farmland, industrial zones, commercial zones, retail zones, residential areas and recreational zones;
- Natural corridors: Meadows, open grass, forests, woodlands and wetlands;
- Water Features: Rivers, canals, streams, lakes and ponds.

2.10. Balancing risk and operational efficiency in corridor design

While structured corridor designs enhance safety, they often involve trade-offs with operational efficiency. Increased structure may reduce capacity and flexibility, requiring careful balancing of safety, throughput, and performance objectives [67, 63]. Bauranov and Rakas [4] highlight that unstructured designs such as Free Flight offer maximum capacity and flexibility but require highly advanced technology and pose increased safety risks. In contrast, more restrictive structures like tubes and fixed lanes accommodate less-equipped UAVs and simplify safety management but limit throughput and may increase delays. Therefore, corridor design must carefully balance operational efficiency with acceptable safety levels.

2.10.1. Maintaining operational efficiency

The optimal flight altitude for UAV corridors is highly context-dependent. Ren and Cheng [48] demonstrate that for open or sparsely populated environments (e.g., lakes, forests, roads, and meadows), the lowest aggregate third-party risk occurs at approximately 30 meters altitude. In these areas, lower flight heights reduce both crash severity and noise exposure due to limited human presence and minimal physical obstacles. Conversely, for built-up urban areas, higher altitudes are safer. Buildings provide structural shielding for people, while higher flight levels also reduce privacy and noise concerns for residents. As a result, Ren and Cheng [48] recommend flight altitudes around 60 meters when operating above densely built-up zones. However, transitions between these altitudes require vertical ascent and descent, which are energetically costly. Simulation experiments by Michel et al. [35] confirm that vertical movements demand significantly more energy than horizontal flight. This further complicates the efficiency-performance trade-off, as routes may require detours to avoid no-fly zones and altitude adjustments depending on the underlying land use.

2.10.2. Maintaining acceptable levels of safety

While maximizing operational efficiency remains important, corridors cannot simply be placed over linear infrastructure without consideration of safety concerns. Besides legal restrictions such as no-fly zones and U-space requirements, UAV routing must account for ground-level risks, including noise, privacy, and societal acceptance. Not all technically feasible routes are equally desirable: flying over densely populated areas inherently exposes UAVs to higher third-party risks, even if allowed under current regulations. To address this, risk-based path planning becomes essential. The study by Hu et al. [21] demonstrates that explicitly incorporating risk considerations into route optimization yields safer trajectories compared to purely distance-minimizing solutions. Similarly, Tang et al. [64] show how risk-avoidance objectives can be directly integrated into path planning models. These principles will be further operationalized in Chapter 4 and Chapter 5, where third-party risk is systematically integrated into the corridor network model and applied to optimize UAV routing.

2.11. Conclusion on Literature Synthesis and Design Context for UAV Corridor Planning

The literature reviewed in this chapter shows that UAV-based freight transport holds considerable societal and operational promise. However, the integration of UAV corridors into national airspace systems remains a significant challenge, especially in countries like the Netherlands, where airspace is both densely regulated and spatially constrained.

The literature further emphasises that the competitive feasibility of UAV-based freight transport depends on the scenario and is not a universal substitute for traditional delivery by vans. Instead, UAVs are particularly promising in niche applications characterised by fast delivery needs, low route density, short to medium transport distances and high utilisation rates to achieve cost-effectiveness. In this context, fixed-point delivery networks (variant A) emerge as the most practically viable model, as they allow for regulatory simplification, operational predictability and automation-driven cost reductions. Moreover, both urban and rural contexts can offer different but relevant opportunities for UAV deployment, reinforcing the need for spatially adaptable corridor designs within Dutch airspace.

UAV operations typically occur within Very Low Level (VLL) airspace, defined as extending up to 500 feet above ground level. In the Dutch context, corridor planning must account for the presence of Control Zones (CTRs), special use airspace, and uncontrolled Class G airspace. UAVs must either navigate around or comply with the operational requirements of these zones. Furthermore, operating UAVs over densely populated urban areas introduces additional risks related to noise, privacy, and ground safety. To address these concerns, the Rijksoverheid [49] has issued national regulations that designate specific no-fly zones for UAV operations across all UAV categories, further constraining the available airspace.

In addition to national restrictions, European regulations further impact UAV freight operations. As such operations typically involve BVLOS flights, they fall under the EU's 'specific' category. To ensure safe UAV integration, the U-space framework has been introduced, providing Member States with guidelines for organizing UAV airspace. While Member States can determine the geographic scope of U-space corridors, their designation must be supported by formal airspace risk assessments and the implementation of key U-space services, including network identification, geo-awareness, flight authorization, and traffic information.

Unstructured airspace concepts such as Full Mix and Layers have been shown to be insufficient in ensuring safety in VLL environments. Conversely, structured corridor networks enhance both operational predictability and regulatory enforceability. For instance, Doole [8] demonstrate that fixed, one-way air routes contribute significantly to safety and conflict avoidance in urban airspace. Building on this insight, this thesis proposes a structured network of fixed one-way UAV flight paths as the most effective strategy for managing high-density drone traffic.

While structured UAV corridors improve airspace safety and regulatory enforceability, they also introduce trade-offs that affect operational efficiency. Increased structure can reduce throughput and flexibility, especially when altitude shifts or detours are needed to avoid no-fly zones or high-risk ground areas. Literature shows that optimal corridor placement and altitude vary by environment: lower alti-

tudes minimize risk and noise over sparsely populated zones, while higher altitudes are preferred in urban contexts to reduce ground-level exposure. However, altitude changes increase energy costs, complicating the design of efficient and safe flight routes. These trade-offs reinforce the need for corridor designs that balance third-party risk with performance considerations, a principle that underlies the spatial risk assessment and routing models developed in subsequent chapters.

Unlike Doole's work, which focuses solely on urban road networks, this study extends the spatial foundation to incorporate a broader set of linear infrastructures, including railways, canals, and natural corridors. While such structuring enhances safety and regulatory compliance, it also introduces trade-offs in operational efficiency and maintaining acceptable levels of safety that must be carefully managed.

From this synthesis, three key design requirements can be identified that must be addressed to enable operationally viable UAV freight operations:

- Competitive feasibility relative to conventional delivery systems: UAV freight transport must achieve sufficient operational efficiency to remain competitive with traditional delivery methods. This requires focusing on scenarios characterized by high-speed delivery demands, low route density, short to medium travel distances, and fixed-to-fixed delivery models that support automation and cost control.
- Spatial limitations due to Dutch regulatory airspace constraints: UAV corridors must navigate around extensive no-fly zones arising from national regulations, including controlled airspace segments, protected environmental zones, and critical ground-based infrastructure.
- Regulatory compliance with U-space requirements: The design of UAV corridors must facilitate airspace risk assessments and enable the technical deployment of U-space services such as geo-awareness, network identification, flight authorisation, and airspace monitoring, as mandated under European regulation.

Beyond these three design requirements, UAV corridor planning must also navigate the inherent trade-off between operational efficiency and third-party safety. While the regulatory and spatial constraints define the solution space, the corridor design itself must strike a careful balance within that space. Altitude profiles, route detours, and area-specific avoidance strategies all influence how efficiently UAVs can operate without compromising safety.

Contribution of this study

Despite recent advances, a critical gap remains in the literature: no existing framework integrates (1) the full spectrum of regulatory constraints specific to Dutch and European airspace, (2) systematic third-party risk assessment embedded within the design process rather than applied retrospectively, (3) diverse infrastructure types beyond urban road networks, and (4) flexible trade-offs between safety and efficiency within a single, implementable model. This gap is particularly pronounced in the Dutch context, where dense airspace usage, extensive no-fly zones, and evolving U-space requirements create a uniquely constrained operational environment not holistically addressed in existing studies.

3

Corridor design

This chapter develops the spatial design methodology for constructing UAV freight corridors within Dutch airspace. Building on the literature review in Chapter 2, a fixed one-way corridor network is identified as the most viable design approach for enabling safe, scalable UAV operations at low altitudes. While such fixed corridors may require longer flight paths compared to fully flexible routing, they offer significant safety, regulatory, and operational advantages, particularly in complex and densely populated environments. The design framework developed in this chapter integrates regulatory constraints, spatial no-fly zones, infrastructural opportunities, and operational feasibility into a coherent network model that forms the foundation for subsequent risk assessment and corridor optimization.

The chapter is structured as follows. Section 3.1 introduces the overall design principles and explains how fixed one-way corridors address key design requirements related to feasibility, safety, and regulatory compliance. Section 3.2 formalizes the corridor system as a graph-based network model, describing how infrastructural and land-use data are translated into a spatial routing framework. The subsequent sections detail how official no-fly zones, ground-based restrictions, and socially sensitive areas are incorporated as constraints into the network design, ensuring that proposed corridors remain both operationally viable and aligned with regulatory and societal expectations. Finally, Section 3.3 summarizes the key modeling assumptions and limitations that underlie the proposed corridor design.

3.1. Design framework

Section 2.8 established that adding structure to the design of UAV corridors can increase safety and that focusing on a fixed-to-fixed UAV network provides the most intrinsically safe operational configuration. Doole [8] proposed an initial design approach by using the existing road network as the spatial basis for UAV corridors. This study builds on that concept and extends it to include a wider range of linear infrastructure and natural corridors, as elaborated in Section 2.9.

However, as discussed in Chapter 2, the integration of UAV freight transport in Dutch airspace ultimately faces three main challenges: (i) ensuring competitive feasibility compared to conventional delivery modes, (ii) dealing with limited airspace availability due to airspace zones with restricted access, and (iii) complying with European regulatory frameworks, particularly U-space requirements. The design approach proposed in this study addresses these three challenges simultaneously.

First, fixed-point delivery networks (Variant A) emerge as the most practically viable operating model, as they allow for regulatory simplification, operational predictability, and automation-driven cost advantages Kennisinstituut voor Mobiliteitsbeleid [27]. The proposed design employs one-way fixed corridors, which align with this model. In such systems, the aviation authority only needs to authorize specific, predefined routes rather than a broad area, substantially reducing both administrative complexity and regulatory burden. Furthermore, autonomous flight operations are technically more straightforward within these predefined corridors. Autonomous operations in these tightly controlled environments additionally support reductions in labour costs. Moreover, Kennisinstituut voor Mobiliteitsbeleid [27] in-

dicates that, for non-fixed-to-fixed delivery models, vans currently remain the more economically viable alternative, reinforcing the relevance of focusing on Variant A for UAV-based freight transport.

Second, within Very Low Level (VLL) airspace, the available flight space is already heavily constrained by controlled airspace segments and numerous no-fly zones. The use of fixed predefined corridors simplifies UAV network implementation, as routes can be planned in advance to avoid restricted areas. Once established, corridors also provide a clear and predictable structure for other airspace users, who can easily avoid these zones when planning their own operations.

Third, compliance with U-space regulation requires several mandatory services, including network identification, geo-awareness, and flight authorization. The use of fixed corridors aligned with existing linear infrastructures significantly facilitates the technical deployment of these services, as essential components such as 5G antennas, monitoring systems, and management equipment can be more easily installed and maintained along accessible ground infrastructure, avoiding the complexities of dense urban environments. Another benefit is that the airspace risk assessment required by Article 3(4) of Regulation 2021/664 can be conducted at the corridor level rather than for individual flights, improving both consistency and administrative efficiency, particularly when applied at scale. Additionally, as will be further discussed in Chapter **??**, the proposed framework enables corridor planning to explicitly account for area-specific risk variations through a structured risk-based assessment approach.

Figure 3.1 provides a simplified conceptual representation of how a UAV corridor can be designed based on the proposed design framework. The illustration shows a one-way fixed-to-fixed corridor that follows a sequence of linear infrastructures and natural corridors, connecting two predefined points while avoiding restricted areas. This visualisation reflects the key design principles described in chapter sec:infra-as-design and shows how multiple types of infrastructure can be combined within a single corridor structure.

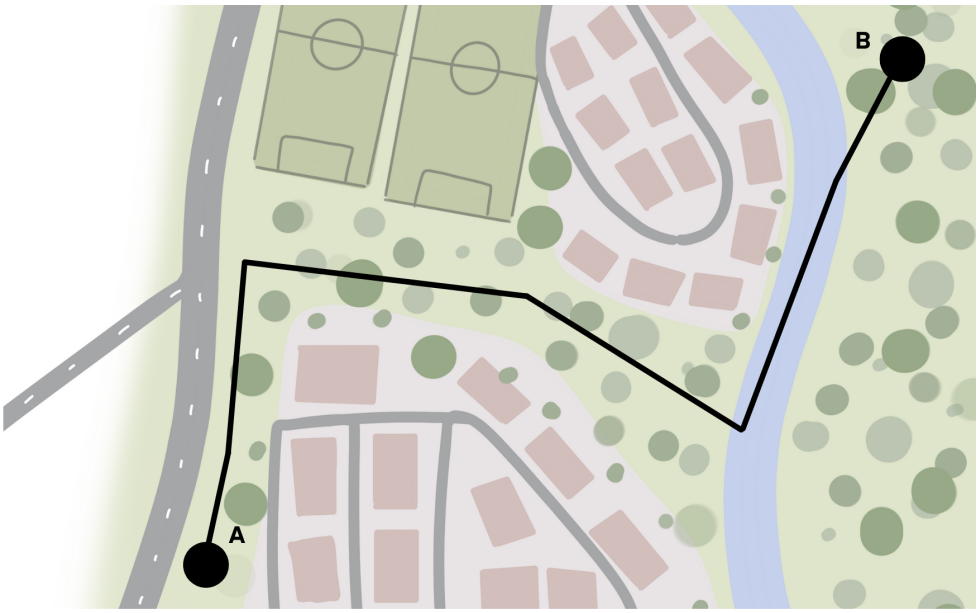

Figure 3.1: Representation of one-way corridor design based on linear infrastructures in Dutch airspace

The previous sections set out the conceptual principles for corridor design based on regulatory, spatial and operational considerations, but these principles need to be translated into a formal network structure to enable systematic modelling, risk assessment and path planning. The following section describes how these principles are operationalized into a formal network model.

3.2. Applications of graph-based routing

The following section formalizes the corridor design framework by representing the system as a transportation network. A transportation network can refer to either a scheduled service (such as an airline,

bus, or train) or a permanent track (such as roads, railroads, and canals). In general, they include various forms of connections that facilitate movement between locations. A network is composed of nodes (e.g., ports, substations, or airports) and edges (e.g., shipping lanes, power lines, or air routes), which together form an interconnected system [50]. All transportation networks can be depicted using graph theory [50], as exemplified in the tube concept by Sunil et al. [63]. Rodrigue, Comtois, and Slack [50] describes graph theory as *"a branch of mathematics concerned with how networks can be encoded and their properties measured."*

This representation is particularly relevant for UAV corridors, which can naturally be modelled as graphs composed of nodes (e.g., starting points, intersections, transition points, or waypoints) and edges (e.g., corridor segments connecting two locations). The proposed design of fixed UAV corridors forms a structured network, representing a system of predefined routes between origin and destination points [50]. TUAV operations include multiple departure, delivery and intersection locations, which together define the collection of nodes within the network. The individual corridor segments connecting these locations correspond to the network edges and represent the physical flight paths.

Graph-based path planning is widely applied across numerous domains beyond UAV corridor design. For example, Pang et al. [44] explicitly frame UAV risk-based path planning as a graph problem, where routing is performed over a weighted network in which edge weights represent third-party ground risk. Similarly, Silveira, Teixeira, and Guedes Soares [60] apply Dijkstra's algorithm to maritime navigation, constructing a route network from historical ship movement data, where edge weights reflect route frequency to prioritize safer and more commonly traveled paths. In infrastructure networks, Mu et al. [40] employed Dijkstra's algorithm to optimize water distribution systems by identifying efficient shortest paths between reservoirs and demand points. In the field of communication networks, Karve and Kapadia [25] proposed a multi-criteria version of Dijkstra's algorithm that simultaneously considers energy consumption, area coverage, and operational cost. Collectively, these studies demonstrate the versatility of graph-based approaches for solving path planning problems in a wide range of structured spatial networks, including linear infrastructures and natural corridors as examined in this study.

3.2.1. Graph construction from infrastructural data

To construct a realistic and context-sensitive UAV routing network, this study integrates a broad set of infrastructural and natural corridors as candidate flight paths. These include major transportation infrastructures such as highways, railways, pedestrian and cycling paths, waterways, as well as selected categories of open land use. The inclusion of these corridors was guided by criteria of spatial continuity, UAV operational suitability, and data availability within open-access geospatial datasets. A comprehensive overview of all corridor types and the associated extraction procedures is provided in Appendix A and Appendix B.

All data collection, cleaning, and analysis procedures were implemented using Python notebooks, which are publicly available in the accompanying GitHub repository at `https://github.com/catomartens/thesis_cf_martens`. The repository includes comprehensive documentation and a README file with step-by-step instructions for reproducing the analysis.

To model the spatial structure of these corridors, a graph-based representation was constructed based on geospatial data. All relevant data on infrastructure types and land use categories were extracted from OSMnx for the designated study areas; the detailed data retrieval and processing procedures are described in Appendix B. To test and evaluate corridor planning under realistic conditions, origin and destination locations are required to simulate freight transport routes. In this study, origin-destination pairs are based on real-world PostNL data, with distribution centres serving as origin points and PostNL pickup points functioning as delivery destinations. The rationale for selecting this case study and data source is further elaborated in Chapter 6.

Once the geographic data have been retrieved from OSMnx and preprocessed for analysis, the construction of a graph-based network representing all feasible UAV routes within the study area can be initiated. The graph construction process consists of several key steps, assuming that all necessary data have already been collected:

1. Node initialization: All PostNL locations (distribution and pickup points) are first added to the

graph as nodes.

2. Edge insertion from linear infrastructure: All available LineString and MultiLineString geometries from OSM are added as edges. These linear features are spatially connected by snapping nearby endpoints to ensure a coherent network structure. Each added edge is assigned a risk score and height attribute based on its corresponding area type.
3. Integration of polygons: The third step, involving the integration of polygon-based areas, is more complex and is therefore described in greater detail below.
4. Connector placement: Each PostNL point is linked to its nearest graph location via a postnl connector edge, provided the distance does not exceed a specified maximum connection threshold.
5. No-fly zone enforcement: All edges in the network that lie within designated no-fly zones are identified and removed from the graph to ensure compliance with airspace restrictions.
6. Data completeness adjustments: Where necessary, minor corrections are applied to supplement missing attributes or resolve local inconsistencies in the graph structure.

In the third step, all polygon-based areas suitable for UAV routing are integrated into the corridor network. These areas, which include open land-use zones and other spatially contiguous regions defined in the GeoDataFrame, are modelled differently from linear infrastructure by employing grid-based internal structures combined with sampled boundaries. First, for each polygon not classified as a no-fly zone, a uniform grid with 50-metre cells is generated, retaining only the cells that intersect the polygon and placing nodes at their centroids. These nodes are labelled as grid points and connected to neighbouring grid points where connections do not cross polygon boundaries, with all nodes and edges inheriting the polygon's risk, flight height, and other attributes. Next, the polygon boundaries are sampled every 50 metres to create boundary nodes that are connected to form a perimeter, and internal grid points are linked to their nearest boundary node within a 35-metre snapping tolerance to ensure internal-perimeter accessibility. Finally, integration with the existing linestring-based network is achieved by detecting and snapping all pairwise intersections within a 2-metre tolerance, introducing new intersection nodes where necessary and splitting edges accordingly. The last step connects all line intersection and waypoint nodes to the nearest polygon boundary node within a maximum distance of 15 metres, creating connector edges with a fixed height of 30 metres. This multi-stage approach ensures that polygon-based areas are seamlessly incorporated into the overall corridor network, enabling safe and efficient UAV routing both across and around these zones.

The resulting network captures all feasible paths along both linear infrastructures and polygon-based natural corridors, providing a spatially consistent basis for subsequent routing, risk assessment, and corridor evaluation. For a comprehensive overview of the underlying data sources, spatial processing pipeline, snapping procedures, and parameter settings applied during graph construction, the reader is referred to Appendix B. An illustrative close-up of the resulting corridor network in Breda is provided in Figure 3.2, while a broader representation is shown in Figure 3.3.

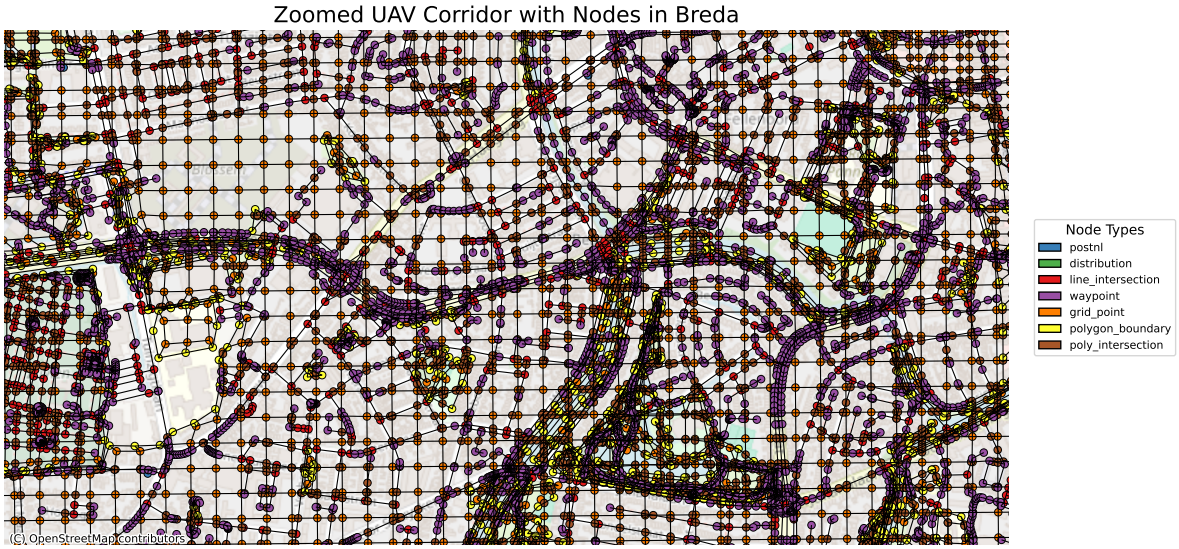

Figure 3.2: A visual example of a zoomed-in section of the resulting corridor network in Breda, the Netherlands

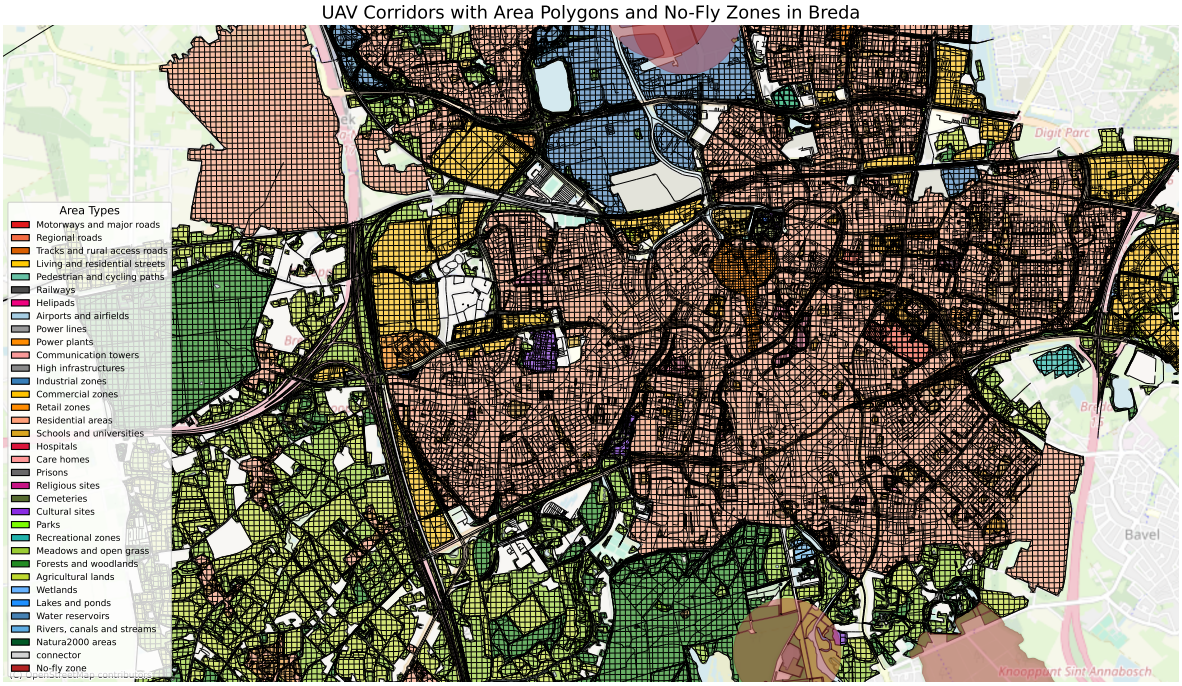

Figure 3.3: A visual example of a section of the resulting corridor network in Breda, the Netherlands

3.2.2. Definition and handling of no-fly zones

In line with the definitions proposed by Palmerius et al. [43], this study defines no-fly zones as three-dimensional areas that must be fully avoided during UAV route planning. Two distinct categories of no-fly zones are considered: areas posing risks to safe UAV flight operations due to potential mid-air conflicts, and areas where ground-based hazards introduce unacceptable safety, security, or societal risks. Both categories and their implications for corridor design are discussed below.

1. **Authoritative no-fly zones in VLL airspace**

The first category of no-fly zones consists of all airborne areas within Very Low Level (VLL) airspace that are officially designated by the Dutch government as fully restricted for UAV operations. These

zones reflect locations where UAV presence is prohibited due to elevated airspace risks and were obtained from authoritative government geospatial datasets on UAV flight restrictions. However, not all no-fly zones contained in the dataset have been incorporated in this study, as some may allow conditional UAV access or are subject to dynamic airspace management. As stated by EASA [13], within controlled airspace, Air Navigation Service Providers (ANSPs) could dynamically coordinate manned and unmanned operations to ensure safe separation, implying that UAV access may be permitted in certain cases under specific operational conditions.

A detailed overview of the included and excluded no-fly zone categories, along with the rationale for these selections, is provided in Appendix A. All official designated no-fly zones are represented as polygonal geometries, making them directly compatible with the corridor modelling framework. The following categories from the governmental dataset are considered as no-fly zones in this study:

- Low-Flying Area (GLV) and low flight routes
- Firefighting Training Areas
- General Training Area (EHD)
- Prohibited Airspace (EHP)
- Permanently Reserved Airspace (EHR)

2. **Ground-based exclusion zones**

In addition to airborne restrictions within VLL airspace, several ground-based areas impose absolute constraints on UAV operations. These zones include locations that are legally protected due to security, privacy, safety, or environmental considerations. Although these restrictions do not originate directly from airspace classification, they effectively render the airspace directly above these areas inaccessible for UAV flights and must therefore be treated as absolute spatial exclusions during corridor planning. The identification of these zones was based on a combination of authoritative government geospatial datasets and supplementary regulatory information on nationally protected locations in the Netherlands. For this, relevant features were extracted from OpenStreetMap using OSMnx to capture specific infrastructure objects, and a dedicated dataset containing all Natura 2000 protected areas within Dutch territory was incorporated to ensure full coverage of environmental exclusions. A complete overview of all data sources, included categories, and selection criteria is provided in Appendix B.

Within this category of ground-based no-fly zones, three types of restrictions have been deliberately excluded from full no-fly classification: roads with speed limits exceeding 80 km/h, railways and powerlines. Complete avoidance of these infrastructures could be practically infeasible within the Dutch context due to their high spatial coverage and central role in the national transportation network. Instead, these features are incorporated into the network but assigned elevated risk scores to reflect their increased safety concerns.

The ground-based no-fly zones considered in this study are as follows:

- Protected Nature Reserves
- High-Risk Areas
- Airports
- Air Ambulance Landing Sites
- Restricted Zones
- Prisons
- Water reservoirs

Some of these features are represented as well-defined polygons in the underlying geospatial datasets, while others are point-based locations for which additional spatial assumptions and buffers were applied to define the corresponding exclusion zones. A detailed description of the data sources, processing steps, and assumptions made for each category is provided in Appendix B.

3.2.3. Socially sensitive locations to avoid (non-regulatory constraints)

In addition to the strictly defined no-fly zones discussed in the previous subsection, there exist locations that, while not officially designated as no-fly zones by governmental authorities, are nonetheless undesirable to overfly due to their societal sensitivity, functional importance, or elevated perceived risk. These include institutions and infrastructure types where UAV operations may raise significant privacy concerns, increase public anxiety, or trigger negative societal acceptance. Although these sites are not formally restricted, their avoidance aligns with ethical considerations and principles of public acceptance for UAV deployment in populated environments.

The following categories of socially sensitive locations were included:

- Schools, universities, and kindergartens
- Hospitals
- Religious institutions (e.g., churches, mosques, temples)
- Cultural heritage sites
- Cemeteries
- High-rise infrastructure
- Communication towers
- Power plants

These locations were not classified as hard no-fly zones within the corridor network, but were instead incorporated into the spatial risk assessment framework via appropriately elevated risk scores. The same risk scoring methodology applied to linear infrastructure and land-use polygons, what will be detailed in Chapter 4, was used here to ensure methodological consistency. The assigned risk levels reflect the potential severity of UAV presence or failure over these locations, rather than their formal regulatory status.

The procedures used to retrieve and process these data via OSMnx are described in Appendix B. It is important to emphasize that only a subset of socially sensitive locations could be incorporated into the current model, based on the assumption that these objects are sufficiently represented in OpenStreetMap. Accordingly, this inclusion should be regarded as a first approximation within the broader risk assessment framework. Future work may expand and refine this component by integrating additional data sources to improve the coverage of currently underrepresented object types.

3.3. Design assumptions and limitations

Having established the corridor design principles, network construction methods, and data processing procedures, this section summarizes the key assumptions underlying the proposed UAV corridor model:

- **Fixed one-way corridor structure:** The model assumes a system of fixed corridors connecting predefined origin and destination points. Each corridor is designed as one-way in terms of directional flow but allows bidirectional operation by providing separate lanes for each direction, ensuring spatial separation and conflict avoidance.
- **Spatial grid representation of polygonal areas:** Polygonal land-use areas (e.g., forests, residential zones) are discretized into internal grid structures. UAVs follow grid-aligned paths rather than direct lines across such areas. Although this introduces routing inefficiencies, it enables computational feasibility for network generation.
- **Slot-based traffic management:** UAVs are assumed to operate under a pre-allocated slot system managed by the U-space service provider. Time slots are assigned to ensure spatial and temporal separation between UAVs, minimizing potential airspace conflicts.

- **Simplified intersection modeling:** Potential conflict zones at corridor intersections are simplified in the current model. Detailed intersection management, such as conflict resolution protocols or vertical separation at nodes, is considered beyond the scope of this study.
- **Continuous positioning and communication:** UAVs are assumed to maintain continuous GNSS-based positioning (e.g., GPS) and real-time data exchange through reliable mobile network connectivity (e.g., 5G). This supports ongoing navigation updates, peer-to-peer situational awareness, and coordination with U-space service providers.
- **Centralized U-space coordination:** The model assumes that a centralized U-space management system continuously monitors UAV operations, coordinates slot allocations, enforces traffic rules, and manages real-time adjustments as necessary to maintain safety and operational efficiency.

3.4. Conclusion of Corridor Design

This chapter has developed a spatial design framework for UAV freight corridors that integrates regulatory, operational, and spatial constraints into a structured network model. Building on the one-way fixed corridor concept by Doole [8], the design incorporates both linear infrastructures and polygonal land-use areas to create a graph-based corridor network. In doing so, the proposed design directly addresses the three key integration challenges identified earlier in this study: ensuring competitive feasibility, navigating limited airspace availability, and enabling compliance with European U-space regulations, while simultaneously providing a structured foundation to enhance safety in UAV operations.

The proposed network structure was operationalized through a graph construction process using OSMnx-derived geospatial data for the Dutch context. Distribution centers and pickup points from PostNL served as origin-destination nodes within this network, with linear infrastructure and polygon grids defining the feasible edge structures. Additional design assumptions, such as slot-based traffic management, GNSS-based positioning, continuous connectivity via 5G networks, and simplified intersection handling, ensure that the proposed model remains computationally tractable while still capturing essential safety and feasibility dimensions.

The resulting network model represents a regulation-compliant UAV corridor network that provides the spatial foundation for subsequent analyses. While this model now defines all technically feasible routes within the design space, not all routes are equally desirable from a safety and operational perspective. In the following chapters, area-based third-party ground risks will be incorporated into the network, after which risk-informed path planning will be applied to identify the optimal trade-offs between operational efficiency and safety performance.

4

Risks assessment framework

Having established the spatial design and graph-based structure of the UAV corridor network in Chapter 3, this chapter focuses on the integration of third-party ground risks into the corridor planning framework. While the previous chapter defined all technically feasible routes, it did not yet differentiate between corridors based on safety or societal risk exposure.

Section 4.1 introduces the core ground risk considerations that motivate the integration of risk assessment into UAV corridor design. Section 4.2 provides an overview of existing UAV risk assessment methodologies, distinguishing between quantitative, qualitative, and hybrid frameworks. This includes a detailed discussion of existing Third-Party Risk (TPR) models and integrated risk frameworks relevant to UAV operations. Section 4.3 then introduces the proposed semi-quantitative risk assessment (SQRA) framework developed for this study. In this section, the three core steps of the SQRA process are systematically applied to evaluate area-specific risks across the UAV corridor network.

4.1. Ground risk considerations in UAV corridor design

Studies have shown that UAV operations introduce safety risks both during flight and at critical phases such as take-off and landing, with potential for crashes or package drops that may harm pedestrians, vulnerable individuals, pets, or property on the ground, particularly due to exposed propellers and uncontrolled impacts [24]. In order to enable UAV operations over densely populated environments, minimizing third-party ground risk has become an important prerequisite for safe deployment [53]. As a result, increasing attention has been directed toward structured airspace designs in which UAVs follow predefined routes that limit their exposure to ground-based hazards [67, 8]. Naturally, among the set of possible paths connecting origin and destination points, some routes present higher or lower risk levels depending on the ground context they traverse.

In addition, European regulations mandate that UAV operations must be supported by a formal airspace risk assessment. To safely integrate UAVs into both urban and rural airspace structures, EU Member States are required to delineate UAV corridors based on these assessments [13]. Since performing detailed risk evaluations for each individual UAV flight would be impractical at scale, it is advantageous to incorporate risk assessments already during the corridor design stage. The following sections introduce the methodological framework applied in this study to systematically estimate ground risks across different areas and to integrate these risk assessments directly into UAV corridor design.

4.2. Overview of UAV risk assessment approaches

Risk assessment methodologies are traditionally classified into two broad categories: quantitative and qualitative approaches [34]. Quantitative risk assessments rely on statistical and mathematical models to produce accurate, numerical evaluations of the likelihood and impact of risks. Although these methods are often data-intensive and require considerable modeling effort, they provide results that can be incorporated directly into optimization frameworks and simulation models. Qualitative approaches, on

the other hand, usually rely on expert judgment, checklists or descriptive frameworks. These methods are more accessible and easier to interpret for a wide range of stakeholders, but they can be more subjective and less reproducible [34].

Both approaches have advantages and limitations, and their appropriateness depends on the context of the assessment. Several quantitative methods have already been applied in the UAV domain to support risk-informed decision making. These are discussed below.

4.2.1. Quantitative Third-Party Risk models for UAV operations

Third Party Risk (TPR) refers to the safety concerns and property loss that UAV operations can cause to people, property or infrastructure not directly involved with the UAV itself [64]. In recent years, a wide range of TPR models have been developed to support risk-aware path planning for UAVs, especially in urban environments. Most existing models represent the environment as a 2D or 3D grid, where each cell is assigned a risk score based on one or more components. These components typically include:

- **Fatality risk:** All risk UAV that cause injuries or fatalities to pedestrians and/or people in vehicles.
- **Property damage risk:** Falling UAVs can hit critical infrastructure or collide with tall buildings, causing damage.
- **Societal impact risk:** Broader implications such as noise pollution, privacy implications or public acceptance of drones.

These components are usually combined into a single risk cost value using weighted sum or probabilistic methods. This risk value is then assigned to the associated grid cell.

For instance, Pang et al. [44] integrate fatality risk, property damage risk, and societal impact into an area-based cost model using population and traffic density as inputs. Zhou et al. [69] similarly use a 3D grid to model risk at different altitudes, incorporating explicitly the probabilities of failure of the UAV at spatial exposure levels. Tang et al. [64] define risk in terms of obstacles, fatality risk, and property loss. These are quantified and combined using a weighted sum. As with previous models, the accuracy and relevance of the results depend heavily on the availability of environmental data and UAV specifications, as well as careful calibration of weights and thresholds.

Other models, such as Hu et al. [21], include damage to people on the ground, damage to vehicles and collision risk with manned aircraft near airports. Their framework uses physical parameters (e.g., UAV speed and weight) and derives human impact through kinetic energy and sheltering models. The study shows that risk-based routing can produce safer paths than shortest-distance planning. Primatesta, Guglieri, and Rizzo [47] also deviate from full third-party risk modelling by applying a risk map based solely on population density, aiming only to avoid highly populated areas.

A related but distinct line of work includes Kim and Bae [29], who use third-party risk assessment to compare the safety of predefined corridors without performing active path optimization. Similarly, Zhang et al. [68] and Rudnick-Cohen, Herrmann, and Azarm [53] assess the expected ground fatality risk along fixed routes using probabilistic crash simulations and population data, but do not dynamically adjust the path.

In contrast to these studies, several other works emphasize different types of risk models. De Filippis, Guglieri, and Quagliotti [7], for instance, do not model third-party risk but focus entirely on terrain-based risk. In this study risk is defined as the UAV's proximity to the ground or obstacles based on the orography of the area. Finally, López et al. [31] focus exclusively on collision risk between UAVs in shared airspace. In their model, risk is defined as the distance between UAVs, with conflict avoidance strategies ensuring minimum safe separation. Ground risk or third-party risk factors are not considered in their approach. While these studies provide valuable perspectives on UAV risk, they are not directly applicable to the research objective of this study. The current work focuses exclusively on third-party ground risk for the design of spatially fixed UAV corridors. Models that emphasize terrain risk or aerial conflict avoidance are outside this scope.

4.2.2. Integrated UAV risk framework

Besides purely qualitative approaches to UAV risk assessment, frameworks can also be adopted that integrate both qualitative and quantitative components. For example, Allouch et al. [1] propose a systematic UAV safety framework to ensure the safe operation of UAV missions over the Internet of Drones (IoD), combining two complementary layers. The first layer consists of a qualitative risk assessment grounded in ISO 12100 and ISO 13849 standards, while the second layer introduces a quantitative Bayesian Network model.

The quantitative component begins with the definition of system boundaries, followed by a systematic hazard identification process based on accident databases, safety reports, and expert judgement. This process results in a comprehensive classification of both internal and external risk factors, several of which directly contribute to the overall crash probability of UAV operations. Internal risks primarily involve technical failures, such as battery malfunctions, autopilot software errors, sensor inaccuracies, and mechanical breakdowns of components like motors or propellers. External risks, in contrast, are shaped by the operational environment and are inherently location-dependent. These include interference, environmental conditions, obstacles, the navigational environment, air traffic conditions, electrical hazards, communication reliability, and human factors, as summarized in Table 4.1.

Table 4.1: External Hazard Sources for UAV Operations [1]

Type	Examples
1. Obstacles	Fixed obstacles (trees, electric cables, buildings), and dynamic obstacles (birds, cars).
2. Interference	Electromagnetic interference (EMI), humans (eavesdropping of radio signals), communication interference.
3. Communication	Network congestion, network unavailability/delays, network jitters.
4. Navigational Environment	GPS signal loss/error, GPS spoofing, ADS-B signal inaccuracy, navigation system error, attitude error, erroneous waypoint.
5. Electrical Environment	Man-made or natural RF fields such as High Intensity Radio Transmission Areas (HIRTAs), electrostatic phenomena.
6. Air Traffic Environment	Other aircraft in close proximity, classes of airspace that may be flown nearby.
7. Environmental conditions	Wind, temperature, atmospheric attenuation, icing, precipitation, visibility (day or night).
8. Human Factor	Lack of safety culture awareness, security attacks (on the ground control station, on the data link, on UAV), pilot error (inexperienced pilots, unfamiliarity with area, fatigue, rush).

The identified risks are evaluated using a standardized risk matrix that combines the likelihood of occurrence with the severity of potential consequences. This matrix-based approach classifies risks along two primary dimensions:

1. *Severity*, representing the potential harm to humans, property, or the UAV system; and
2. *Likelihood*, representing the estimated probability of the hazardous event occurring under specified conditions.

The overall risk score is calculated as the product of these two dimensions. This structured framework offers a transparent and adaptable method for evaluating operational hazards associated with UAV activities. Based on the resulting risk scores, appropriate mitigation measures are proposed, ranging from inherently safe system design and protective measures to operational restrictions and procedures. Finally, ISO 13849 is applied to identify the required safety functions and determine the associated performance levels (PLr) through the use of a risk graph.

As a second layer, a Bayesian Network model is integrated to extend the analysis. This quantitative extension builds upon the hazards identified in the qualitative assessment and models the causal relationships between them. It provides probabilistic estimates of crash likelihood under varying conditions, thereby offering a more dynamic and scenario-specific foundation for safety evaluation and decision-making.

In summary, despite the growing number of applied TPR models proposed in recent studies [69, 64, 44], their complexity and extensive data requirements pose challenges for practical implementation, particularly within the context of exploratory research or early-stage network design. Many of these models depend on detailed UAV specifications, probabilistic crash simulations, or high-resolution environmental datasets that may not be readily available or feasible at larger spatial scales. Given these limitations, a more pragmatic two-layered approach, such as that introduced by Allouch et al. [1], is better suited to the objectives of this study. The following section therefore builds upon this perspective by introducing a Semi-Quantitative Risk Assessment (SQRA) framework to evaluate ground risks within the corridor design process.

4.2.3. Semi-Quantitative Risk Assessment (SQRA)

Qualitative risk assessments, such as the initial layer proposed by Allouch et al. [1], rely on expert judgement and descriptive categories to evaluate risk in non-numerical terms. In contrast, fully quantitative methods apply statistical and mathematical modelling to estimate probabilities and consequences with high precision [34], as demonstrated by several advanced TPR models (e.g., Zhou et al. [69] and Tang et al. [64]). Both approaches have inherent limitations. Qualitative methods are accessible and straightforward to apply, but often lack precision and may be subject to expert bias and subjectivity. Quantitative models, while providing detailed and accurate risk evaluations, typically require large amounts of data, advanced modelling techniques, and domain-specific knowledge, as outlined in Section 4.2.1.

Semi-Quantitative Risk Assessment (SQRA) offers an intermediate approach, combining the structured reasoning of qualitative assessments with the comparability and transparency of numerical scoring systems. This hybrid framework allows for systematic comparison and prioritisation of risks without the full data demands of fully quantitative modelling [34]. According to Majka [34], a standard SQRA process consists of six steps. In this study, the first three steps are applied:

1. **Risk Identification:** Identify potential hazards or risk factors using expert input, brainstorming sessions, and historical incident data.
2. **Risk Analysis:** Evaluate each identified risk by assessing its likelihood and potential consequences. Use predefined qualitative categories and assign a corresponding numerical score to each (typically on a 1–5 scale).
3. **Risk Rating Calculation:** Calculate the risk rating for each hazard by multiplying its likelihood and consequence scores. Use a risk matrix to classify the overall risk level (e.g., low, medium, high).

These first two steps also directly correspond to the Hazard Identification and Risk Analysis stages prescribed by EASA for the airspace risk assessment required under U-space regulations [11], thereby ensuring that the proposed framework remains well-aligned with regulatory expectations and facilitates compliance in practical corridor planning.

The remaining SQRA steps, Risk Prioritization, Risk Control and Mitigation, and Monitoring and Review, fall outside the scope of this thesis, but offer valuable directions for future work. Notably, the definition of mitigation measures to ensure an acceptable level of risk is also a required component of the full U-space airspace risk assessment. Extending the current framework to include these additional steps would enable more advanced comparisons between alternative corridor designs, the formulation of risk mitigation strategies, and the continuous monitoring of operational safety.

4.3. Operationalizing Semi-Quantitative Risk Assessment for UAV corridor design

One of the key advantages of designing UAV operations around fixed predefined corridors is that risk assessment can be performed corridor-by-corridor instead of flight-by-flight. Once a corridor is assessed and approved, multiple UAVs can operate within the same safety constraints, improving regulatory compliance and operational scalability. To support this, risks must be represented spatially, capturing how UAV-related risks vary across different ground environments.

Most existing studies adopt grid-based representations, where UAVs are assumed to navigate freely through a continuous 2D or 3D space. In such models, risk scores are assigned to individual grid cells [44, 69, 68, 47, 26, 64]. In contrast, this study focuses on structured airspace, where UAVs follow fixed, infrastructure-aligned paths. Accordingly, the risk assessment is conducted at the corridor level based on the ground environment, rather than at the grid-cell level. This corridor-level assessment aligns better with regulatory requirements and simplifies risk evaluation for predefined routes.

The emphasis in this study is placed on third-party ground risk, as this represents the most immediate safety concern for UAV freight operations within the Netherlands, where densely populated, cultivated, and urbanized areas dominate the landscape. While airspace risks such as airborne conflicts or separation management are critical to dynamic air traffic management, they are usually addressed through U-space services, strategic deconfliction protocols and designated airspace structures.

This study proposes a semi-quantitative risk assessment (SQRA) approach tailored to UAV corridor design. Inspired by both qualitative frameworks and quantitative Third-Party Risk (TPR) models, the method bridges interpretability and analytical usability. It retains the transparency and policy relevance of qualitative risk matrices, while enabling spatial quantification that supports routing optimization. The proposed method is both practical and scalable, as it avoids complex models such as collision simulations, detailed UAV failure modeling, or high-resolution crash probability estimations, which are typically required in full TPR models [69, 64]. At the same time, it is modular and extensible, allowing for the inclusion of more advanced sub-models if required, without compromising transparency or regulatory applicability.

The model incorporates the three most commonly used risk domains from the UAV safety literature (fatality risk, property damage risk, and societal impact) into a structured scoring scheme. These risks are assessed for predefined corridor segments aligned with infrastructure such as roads, railways, canals, or natural infrastructure, as introduced in Chapter 3. This provides a structured and interpretable foundation for risk-based UAV routing. The following subsections describe each step of the proposed risk assessment framework in further detail.

4.3.1. Application of SQRA to spatial UAV risk

This section applies a Semi-Quantitative Risk Assessment (SQRA) framework to estimate the spatial third-party risks of UAV corridors. The approach combines different external risk factors, assigns likelihood and severity scores, and calculates overall risk levels for each area. The SQRA method is adapted to reflect how risks vary across different ground environments that UAVs may fly over. The following subsections describe each step of the process, from identifying relevant hazards to calculating the final risk scores used for corridor evaluation.

Step 1: Risk Identification

Risk identification corresponds directly to the hazard identification stage required under the U-space airspace risk assessment framework [11]. In this study, however, the scope is deliberately limited to external risks, as internal risks are typically system-specific and independent of spatial corridor design. Since the objective of this study is to assess spatial variation in risk across different corridor routes, internal system failures (e.g., mechanical faults, battery failures, software errors) are excluded from the analysis.

Similarly, broad environmental and human factors that primarily vary over time rather than space (such as weather fluctuations or operator fatigue) are not incorporated. In addition, air traffic environment risks are excluded, as these are determined by proximity to controlled airspace and are already addressed

through the no-fly zone exclusions established in Chapter 3.

The focus of this study is therefore on five spatially-defined external risk factors, adapted from the comprehensive hazard identification conducted by Allouch et al. [1]:

1. **Obstacles:** Physical obstructions that may interfere with UAV flight, including both static obstacles (e.g., buildings, trees, power lines) and dynamic obstacles (e.g., vehicles, wildlife).
2. **Interference:** Risk of unintentional signal interference caused by dense human activity and overlapping wireless transmissions (e.g., Wi-Fi, mobile phones, RF consumer electronics), degrading UAV communication quality.
3. **Communication:** Risk of UAV communication loss or degradation due to limited or unstable mobile network coverage (e.g., 4G/5G availability, network saturation), affecting real-time control and coordination functions.
4. **Navigational Environment:** Risk of degraded GNSS (GPS) positioning caused by obstruction or multipath reflection from dense built environments, such as narrow streets, tall buildings, or enclosed industrial zones.
5. **Electrical Environment:** Exposure to strong electromagnetic fields from surrounding infrastructure (e.g., high-voltage power lines, radar installations, substations, industrial machinery), potentially disturbing UAV electronic systems.

Step 2: Risk Analysis

In this study, risk analysis follows the semi-quantitative structure introduced earlier, whereby each identified external risk factor is evaluated along two dimensions: likelihood of occurrence and severity of consequences. When compared to the third-party risk (TPR) models proposed by studies such as Pang et al. [44], Zhou et al. [69], and Tang et al. [64], the likelihood dimension applied here can be seen as a simplified proxy for the UAV failure probability component modeled in those frameworks. Similarly, the severity dimension aligns with outcome-specific consequence categories frequently used in TPR literature, specifically fatality risk, property damage, and societal impact.

1. **Likelihood scoring criteria**

The probability of a risk is defined as the likelihood of its occurrence. In this study, likelihood is categorized into five ordinal categories: frequent, probable, occasional, remote, and improbable [1]:

- Frequent: A frequent event occurs repeatedly or has been observed often.
- Probable: A probable event is expected to occur regularly, though not as often as frequent ones.
- Occasional: An occasional event occurs sporadically. A remote event is unlikely but possible,
- Remote: Unlikely to occur but possible or has occurred rarely
- Improbable: Event is extremely rare or has never been observed.

To assess the likelihood of different external risk factors, each area type is assigned a likelihood score reflecting the estimated frequency of occurrence. In order to enhance objectivity and consistency, a standardized scoring procedure is applied. This structured approach ensures that likelihood estimations remain transparent, replicable, and systematically grounded in the spatial characteristics of each corridor segment. For each external risk factor, a proxy indicator was defined to serve as the basis for estimation. Based on spatial approximations and assumptions, likelihood scores ranging from 1 (Improbable) to 5 (Frequent) were assigned to each area type. An overview of the proxies used for each external risk factor is provided in Table 4.2.

It is important to emphasize that this scoring framework represents an approximation. While it provides a systematic structure for likelihood assignment, the underlying scores rely on assumptions rather than direct empirical measurement. Future research should aim to refine and validate these estimates through observation, empirical data collection, or model-based quantification. A detailed explanation of the proxies, scoring logic, and area-specific assignments is provided in Appendix D.

Table 4.2: Summary of likelihood scoring proxies and thresholds per external risk factor

External Risk Factor	Proxy Used	Freq. (5)	Prob. (4)	Occas. (3)	Rem. (2)	Improb. (1)
Obstacles	Obstacle density (objects/km^2)	>150	50–150	20–50	5–20	<5
Interference	RF emitter density (sources/km^2)	>1000	200–1000	50–200	10–50	<10
Communication	4G/5G coverage (% area with stable signal)	<50%	50–70%	70–85%	85–95%	>95%
Navigational Env.	Building count (buildings/km^2)	>400	300–400	200–300	100–200	<100
Electrical Env.	EM source density (sources/km^2)	>20	10–20	5–10	2–5	<2

2. **Severity scoring criteria**

Severity is assessed in terms of the potential consequences for human safety and the surrounding environment. Following the methodology of Allouch et al. [1], this study adopts a four-level qualitative categorization of severity, which will serve as the basis for all subsequent severity assessments:

- Catastrophic: A catastrophic event involves serious injury or loss of human life and must be addressed with the highest priority.
- Critical: A critical event causes significant damage to third-party assets or infrastructure, but does not pose a direct threat to human safety.
- Marginal: A marginal event results in damage to the drone system itself without affecting external entities.
- Negligible: A negligible event has minimal operational impact, potentially causing minor degradation in system performance without leading to mission failure.

To ensure consistency with the UAV safety literature, this study adopts the three most commonly applied consequence domains in third-party risk models: fatality risk, property damage risk, and societal impact. This structure mirrors the widely used frameworks proposed by Pang et al. [44], Zhou et al. [69], and Tang et al. [64], each of which operationalizes these dimensions in grid-based or area-specific risk models. By adopting this widely recognized structure, the present study remains aligned with existing research while ensuring comparability across both grid-based and fixed-corridor UAV routing models.

Building on this structure, the current study assesses severity as a function of the consequences of UAV operations in different spatial areas. This explicitly distinguishes between risk domains that are dependent on a crash event and those that are not. For risks of fatalities and property, severity reflects the extent of damage dependent on a UAV crash. In these domains, the probability of external hazards contributes to the probability of a crash, while severity quantifies the potential impact of that crash in the affected area. Conversely, societal severity is assessed independently of the occurrence of a crash, reflecting the level of public disruption, privacy concerns or social discomfort associated with UAV overflights. Social severity thus reflects ground-level sensitivity to the presence of UAVs, even in the absence of incidents. This separation allows for a more nuanced representation of how different area types translate UAV activity into differentiated consequence profiles across the three risk domains.

- **Fatality risk:** Fatality risk reflects the potential for direct harm to humans in the event of a UAV crash. In line with the approaches of Pang et al. [44], Zhou et al. [69], and Rudnick-Cohen, Herrmann, and Azarm [53], human exposure is proxied by spatial characteristics such as pedestrian density, land use type, and vulnerability of ground-level activities. Areas with high and unprotected public presence (e.g., residential streets, schools, playgrounds) are assigned higher severity scores due to the elevated likelihood that a UAV crash may strike individuals directly.
- **Property damage risk:** Property risk represents the potential for damage to third-party infrastructure or critical facilities resulting from a UAV crash. Similar to the models by Pang et al. [44]

and Zhou et al. [69], this study applies proxies based on the presence, density, and criticality of infrastructure. High-value or operationally sensitive facilities such as power plants, communication towers, and industrial zones receive higher severity scores because a crash at these locations may cause significant technical, financial, or cascading disruptions.

- **Societal impact risk:** Societal risk captures public acceptance concerns associated with UAV operations, independent of crash events. Following the logic of prior studies such as Pang et al. [44] and Zhou et al. [69], this dimension reflects potential social discomfort, privacy concerns, and noise disturbances caused by UAV presence itself. Higher severity scores are assigned to socially sensitive environments such as residential neighborhoods, cultural sites, and schools, where UAV overflight is most likely to generate negative public sentiment.

To ensure consistency and reduce subjectivity in the risk assessment, the severity classifications for fatality, property damage, and societal impact have been systematically standardized across area types. This standardization is grounded in a set of structured assumptions and approximations regarding spatial characteristics and operational contexts, as detailed in Appendix D and summarized in Table 4.3. By translating qualitative considerations into predefined severity scores, the framework aims to minimize interpretive variability and improve reproducibility across diverse environments. This structured scoring approach provides a transparent and replicable basis for evaluating UAV-related risks. In accordance with the SQRA methodology, each severity category is assigned a numerical score ranging from 1 to 4, which facilitates quantitative risk aggregation in the subsequent steps of the assessment.

Table 4.3: Summary of severity scoring proxies and definitions per consequence domain

Severity Dimension	Proxy Used	Catastrophic (4)	Critical (3)	Marginal (2)	Negligible (1)
Fatality risk	Direct exposure of humans to UAV crash (based on activity type and density)	High pedestrian presence in open areas	Frequent outdoor human presence	Mostly indoor population	No expected human presence
Property damage risk	Presence and criticality of third-party infrastructure	Critical infrastructure	Essential services	Minor structures	No infrastructure
Societal impact	Public sensitivity to UAV visibility/noise	Residential and private zones	Culturally and publicly sensitive	Transitional and public infrastructure	Remote, uninhabited, or water areas

Step 3: Risk rating calculation

Each spatial area can be exposed to multiple external risk factors (e.g., obstacles, interference), which may lead to fatalities or property damage. In addition, the societal consequences of UAV operations vary across areas and are not tied to failure likelihood. To aggregate these impacts into a meaningful overall score, this study applies a domain-weighted summation rather than a matrix-based classification.

In line with approaches used in Zhou et al. [69], Tang et al. [64], and Pang et al. [44], the total risk per corridor segment is computed by combining severity and likelihood scores through weighted aggregation. Importantly, severity scores are assigned at the area level, independent of individual external risk factors. This reflects the assumption that the severity of a UAV crash or presence is primarily driven by the characteristics of the area itself, rather than by the specific external hazard that triggers the event. For example, a UAV crash over a densely populated residential street will consistently yield more severe consequences than a crash over an open field, regardless of whether the failure was caused by obstacles, interference, or navigation errors.

1. Crash-related components
For each external risk factor $i \in \{1, \ldots, n\}$, the contributions to fatality and property risk are computed as:

$$R_{if} = S_f \cdot L_{if}, \qquad R_{ip} = S_p \cdot L_{ip} \tag{4.1}$$

where:

- S_f, S_p: severity scores for fatality and property damage, respectively;
- L_{if}, L_{ip}: likelihood scores for external risk factor i contributing to fatality or property damage.

Since multiple external risk factors may simultaneously contribute to crash probability, their effects are aggregated additively:

$$R_f = \sum_{i=1}^{n} R_{if}, \qquad R_p = \sum_{i=1}^{n} R_{ip} \tag{4.2}$$

The resulting cumulative scores are subsequently normalized to the range $[0, 1]$ across all area types to ensure comparability with the societal impact component.

2. Societal component
Societal risk is modeled independently, as it reflects persistent public concerns (e.g., privacy intrusion, noise disturbance) that are not conditional on UAV failure. It is captured directly as:

$$R_s = S_s \tag{4.3}$$

where:

- S_s: societal severity score assigned to the area.

This score is similarly normalized to the range $[0, 1]$ across all area types.

3. Total risk score
The final composite risk score for each area segment is computed as a weighted sum of the normalized risk components:

$$\text{Risk}_{\text{area}} = \alpha_f \cdot R_f^{\text{norm}} + \alpha_p \cdot R_p^{\text{norm}} + \alpha_s \cdot R_s^{\text{norm}} \tag{4.4}$$

where:

- α_f, α_p, α_s: domain weights reflecting the relative importance assigned to fatality, property, and societal risk;
- R_f^{norm}, R_p^{norm}, R_s^{norm}: normalized domain scores for fatality, property, and societal components, respectively.

This structure ensures that each risk domain contributes proportionally to the overall risk score, while maintaining flexibility to adjust the weights of the domains according to stakeholder priorities or policy objectives. In this study, the weights are set at $\alpha_f = 0.4$, $\alpha_p = 0.3$, and $\alpha_s = 0.3$, reflecting a normative policy assumption that preventing fatalities deserves slightly more attention than property damage and social disruption. These weights are not empirically derived, but represent a pragmatic starting point for balancing the different consequence areas within the context of risk management of UAV corridors.

The proposed framework allows for the systematic integration of multiple external, area-specific risk factors into a single composite risk score for each corridor section. By explicitly distinguishing between types of impacts, the model remains both interpretable and aligned with established third-party risk assessment frameworks. Moreover, the weighting scheme can be easily adapted to changing

policy preferences or new risk dimensions. For example, additional risk categories or shifts in social acceptance can be incorporated by adjusting the domain weights α_f, α_p and α_s, without requiring fundamental structural revisions to the model. Similarly, more sophisticated domain-specific sub-models can be integrated to refine individual risk estimates, while maintaining the transparency and modularity of the overall aggregation procedure.

4.4. Conclusion of risk assessment framework

In this chapter, a semi-quantitative risk assessment framework (SQRA) has been developed to systematically incorporate third-party ground risks into the design of UAV corridors. Building on both legal requirements and existing literature on UAV risks, the proposed approach balances analytical tractability with legal transparency by combining qualitative hazard identification with structured numerical scoring. Fully quantitative third-party risk (TPR) models often rely on UAV-specific crash simulations and high-resolution environmental datasets. In contrast, this developed framework remains practically applicable for early corridor design by using structured assumptions, area-based proxies and spatially defined probability factors.

The SQRA methodology introduced explicitly distinguishes between crash-related risks and societal risks that reflect the public's ongoing sensitivity to the presence of UAVs. Severity ratings are standardised for three consequence domains: risk of fatalities, risk of property damage and societal consequences. Score proxies were developed for each domain to allow spatial differentiation between corridor segments. The resulting multi-domain risk scores are then aggregated into a single composite corridor risk score by weighted summation. This allows flexible adjustment of policy priorities while maintaining model transparency and consistency with established risk frameworks.

The framework presented in this chapter provides a structured basis for integrating risk-based considerations into the corridor network model developed in Chapter 3. In the next chapter, this risk assessment approach is further extended to incorporate operational efficiency, enabling a combined evaluation of safety and performance in UAV corridor path planning.

5

Path Planning Methodology

The previous chapters introduced UAV corridor design, formalised the operational environment and assigned spatial risk scores to infrastructure segments. This chapter describes how these components are integrated into a risk-aware corridor planning framework that balances safety and operational efficiency. The primary objective is to identify low-risk, energy-efficient UAV routes within a predefined network structure based on existing linear infrastructures such as roads, waterways and railways.

The chapter is structured as follows. First, Section 5.1 reviews existing approaches to UAV path planning and highlights key methodological choices relevant to risk-based routing. Section 5.2 defines how operational efficiency is modelled through energy consumption, considering both horizontal flight and vertical transitions. Section 5.3 presents the selected routing algorithm. Finally, Section 5.4 discusses the main modelling assumptions and limitations underlying the proposed approach.

5.1. Existing approaches

UAV operations in urban environments raise significant safety and societal concerns due to the potential consequences of operational failures, which may result in fatalities, property damage, or broader societal impacts [44]. However, such risks can be reduced by minimizing UAV exposure to high-risk areas during flight. The study by Hu et al. [21] demonstrates that risk-aware routing can yield safer trajectories compared to purely distance-minimizing paths. This can be achieved by explicitly incorporating risk avoidance objectives into the path planning process [64].

Several studies have applied risk-based approaches to UAV path planning using grid-based representations of the environment. In these models, the operating area is divided into grid cells, each assigned a composite risk score based on factors such as population density, infrastructure sensitivity, and societal impact. The resulting risk map serves as input for a path planning algorithm, which determines low-risk routes through the grid. UAVs are assumed to navigate freely across the grid space without being limited to predefined corridors.

For example, Tang et al. [64] apply a modified A* algorithm to a two-dimensional risk map, formulating path planning as a single-objective optimization problem focused on minimizing cumulative risk along the route. Pang et al. [44] extend this to a three-dimensional grid and propose a custom search algorithm (EDA-FRA*) that identifies globally optimal paths minimizing aggregate third-party risk. Similarly, Zhou et al. [69] use a 3D grid structure and apply a modified A* algorithm that balances third-party risk and flight distance via a weighted objective function, enabling trade-offs between safety and operational efficiency. Hu et al. [21] adopt a unified risk cost map and compare the performance of three algorithms: standard Dijkstra, a modified A*, and an adapted Ant Colony Optimization (ACO) approach, each aiming to identify cost-effective low-risk trajectories. Zhang et al. [68] apply a comparable static grid-based method, estimating cell-level risk based on crash probability, population density, and impact footprint, and subsequently use a modified Particle Swarm Optimization (PSO-BAS) algorithm to determine the safest path.

In contrast, Primatesta, Guglieri, and Rizzo [47] employ a two-stage approach combining both offline and online components. An initial path is generated offline using the RiskA* algorithm on a static risk map, after which real-time adjustments are made during flight using the Borderland algorithm, which modifies the route locally in response to emerging threats. Other studies, such as Kim and Bae [29] and Rudnick-Cohen, Herrmann, and Azarm [53], do not compute optimized paths but rather evaluate and compare predefined corridor alternatives. For example, Rudnick-Cohen, Herrmann, and Azarm [53] employ Monte Carlo simulations to estimate the expected number of fatalities per corridor based on UAV failure rates and population distributions, while Kim and Bae [29] calculate risk scores by combining failure probabilities, population exposure, and impact area for each candidate corridor.

Finally, López et al. [31] focus on minimizing the risk of airborne collisions between UAVs in shared airspace. Their model dynamically updates trajectories in real-time using a decentralized control strategy, without recalculating the entire route. The UAV reacts locally to potential conflicts and adjusts its trajectory while maintaining safety margins relative to other vehicles.

5.1.1. Path planning methodologies

From the studies discussed above, four distinct design choices can be identified in the application of risk-based UAV path planning:

1. **Ground risk vs. airspace conflict risk**

   Studies differ in the type of risk they prioritize. Most focus on minimizing third-party ground risks (TPR), which include the safety of people, property, and societal concerns on the ground. Other studies, such as López et al. [31], instead emphasize airborne conflict avoidance, focusing on collision risks between UAVs operating within shared airspace. As discussed in Chapter 4, the present study focuses exclusively on third-party ground risks. This is motivated by the specific context of UAV freight operations within the Netherlands, where densely populated, cultivated, and urbanized ground environments dominate the operational landscape. While airborne conflict management remains a key component of overall airspace safety, it is typically addressed through U-space services, deconfliction protocols, and regulated airspace structures, and is therefore outside the scope for this framework.

2. **Risk-only objective vs. risk–efficiency trade-off**

   A second distinction concerns the objective function used in path planning. Some studies focus exclusively on minimizing safety risks, selecting the path with the lowest aggregate risk score (e.g., Tang et al. [64], Pang et al. [44]). Others, such as Zhou et al. [69] and Hu et al. [21], balance risk reduction against operational efficiency by combining safety and distance into a weighted objective function. These trade-offs enable more realistic and commercially viable routing solutions, particularly in contexts where efficiency is a key operational driver. Since maintaining cost-competitive and operationally feasible UAV operations remains one of the central challenges for integration, this study adopts a weighted single-objective approach, where safety risk and route efficiency are combined into a composite cost function controlled by a weighting parameter α. By adjusting α, policy-makers can flexibly balance safety and efficiency priorities [67, 63, 4].

3. **Static vs. dynamic planning**

   Another key distinction is whether path planning is performed statically before flight or dynamically updated during flight. Most studies, such as Zhou et al. [69], Tang et al. [64], and Hu et al. [21], apply static pre-flight planning based on predefined risk maps, resulting in fixed routes. In contrast, dynamic approaches as seen in Primatesta, Guglieri, and Rizzo [47] and López et al. [31], continuously adjust UAV trajectories in response to real-time environmental changes, such as emerging threats or proximity to other aircraft. In this study, dynamic path adaptation is not considered, as the objective is to design fixed UAV corridors rather than flexible flight paths. Once these corridors are determined, UAVs operate along them under constant route configurations, with safety assurance embedded directly in the static corridor design rather than in-flight adjustments.

4. **Grid-based free navigation vs. network-based fixed routing**

   Finally, a methodological distinction lies in how the environment is represented. Most studies apply a 2D or 3D grid-based approach, allowing UAVs to navigate freely between grid cells based on local risk costs (e.g., Tang et al. [64], Pang et al. [44]). In contrast, other studies evaluate risk along predefined corridor structures or network paths where UAVs follow fixed, infrastructure-aligned routes (e.g., Kim and Bae [29], Rudnick-Cohen, Herrmann, and Azarm [53], Doole [8]). This study adopts the latter approach, building on Doole [8]. Rather than using grid cells, existing linear infrastructures such as roads, railways, and waterways are used to construct a network of candidate UAV corridors. Each network segment represents a potential corridor section, and the path planning process consists of selecting the sequence of segments that minimizes total risk (and travel cost) between origin and destination nodes. This graph-based formulation is further detailed in Chapter 3.

Studies such as Kim and Bae [29] and Rudnick-Cohen, Herrmann, and Azarm [53] assess corridor safety only after candidate routes have been defined. In contrast to this, this study integrates risk assessment directly into the corridor design process. By assigning spatially explicit third-party risk scores in advance, risk becomes part of the route optimisation itself rather than a later evaluation step. This approach makes it possible to identify corridor configurations from the outset that offer an optimal balance between safety and operational efficiency. In addition, the model explicitly takes into account regulatory constraints, including no-fly zones and airspace classifications. This makes it practically applicable for regulated UAV integration within structured airspace systems such as the Dutch U-space.

5.2. Operational efficiency

To enable competitive UAV freight operations and ensure competitive feasibility alongside traditional delivery systems, routes must not only minimise risk but also maintain sufficient operational efficiency. This requires a balance between safety and efficiency. In this study, path efficiency is defined as the total energy required by a UAV to complete its route. This energy consumption is determined by two primary factors: (i) changes in flight altitude, since ascent and descent require additional energy, and (ii) the total distance travelled horizontally. The following sections describe how both components are taken into account when assessing efficiency.

5.2.1. Recommended operating altitudes

The optimal flight altitude for UAV corridors is not universally fixed, but highly dependent on the surrounding ground environment. Ren and Cheng [48] developed a third-party risk model for urban UAV delivery that explicitly quantifies risks to uninvolved individuals on the ground. Their model incorporates three key dimensions: the likelihood of injury from a crash, noise nuisance caused by UAV propellers, and privacy concerns resulting from onboard cameras. These risks were assessed using a 3D grid-based model applied to an urban case study in China, evaluating third-party risks at altitudes ranging from 30 to 60 meters.

The results show that for most open or unshielded environments, such as lakes, forests, roads and open spaces, the lowest aggregate third-party risk occurs at about 30 metres height. In these areas, lower heights reduce both crash severity and noise exposure through the combination of limited human presence and minimal physical obstacles. Consequently, 30 metres can be considered the preferred operational height for UAV corridors traversing open or sparsely populated areas.

Conversely, in built-up environments, risk levels were found to decrease as flight heights increased. Above buildings, higher flight altitudes reduce the likelihood of collisions between UAVs and people, as occupants are sheltered and structurally protected from falling UAVs. Moreover, flying at higher altitudes reduces privacy issues and reduces noise exposure at ground level. Moreover, flying at higher altitudes reduces privacy issues and reduces ground-level noise exposure. On this basis, Ren and Cheng [48] recommend flight altitudes of approximately 60 meters for corridors operating above densely built-up areas.

Building on these findings and the underlying mechanisms of third-party risk, this study proposes the following height guidelines based on area type:

- 30 meters: Lakes and ponds, Rivers, canals and streams, Water reservoirs, Forests and woodlands, Meadows and open grass, Agricultural lands, Wetlands, Cemeteries, Parks, Recreational zones, Tracks and rural access roads, Pedestrian and cycling paths, Natura2000 areas, Living and residential streets, Regional roads, Motorways and major roads, Railways, Bridges, Cultural sites, Religious sites, Schools and universities
- 60 meters: Residential areas, Commercial zones, Retail zones, Industrial zones, Hospitals, Care homes, Prisons, Power plants, Communication towers, Wind turbines, Power lines, High infrastructures

In summary, lower altitudes (30 meters) are preferred for corridors over open spaces, roads, or areas without substantial shielding, as they minimize crash energy upon impact. For flights above built-up environments or sensitive institutions (e.g., healthcare facilities, schools), higher altitudes (60 meters) are recommended to reduce both safety risks and societal concerns related to noise and privacy.

5.2.2. Energy consumption assumptions

UAV-based parcel delivery is primarily expected to target lightweight consumer goods such as toiletries, medications, documents, clothing, and small electronics. These types of items are commonly ordered through e-commerce platforms like Bol.com and Amazon. Existing commercial UAV services support this profile: Amazon Prime Air reports that its drones can deliver packages up to five pounds (approximately 2.3 kg) within one hour [2], while Zipline's Zips UAVs are designed to carry payloads up to 8 pounds (approximately 3.6 kg) [70]. Trials in the Netherlands have demonstrated that heavier payloads are technically feasible as well; for example, Odido successfully completed UAV delivery tests with a 5 kg payload over a 50-kilometer distance [42]. Nevertheless, these figures represent maximum payload capacities, and actual average delivery weights are expected to be considerably lower in routine operations.

Based on the typical characteristics of small-parcel deliveries and payload specifications of current UAV platforms, this study assumes an average payload weight of 1.0 kg, representative of common e-commerce shipments, medical supplies, and lightweight consumer goods. Following Rodrigues et al. [51] and Rodrigues et al. [52], who report energy consumption values between 0.05 and 0.08 MJ/km (equivalent to 13.9–22.2 Wh/km) for electric quadcopters carrying 0.5–1.0 kg payloads, a representative horizontal energy consumption of 25 Wh/km is adopted for modeling purposes.

In addition to horizontal flight, altitude changes impose substantial additional energy costs. Simulation experiments by Michel et al. [35] indicate that a 30-meter vertical climb requires approximately 2.03 Wh, while descending 30 meters consumes 0.93 Wh. In comparison, forward horizontal flight over 50 meters consumes approximately 5.2 kJ (1.45 Wh), confirming that vertical movements are substantially more energy-intensive per meter than horizontal travel. To capture these effects, the energy model applied in this study adds 2.03 Wh for every 30 meters of ascent and 0.93 Wh for every 30 meters of descent to the baseline horizontal energy consumption of 25 Wh/km. A sensitivity analysis, presented in Section **??**, examines how sensitive the model is to the energy consumption assumptions.

Although Michel et al. [35] also demonstrate that combining vertical and horizontal motion can improve energy efficiency, such diagonal movement is not explicitly modeled in this study. In the context of the present network-based routing approach, altitude adjustments typically occur over short segments where sustained diagonal flight is operationally impractical. The adopted energy model therefore provides a simplified yet representative approximation of UAV energy requirements across varying corridor profiles.

5.3. Path planning algorithm

This study focuses on the design of risk-aware structured UAV corridors within a spatial network. As a basis for this, the path planning algorithm must be able to (i) handle weighted graphs, (ii) incorporate risk and efficiency into routing decisions, and (iii) remain computationally efficient for large-scale applications.

Given these requirements, Dijkstra's algorithm was selected as the primary path planning method. Dijkstra's Shortest Path First (SPF) algorithm is one of the most widely applied methods for determining

the shortest paths in weighted graphs. It follows a greedy approach, selecting the locally optimal solution at each step and iteratively building the globally optimal path [25]. In contrast to A*-based algorithms, which are commonly used in UAV path planning [69, 64, 47], Dijkstra's algorithm offers superior computational performance in the current setting. Preliminary tests showed that A* with Euclidean heuristics resulted in significantly higher computation times without substantial gains in path quality for this corridor-based network structure. Therefore, Dijkstra's algorithm was adopted to ensure scalability and practical runtime for the large spatial networks used in this study. A detailed description of the implementation and computational workflow is provided in Appendix C.

5.3.1. Dijkstra's Algorithm for shortest path

To find the optimal UAV corridors between distribution points and PostNL delivery points, Dijkstra's algorithm is applied to the weighted network graph. Dijkstra's algorithm computes the shortest path by minimizing a cumulative cost across the network. Instead of minimizing distance alone, each edge in the network is assigned a composite cost that integrates both third-party safety risk and energy consumption. First, the cost of each edge is calculated based on the risk-energy trade-off parameter α, after which the optimal route is identified using Dijkstra's algorithm on the resulting weighted graph.

The cost c_{uv} for a given edge (u, v) is defined as:

$$c_{uv} = \alpha \cdot (r_{uv} \cdot \ell_{uv}) + (1 - \alpha) \cdot \left(e_{uv}^{\text{hor}} + e_{uv}^{\text{vert}}\right)$$

where:

- $\alpha \in [0, 1]$ is a weighting factor that balances safety versus energy efficiency;
- r_{uv} is the risk score associated with the edge;
- ℓ_{uv} is the horizontal length of the edge in meters;
- $e_{uv}^{\text{hor}} = \ell_{uv} \cdot \varepsilon$, with $\varepsilon = 0.025$ Wh/m representing the horizontal energy cost;
- e_{uv}^{vert} is a fixed vertical energy penalty depending on ascent or descent.

Since UAV altitude changes are discretized into steps of 30 meters (corresponding to the assumed operational flight levels of 30 and 60 meters) the vertical energy cost is modelled as:

$$e_{uv}^{\text{vert}} = \begin{cases} 2.03 \text{ Wh} & \text{if } \Delta h_{uv} = +30 \text{ m (ascent)} \\ 0.93 \text{ Wh} & \text{if } \Delta h_{uv} = -30 \text{ m (descent)} \\ 0 & \text{if } \Delta h_{uv} = 0 \end{cases}$$

Here, $\Delta h_{uv} = h_{uv}^{\text{edge}} - h_u^{\text{node}}$ denotes the difference between the edge's target altitude and the source node's altitude. This fixed-step approach simplifies the energy estimation while still capturing the asymmetry in energy use between ascent and descent. The resulting cost c_{uv} is used as the edge weight in the shortest path computation using `networkx.shortest_path()` with `weight='cost'`. The complete routing workflow, including distribution center assignment and performance metric calculation, is detailed in Appendix C.

5.4. Constraints and Considerations

When computing the routes and interpreting the resulting paths, several constraints and methodological assumptions underlying this framework must be acknowledged:

- Altitude approximations: The UAV flight altitudes are discretized into two general levels (30 m and 60 m) based on broad operational guidelines, rather than on detailed airspace restrictions or fine-grained area-specific studies.
- Connector edge altitude: All connector edges are assigned a fixed altitude of 30 meters. These connector segments serve to ensure network connectivity between natural corridors and linear infrastructures, enabling a fully connected routing graph.
- Turning energy not modeled: The additional energy consumption associated with turning maneu-

vers is not incorporated. As shown by Jacewicz et al. [23], UAVs experience increased energy demand during turns, but this effect is omitted in this study for modeling simplicity.

5.5. Conclusion of path planning methodology

Several design choices must be made when developing UAV corridor planning models. In this study, the methodology is focused exclusively on third-party ground risk, as airspace conflict risks are addressed through separate air traffic management mechanisms and U-space services. A trade-off between safety risk and operational efficiency is incorporated to ensure that resulting corridor designs remain practically feasible for commercial UAV operations. Static pre-flight planning is applied rather than dynamic in-flight adjustments, as the objective is to develop fixed corridor structures that can be pre-approved, regulated, and integrated into national airspace systems. Consistent with this objective, a network-based fixed routing approach is adopted, in which predefined infrastructure-based edges serve as the candidate corridor network.

Operational efficiency is modelled through energy consumption, simplified into horizontal and vertical components. Based on typical payload profiles for small parcel delivery, an average payload of 1.0 kg is assumed, with a horizontal energy consumption of 25 Wh/km. Vertical movements impose additional fixed energy costs of 2.03 Wh for every 30 meters of ascent and 0.93 Wh for every 30 meters of descent. A sensitivity analysis, presented in Section 6.5.2, further evaluates how variations in these energy parameters affect routing outcomes.

To perform risk-based path planning, Dijkstra's algorithm is applied to the network, using a composite cost function that balances third-party risk exposure and energy consumption. The trade-off between these objectives is controlled through the weighting parameter α, which allows policymakers to flexibly prioritize safety or efficiency depending on the regulatory or operational context. The practical implementation of this algorithm is described in Appendix C. In the following chapter, this framework will be applied in a series of experiments to explore corridor design solutions across rural, urban, and suburban environments.

6

Experiments

This chapter presents a comprehensive experimental evaluation of the proposed UAV corridor planning framework across diverse geographic contexts in the Netherlands. Having established the theoretical foundation for infrastructure-based corridor design (Chapter 3), risk assessment methodology (Chapter 4), and integrated path planning approach (Chapter 5), this chapter examines whether the framework can generate practically viable corridors that balance safety and efficiency under real-world constraints.

This chapter is structured as follows: Section 6.1 introduces the three case study regions and justifies their selection based on competitive feasibility criteria for UAV freight transport. Section 6.2 examines how regulatory constraints on overflight of populated areas motivate the network-based approach. Section 6.3 presents the comparative analysis across geographic contexts, including detailed examination of risk reduction patterns, infrastructure utilization shifts, and multi-depot dynamics. Section 6.3.2 analyzes the Pareto-optimal trade-offs between risk and efficiency across all three regions. Section 6.4 investigates how battery capacity constraints affect service coverage under different routing strategies. Section 6.5 presents comprehensive sensitivity analyses for risk parameters (severity and likelihood scores) and energy consumption parameters. Finally, Section 6.6 synthesizes these findings to conclude on the framework's practical viability and readiness for real-world implementation.

6.1. Case studies applied

To ensure competitive feasibility, UAV-based freight transport must meet several key operational conditions: fast delivery capability, low route density, suitability for short to medium distances, and a high utilization rate, as discussed in Section 2.2. To evaluate these criteria, the selected use case involves parcel delivery from PostNL distribution centres to designated PostNL pick-up points. This scenario aligns closely with the identified feasibility requirements for the following reasons:

- **Fast delivery:** Although not inherently guaranteed, fast delivery can be achieved through the prioritization of express parcel streams.
- **Low route density:** PostNL pickup points are usually deliberately placed to optimize area coverage, resulting in a naturally low-density delivery network.
- **Short to medium distances:** PostNL has 39 distribution centres across the Netherlands. While this network provides a relatively solid foundation for the regional deployment of UAVs, further optimisation is possible. To ensure a meaningful evaluation, the use-case areas were therefore strategically selected to provide realistic delivery scenarios in regions adjacent to a distribution centre.
- **High utilization rate:** Pickup points experience consistent parcel flows and increasing consumer adoption, supporting a high potential for UAV deployment efficiency.

To capture variation in geography, infrastructure, and ground-based risk exposure, three representative regions were selected: one urban (Breda), one rural (Borsele), and one suburban (Alphen aan den Rijn

and Waddinxveen). The suburban case is of particular interest, as it includes two distribution centres within operational range. This configuration enables an assessment of how corridor behaviour changes when multiple origin points are available.

The selection of these study areas was guided by the principle of minimizing confounding airspace complexity factors. Specifically, regions were chosen to avoid proximity to complex airspace structures such as CTRs associated with major airports (e.g., Schiphol and Rotterdam Airport) and military training areas. As discussed in Section 3.2.2 and Appendix B.1, UAV operations within such complex airspace environments can remain technically feasible through proper coordination and authorisation procedures. However, the emerging state of policy-related research on UAV corridor planning requires initial validation in less constrained operational environments. This methodological approach makes it possible to establish a proof-of-concept for the proposed risk-aware routing framework before addressing the additional layers of complexity inherent in highly regulated airspace. By demonstrating the effectiveness of the approach in these representative but operationally simple environments, this research provides a basis for future extensions to more complex urban airspace scenarios.

Table 6.1 summarises how each of these regions aligns with the competitive feasibility criteria for UAV-based freight transport.

Table 6.1: Assessment of competitive feasibility criteria across case study areas

Feasibility Criteria	Breda (Urban)	Borsele (Rural)	Alphen–Waddinxveen (Suburban)
Fast delivery requirements	Enabled through prioritisation of express parcels within a dense urban setting	Express delivery feasible despite longer rural distances due to low congestion	Achievable via prioritisation of express parcels across a well-connected suburban network
Low route density	64 pickup points identified in OSM, evenly spread across the city, supporting efficient routing	23 pickup points identified over a large area, reflecting naturally low route density	81 pickup points found across three municipalities, balanced between two distribution centres
Short–medium distances	Distribution centre located centrally within the city limits	Service area located directly adjacent to the Goes distribution centre	Two distribution centres provide localised coverage for each town and the area in between
High utilisation	High population density ensures stable and frequent parcel flows	Lower population density, but rising rural e-commerce adoption supports consistent volumes	Two growing municipalities ensure strong demand; intermediate areas may see moderate use

Urban Case Study: Breda

The urban scenario focuses on the municipality of Breda. As of 2025, Breda has a population of 188,779 inhabitants [5]. The city features a dense urban layout and substantial physical infrastructure, including extensive waterways and multiple intersections of rail, road, and inland shipping networks. Additionally, a major PostNL distribution center is located next to the city center, ensuring that delivery routes to nearby pickup points remain within the targeted short-to-medium range. The built-up character of Breda ensures that the scenario captures typical urban airspace constraints and ground-level complexity. An overview of the Breda region, including area types, infrastructure features, no-fly zones, and the locations of PostNL distribution points, is provided in Figure 6.1.

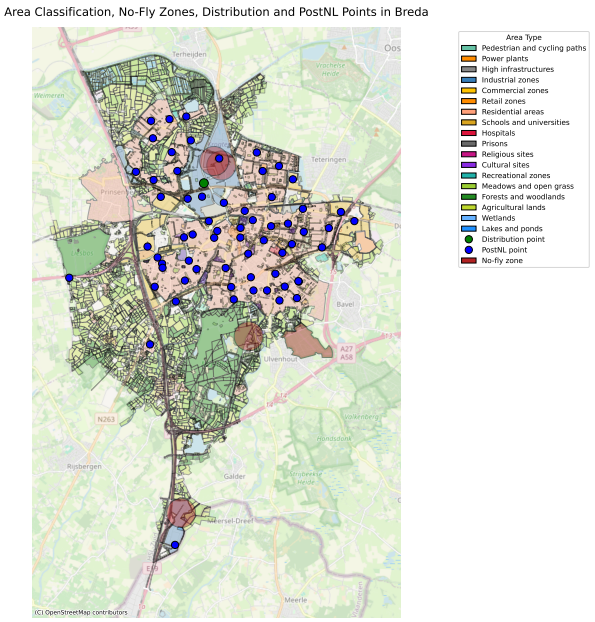

Figure 6.1: Overview of the Breda region.

Rural Case Study: Borsele

The rural scenario centers on the municipality of Borsele, located in the province of Zeeland. As of 2025, Borsele has a population of 23,125 inhabitants [5]. The area is characterized by a sparsely populated, low-density environment and lies adjacent to a major distribution hub in the nearby city of Goes. As illustrated in Figure 6.2, the region consists primarily of open green areas, interspersed with water bodies and limited built-up structures. However, the operational environment also contains critical infrastructure elements that constitute high-risk zones for UAV operations, including a railway line and an extensive power transmission network. This spatial composition, combining open terrain with concentrated linear hazards, differs substantially from the urban complexity observed in Breda and provides a representative rural environment to evaluate the effectiveness of risk-aware UAV corridor planning.

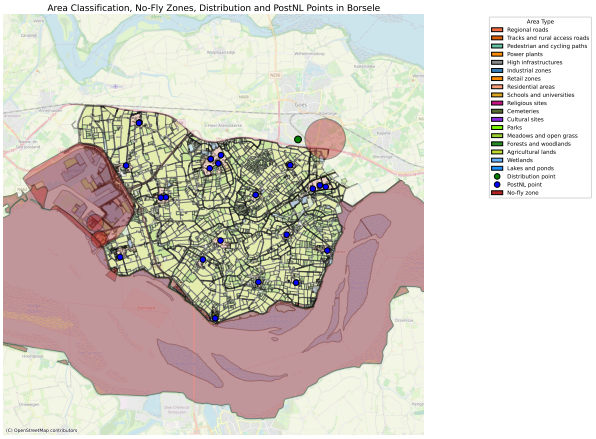

Figure 6.2: Overview of the Borsele region.

Suburban Case Study: Alphen aan den Rijn and Waddinxveen

The suburban scenario focuses on the area encompassing Alphen aan den Rijn and Waddinxveen. These two towns represent mid-sized urban environments separated by semi-rural landscapes, offering a representative mix of suburban conditions. The region contains two PostNL distribution centers, providing relevant logistical nodes for UAV operations. The two municipalities are connected by a railway corridor, which presents both a physical barrier and elevated risk zone for UAV operations. Additionally, the region is characterized by an extensive network of waterways, typical of the Dutch polder landscape, which offer potential low-risk flight corridors but may also constrain routing flexibility. This mixed spatial composition and moderate infrastructure density create a complex operational environment that allows for a nuanced assessment of UAV corridor planning under hybrid urban–rural conditions. The presence of dual distribution centers further enables the analysis of multi-origin routing strategies, where both operational efficiency and spatial ground risk profiles must be jointly considered in the optimization framework.

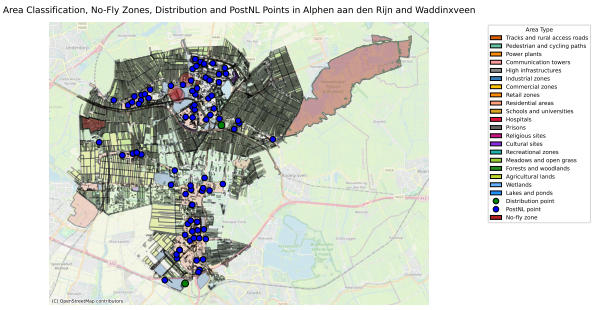

Figure 6.3: Overview of the Alphen aan den Rijn and Waddinxveen region.

6.2. Residential and other high-density areas as hard constraints in UAV routing

In UAV corridor design, the presence of people on the ground is one of the most critical safety constraints. Independent of the applicable airspace classification or UAV geozone, operations are generally prohibited over areas where uninvolved persons may be directly exposed to UAV overflight. Certain locations are particularly sensitive due to the frequent or unpredictable presence of crowds. Examples include residential neighborhoods, shopping streets during business hours, parks and beaches in favorable weather, sporting events, cultural or religious gatherings, political demonstrations, ski resorts, and other locations with high pedestrian or vehicular density [12].

Regulatory frameworks establish specific operational limitations based on UAV classification. Light drones (<250g or Class 0/1) operating in subcategory A1 may overfly uninvolved persons under strict constraints, while heavier UAVs (C3/C4) must maintain minimum distances of 150 meters from urban areas and 30 meters from crowds. The fundamental operational principle requires maintaining a lateral distance from uninvolved persons equal to or greater than the flight altitude (the ”1:1 rule”), with an absolute minimum of 30 meters except in designated low-speed modes. These regulations ensure that UAV operations minimize risk exposure to ground populations while enabling viable commercial applications [12].

In practical implementation of these safety regulations, certain area categories can be treated as hard no-fly constraints within the corridor design model. Specifically, OpenStreetMap classifications such as residential zones, retail zones, recreational areas, cultural sites, and religious sites would be excluded from feasible routing. This results in substantial restrictions for all named regions, as illustrated in Figures E.1, E.2, and E.3. Consequently, the application of these constraints reveals that a significant proportion of PostNL pickup points are located within no-fly zones: 78.46% for Breda, 95.83% for Borsele, and 81.93% for Alphen aan den Rijn and Waddinxveen combined.

These results demonstrate that strict enforcement of safety regulations would make UAV delivery operationally unfeasible, with accessibility limited to less than 22% of delivery destinations in all study areas. However, this strict classification approach does not take into account the heterogeneous risk profiles within these limited zones. Infrastructural elements such as waterways or green corridors crossing residential areas may exhibit significantly lower risk exposure than the surrounding built environment, despite their formal classification in a restricted area. This observation motivates the network-based routing approach developed in this study. As illustrated in Figure 6.4, this methodology allows the identification of lower-risk corridors within otherwise constrained areas. This increases operational feasibility while maintaining safety standards.

Building upon this methodological framework, the following section presents a comprehensive analysis of UAV corridor design across the three study areas, examining how the network-based risk assessment approach performs under varying spatial and infrastructural conditions.

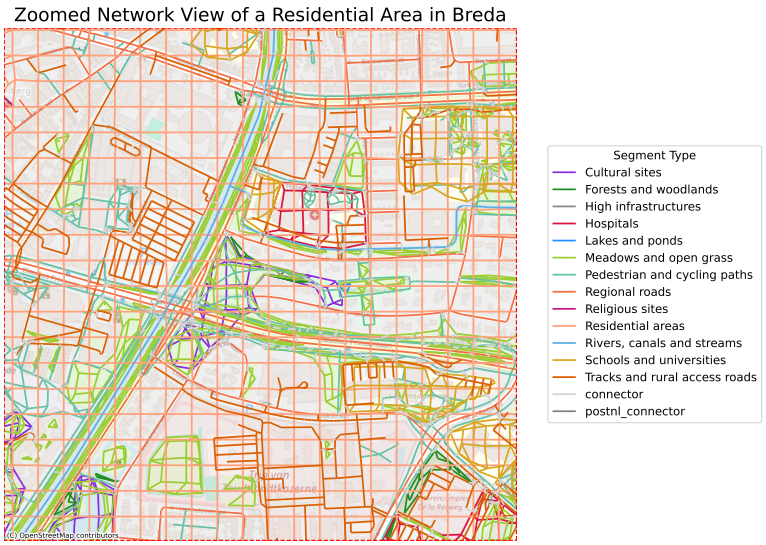

Figure 6.4: Local window example illustrating potential routing opportunities within residential areas.

6.3. Experimental evaluation of proposed design across three regions

Having established a design framework that meets regulatory, spatial, and operational requirements, this chapter examines whether such corridor designs remain practically feasible when balancing safety and efficiency objectives. To assess the feasibility of the proposed framework, it is applied to three distinct regions: an urban area (Breda), a rural area (Borsele), and a suburban area (Alphen aan den Rijn and Waddinxveen). For each region, a range of α values is tested in the single-objective optimization function, where α governs the trade-off between risk aversion and route efficiency. This enables an evaluation of how corridor designs vary depending on whether operational efficiency (α = 0) or risk minimization (α = 1) is prioritized, and whether the framework can generate both efficient and risk-averse routes across diverse geographic contexts.

6.3.1. Comparative analysis across geographic contexts

The three case studies reveal distinct patterns in how geographic context influences the trade-off between operational efficiency and risk mitigation. While all regions exhibit the expected inverse relationship between risk and efficiency metrics, the magnitude and nature of these trade-offs vary significantly, as demonstrated in Tables 6.2, 6.3, and 6.4. Further details on the minimum and maximum of all metrics and full edge selection across all α values can be found in Appendix E.

Risk reduction patterns

All three regions experience substantial risk reductions between $\alpha = 0$ and $\alpha = 0.1$, with diminishing returns thereafter. This finding indicates that significant risk mitigation can be achieved with only modest sacrifices in operational efficiency. The relative impact varies across contexts, with Alphen–Waddinxveen achieving the largest reduction (77%), while Breda (71%) and Borsele (70%) demonstrate similar patterns. These results suggest that the proposed framework can successfully generate low-risk yet reasonably efficient corridor networks across all tested geographic contexts.

The topological differences between efficiency-based networks (Figures 6.5a, 6.6a, and 6.7a) and risk-averse configurations (Figures 6.5b, 6.6b, and 6.7b) clearly illustrate how risk considerations fundamentally reshape corridor structures. Efficiency-based path planning generates predominantly direct trajectories between origin and destination points, prioritizing minimal distance and energy consumption. In contrast, risk-averse routing produces notably circuitous paths characterized by systematic deviations around high-risk zones, accepting increased path length as a trade-off for enhanced ground safety.

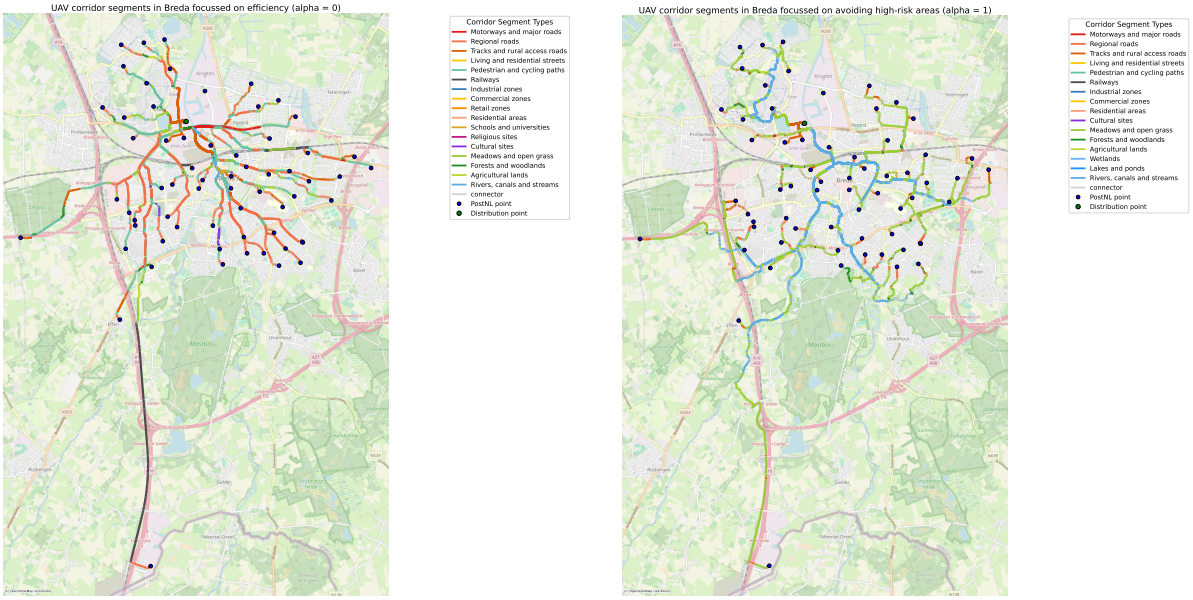

(a) Efficiency-focused planning ($\alpha = 0$) **(b)** Risk-averse planning ($\alpha = 1$)

Figure 6.5: Comparison of corridor networks in Breda under efficiency-oriented and risk-averse planning.

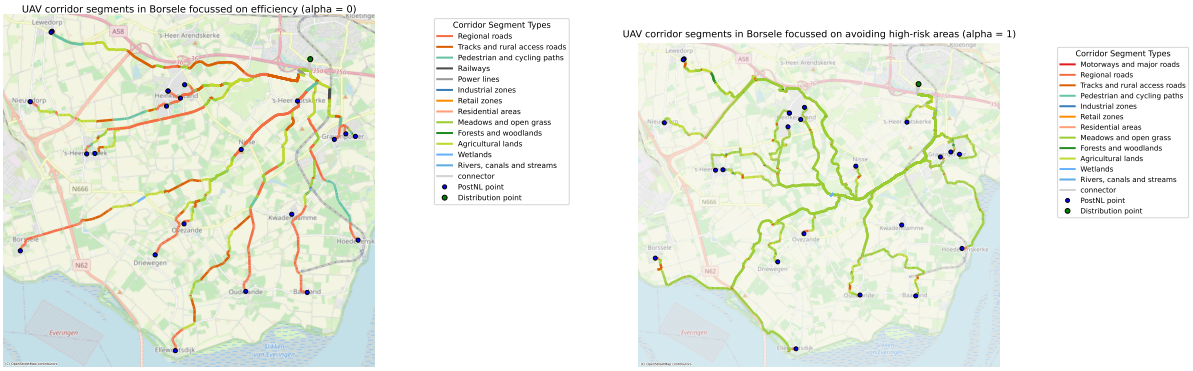

(a) Efficiency-focused planning ($\alpha = 0$) **(b)** Risk-averse planning ($\alpha = 1$)

Figure 6.6: Comparison of corridor networks in Borsele under efficiency-oriented and risk-averse planning.

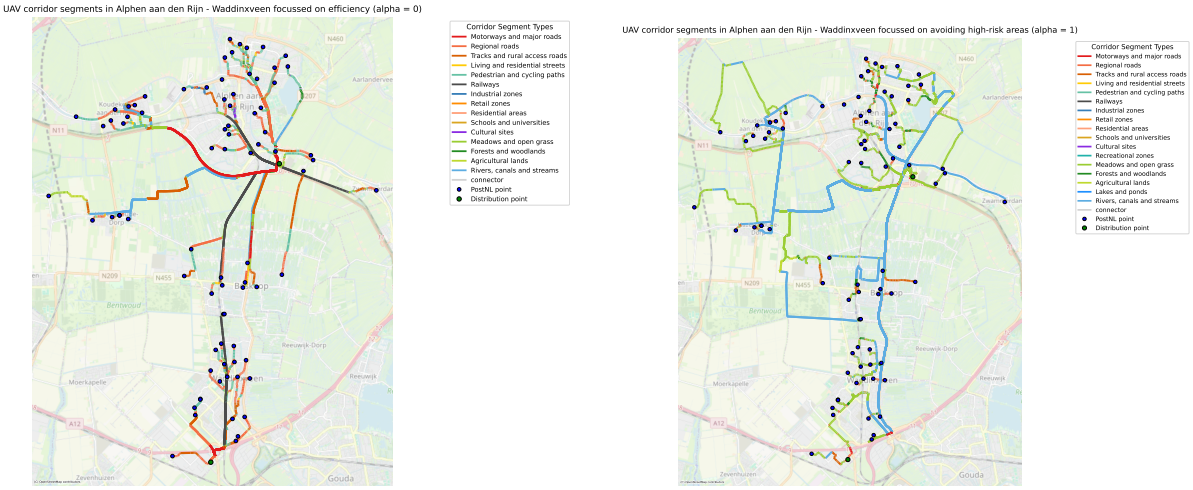

(a) Efficiency-focused planning ($\alpha = 0$)

(b) Risk-averse planning ($\alpha = 1$)

Figure 6.7: Comparison of corridor networks in Alphen aan den Rijn and Waddinxveen under efficiency-oriented and risk-averse planning.

Efficiency penalties and structural adaptations

The cost of risk reduction varies substantially by geographic context:

- **Urban (Breda):** Table 6.2 reveals moderate impacts across all metrics when transitioning from efficiency-based to risk-averse planning: energy consumption increases by 49%, path length by 53%, and turn frequency by 31%. Vertical displacement remains relatively stable, increasing marginally from 5.86 to 6.46 height changes per route. This demonstrates that the urban environment's dense network of alternative corridors enables effective risk mitigation without requiring extreme detours or significant altitude adjustments, as illustrated in Figures 6.5a and 6.5b. Notably, turn frequency initially decreases when transitioning from $\alpha = 0$ to $\alpha = 0.1$, suggesting that preliminary risk considerations may actually streamline routing by eliminating marginally efficient but structurally complex paths.
- **Rural (Borsele):** Table 6.3 demonstrates severe structural adaptations, with path length and energy consumption both increasing by approximately 61% and turn frequency nearly doubling (92%). As illustrated in Figures 6.6a and 6.6b, efficiency-based routing exploits linear infrastructure such as railways, power line corridors and motorways while risk-averse planning systematically avoids these high-risk elements. This avoidance necessitates circuitous navigation through the sparse rural network, explaining the disproportionate increase in path complexity. Despite these horizontal deviations, vertical displacement remains relatively constant (4.9 to 5.29 height changes), indicating that altitude variations are primarily determined by terrain characteristics rather than risk considerations.
- **Suburban (Alphen–Waddinxveen):** Table 6.4 presents an intriguing paradox: despite substantial increases in energy consumption (86%) and path length (88%), topological complexity increases only marginally (22% turn increase). This apparent contradiction can be attributed to the infrastructure substitution pattern observed in the region. As illustrated in Figures 6.7a and 6.7b, risk-averse routing replaces direct railway corridors with extensive waterway networks comprising rivers, canals, and streams. These waterway corridors provide continuous low-risk alternatives that, while longer in absolute distance, maintain directional consistency without requiring frequent course corrections. Vertical displacement metrics remain essentially unchanged, further supporting the hypothesis that waterway-based routing offers topologically efficient risk mitigation.

Table 6.2: Mean values of key performance metrics for UAV corridors in Breda across different α values.

α	Risk	Energy (Wh)	Length (m)	Height Changes	Turns
0.00	1372.6	113.6	4172.7	5.86	123.5
0.10	393.9	130.4	4840.1	6.02	112.1
0.20	362.3	135.8	5042.6	6.21	116.8
0.30	339.2	143.6	5352.1	6.21	123.0
0.40	333.5	146.4	5466.4	6.24	125.5
0.50	329.0	150.2	5609.3	6.27	129.1
0.60	327.1	152.2	5690.3	6.27	133.6
0.70	324.6	157.2	5885.3	6.43	140.1
0.80	323.3	160.5	6016.3	6.43	146.8
0.90	322.1	167.9	6315.2	6.43	157.1
1.00	322.1	169.7	6385.2	6.46	161.4

Table 6.3: Mean values of key performance metrics for UAV corridors in Borsele across different α values.

α	Risk	Energy (Wh)	Length (m)	Height Changes	Turns
0.00	1872.7	222.7	8748.8	4.90	158.1
0.10	564.3	243.7	9559.1	5.38	209.1
0.20	430.7	266.3	10477.8	5.00	228.1
0.30	385.1	280.5	11053.4	4.62	242.4
0.40	377.5	284.9	11231.0	4.62	248.1
0.50	358.3	301.1	11851.5	5.00	260.1
0.60	341.6	320.9	12645.0	5.10	271.5
0.70	340.8	322.3	12701.1	5.10	273.0
0.80	337.5	332.7	13109.5	5.29	280.8
0.90	332.4	354.2	13971.9	5.29	299.8
1.00	332.2	357.5	14103.3	5.29	304.1

Table 6.4: Mean values of key performance metrics for UAV corridors in Alphen–Waddinxveen across different α values.

α	Risk	Energy (Wh)	Length (m)	Height Changes	Turns
0.00	1507.5	117.0	4524.2	4.48	102.2
0.10	352.4	129.3	5004.5	4.43	109.3
0.20	290.4	140.2	5434.7	4.51	102.0
0.30	262.7	149.6	5799.4	4.48	100.6
0.40	247.2	157.1	6088.9	4.51	103.3
0.50	240.5	162.6	6302.4	4.53	102.7
0.60	239.8	163.5	6335.6	4.58	103.4
0.70	233.9	174.1	6765.4	4.48	106.0
0.80	231.5	181.6	7062.2	4.48	108.9
0.90	228.5	200.3	7807.4	4.61	114.5
1.00	227.3	218.2	8524.6	4.56	124.6

Infrastructure utilization shifts

Each geographic context exhibits characteristic infrastructure preferences under risk-averse planning, particularly evident in the transition from $\alpha = 0$ to $\alpha = 0.1$:

- **Breda**: Figure 6.8 demonstrates that efficiency-based planning heavily utilizes pedestrian/cycling paths and regional roads. The introduction of risk considerations triggers a dramatic modal shift toward the city's extensive waterway network and open green spaces, achieving substantial risk reduction through systematic infrastructure substitution. Tracks and rural roads maintain consistent utilization across both planning paradigms, suggesting their balanced risk-efficiency profiles in the urban context.
- **Borsele**: Figure 6.9 reveals that risk-averse planning in rural environments relies predominantly on agricultural corridors and open grasslands, systematically avoiding regional roads and pedestrian/cycling infrastructure that feature prominently in efficiency-based routing. While aggregate statistics indicate minimal railway and motorway usage even under $\alpha = 0$, their strategic exclusion from risk-averse routes significantly impacts network topology at the individual path level, as evidenced by the structural differences between Figures 6.6a and 6.6b.
- **Alphen–Waddinxveen**: Figure 6.10 illustrates a pronounced reduction in high-risk corridor utilization, encompassing motorways, major roads, pedestrian/cycling paths, railways, regional roads, and residential streets. The suburban context compensates for these restrictions by leveraging

natural and semi-natural infrastructure, including meadows, open grasslands, waterway networks (rivers, canals, and streams), and forested areas. This substitution pattern demonstrates how suburban environments can exploit their heterogeneous land-use characteristics to maintain network connectivity while minimizing ground risk exposure.

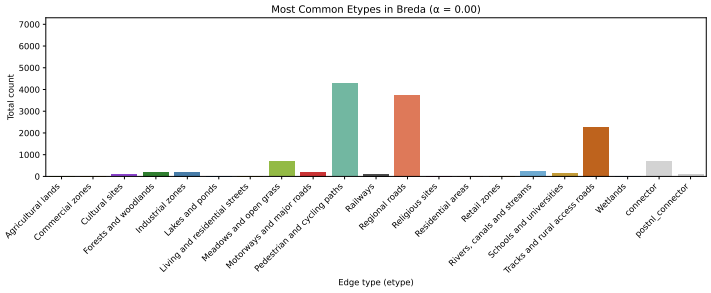

(a) $\alpha = 0.00$

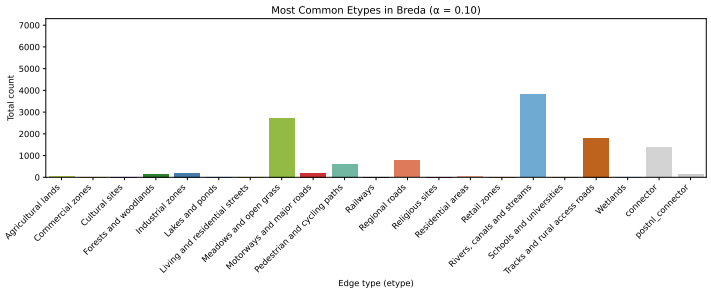

(b) $\alpha = 0.10$

Figure 6.8: Corridor segment usage by edge type in Breda for $\alpha = 0.00 - 0.10$

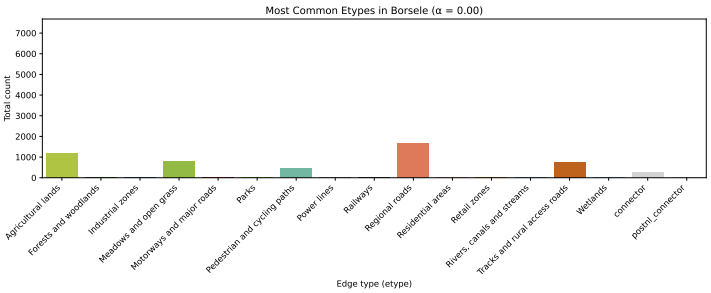

(a) $\alpha = 0.00$

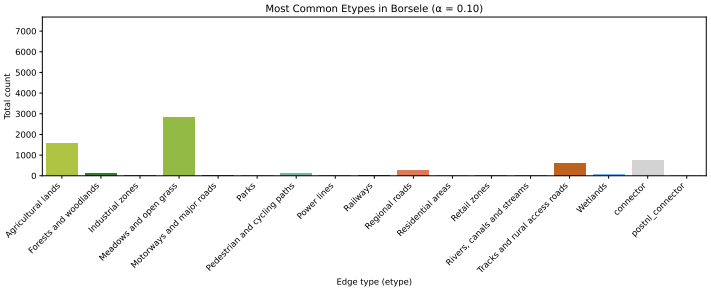

(b) $\alpha = 0.10$

Figure 6.9: Corridor segment usage by edge type in Borsele for $\alpha = 0.00 - 0.10$

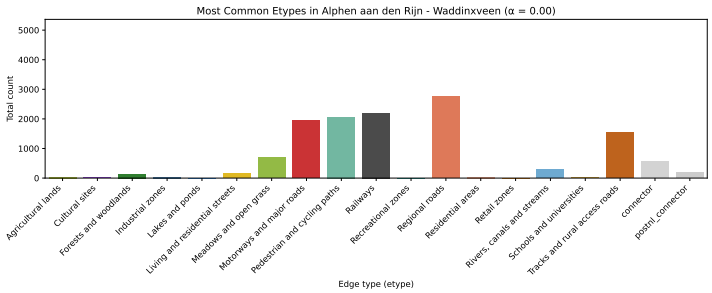

(a) $\alpha = 0.00$

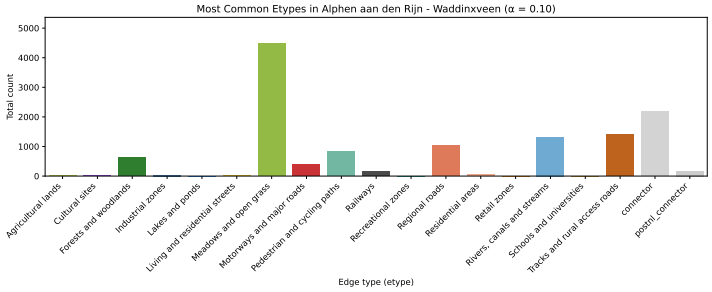

(b) $\alpha = 0.10$

Figure 6.10: Corridor segment usage by edge type in Alphen aan den Rijn and Waddinxveen for $\alpha = 0.00 - 0.10$

Validation of polygon-based area integration

The successful implementation of the polygon integration methodology described in Section 3.2.1 is demonstrated through detailed examination of the Borsele corridor network. Figure 6.11 illustrates the model's capability to plan UAV paths that traverse grid-based polygons, enabling routing both between different polygon areas and from polygons to adjacent linear infrastructure. The visualization shows how the 50-metre grid structure within polygons integrates with the broader corridor network through boundary node connections.

This zoomed-in view of the Borsele region confirms that the multi-stage integration approach successfully generates the 50-metre grid structure within polygons and establishes the necessary connections to enable seamless routing between polygons and to adjacent linear infrastructure. The practical implications of this integration are clear in the rural context, where farmland and open grassland are the dominant corridor type in risk-averse planning (see Figure 6.9).

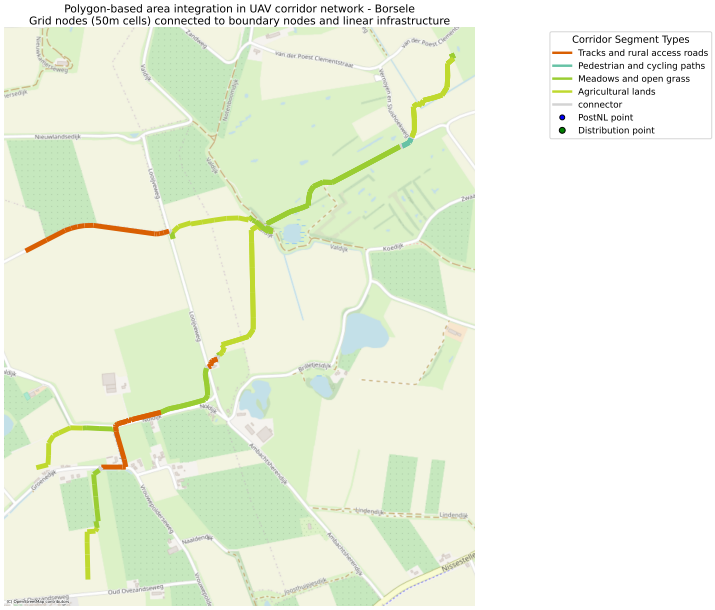

Figure 6.11: Integration of polygon-based areas in the UAV corridor network near Borsele, showing grid-based internal structures connected to linear infrastructure through boundary nodes.

Distribution centre selection

The Alphen aan den Rijn–Waddinxveen case presents an additional dimension of complexity through its dual distribution centre configuration, where risk-efficiency trade-offs influence not only path selection but also origin point allocation. Unlike the single-origin scenarios of Breda and Borsele, route optimization in this multi-origin setting must simultaneously determine both the optimal distribution centre and the subsequent delivery path. As illustrated in Figures 6.12a and 6.12b, efficiency-based routing produces a relatively balanced spatial allocation of delivery points between the two centres, with assignment primarily governed by geographic proximity. Conversely, risk-averse planning generates a highly asymmetric distribution pattern, with the overwhelming majority of delivery points serviced by a single centre in Alphen aan den Rijn. This pronounced spatial bias indicates the presence of significant risk zones that form a barrier between the two distribution centres, making one facility largely inaccessible to risk-aware operations. The phenomenon highlights how security considerations can fundamentally restructure logistics networks.

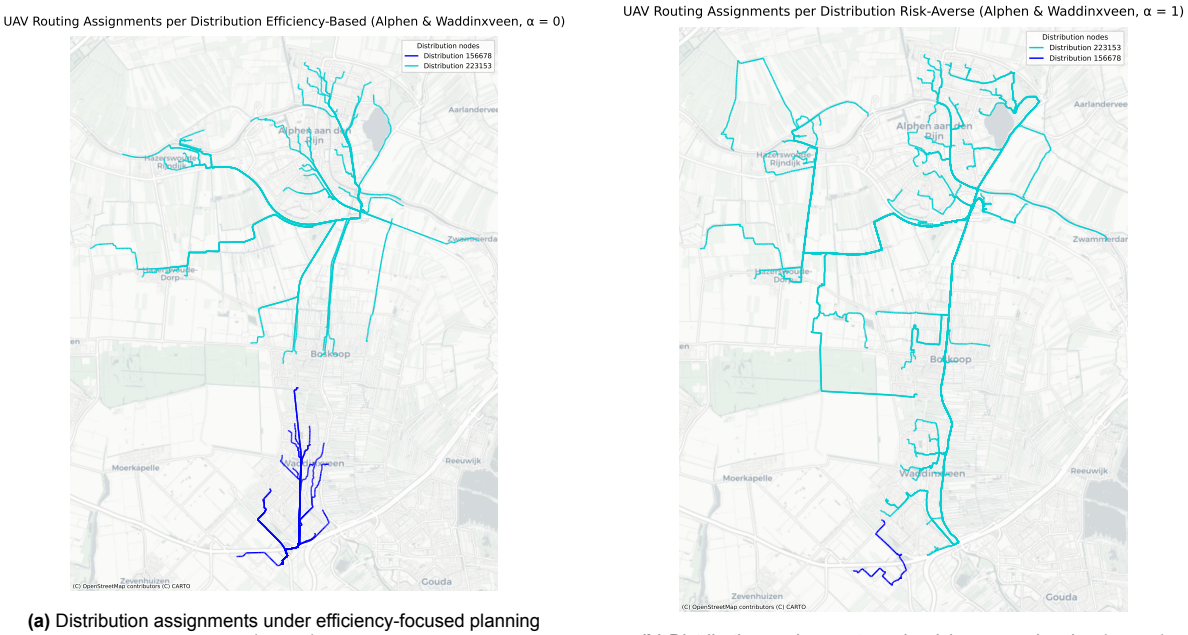

(a) Distribution assignments under efficiency-focused planning ($\alpha = 0$)

(b) Distribution assignments under risk-averse planning ($\alpha = 1$)

Figure 6.12: Distribution center assignment patterns in Alphen aan den Rijn and Waddinxveen under different trade-off settings.

6.3.2. Pareto analysis of risk-efficiency trade-offs

The primary objective of the proposed framework is to determine whether fixed UAV corridors are feasibly designed within the spatial, operational and regulatory constraints of Dutch airspace, while balancing ground risk and energy efficiency. Central to this objective is the identification of acceptable trade-offs between risk reduction and efficiency loss, using a weighting parameter α. Although the framework implements a single-objective optimization using a weighted sum of risk and energy, the set of solutions obtained with different α values can be interpreted as an empirical approximation of a Pareto frontier. Although this set is not strictly Pareto optimal in the multi-objective sense, it illustrates how risk and energy interact for different policy priorities. The resulting trade-off curves for the three case study regions are shown in Figure 6.13.

In all cases, the curves show a clear “hockey-stick” pattern: significant risk reductions can be achieved with relatively modest increases in energy consumption at low values of the α, while decreasing returns occur at higher levels of the α. For example, in the urban case of Breda, average risk reduces with 71% between $\alpha = 0.0$ and $\alpha = 0.1$, while energy consumption rises modestly from 114 Wh to 130 Wh (14%). This shows that a relatively small concession in efficiency yields significant safety gains. Beyond $\alpha = 0.3$, however, further increases in α produce only marginal improvements: risk decreases by merely 5.0% from $\alpha = 0.3$ to $\alpha = 1.0$.

A comparable dynamic is observed in Alphen–Waddinxveen, where risk decreases with 77% between $\alpha = 0.0$ and $\alpha = 0.1$, with energy rising only 10.5%. From $\alpha = 0.4$ to $\alpha = 1.0$, the additional risk reduction amounts to only 8.1%, confirming the presence of diminishing returns. Finally, the rural case of Borsele shows a similar pattern. Here, risk decreases from approximately 70% as α increases from 0.0 to 0.1, while energy consumption rises with 9%. Between $\alpha = 0.5$ and $\alpha = 1.0$, risk declines by only 7.3%, despite a further energy increase of over 50 Wh (18% increase).

These findings confirm that the trade-off structure is consistent in urban, suburban and rural contexts, although the inflection points at which diminishing capital gains occur vary slightly. In the urban context of Breda, a value around $\alpha = 0.10$ seems to offer an efficient compromise. In suburban or rural areas, slightly higher values, closer to $\alpha = 0.20$, may be more appropriate (Figure 6.13). Ultimately, the preferred option remains a policy decision that reflects the acceptable balance between risk exposure and operational efficiency in a given deployment scenario.

A detailed description of the results can be found in Appendix E.

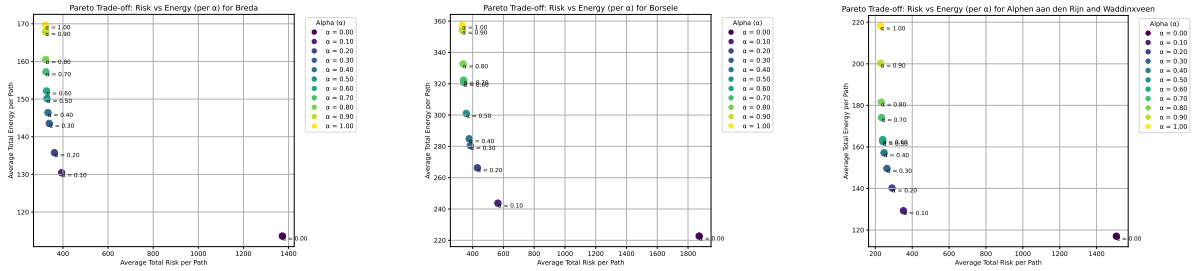

(a) Breda **(b)** Borsele **(c)** Alphen-Waddinxveen

Figure 6.13: Pareto fronts showing trade-offs between energy consumption and risk across the three study areas

6.4. Energy consumption constraints

The findings from the previous section show that within the proposed model only marginal reductions in efficiency are required to achieve substantial risk reduction. However, it is essential to recognize that UAVs have a finite operational range due to limited battery capacity. Therefore, it is also necessary to understand the feasibility of different delivery paths under different battery capacity scenarios.

To address these operational constraints, reachability of delivery destinations is analyzed as a function of battery capacity for two different operational areas: Borsele and Breda, each served by a single distribution center. Included in this analysis is the constraint that UAVs must maintain sufficient energy reserves for the return trip, limiting outbound energy consumption to 50% of total capacity. This approach provides critical insights into the practical implementation of risk-aware route planning within real-world energy limits.

6.4.1. Comparative analysis of urban and rural settings

Figures 6.14 and 6.15 illustrate the network coverage achievable under different battery capacities and routing strategies for both urban and rural study areas. The analysis reveals significant disparities in network coverage between the two environments.

In Borsele, which represents a larger, more dispersed rural setting, efficiency-based routing ($\alpha = 0$) achieves only 39.1% coverage with a 400 Wh battery capacity. Conversely, Breda's more compact urban environment demonstrates near-complete coverage at the same capacity level, with only a single outlying location requiring 800 Wh for accessibility. Under full risk-weighted routing ($\alpha = 1$), Borsele's coverage drops to 17.4% at 400 Wh, and even with maximum battery capacity (900 Wh) achieves only 69.6% coverage. In contrast, Breda demonstrates remarkable resilience under risk-averse planning, maintaining 68.8% coverage at 400 Wh and achieving near-complete coverage at 900 Wh, excluding only locations with no network connectivity.

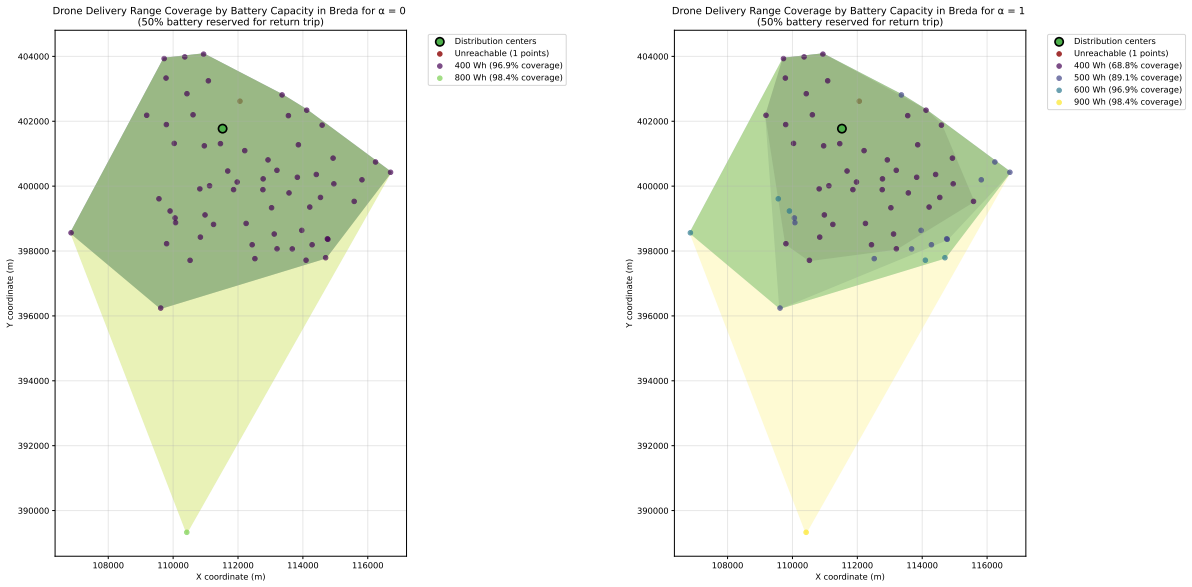

(a) First caption **(b)** Third caption

Figure 6.14: Overall caption for all three subfigures

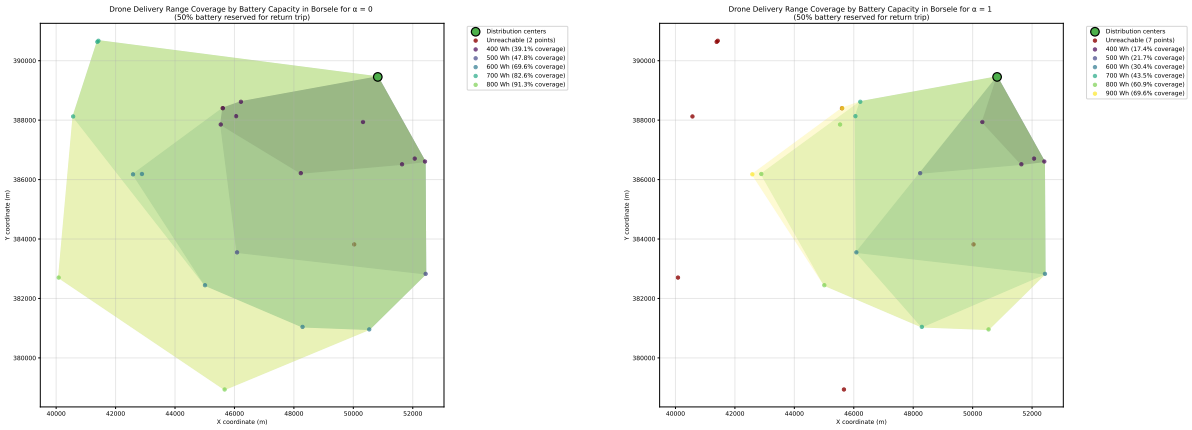

(a) First caption **(b)** Third caption

Figure 6.15: Overall caption for all three subfigures

6.5. Sensitivity analysis on uncertain parameters

This section presents detailed sensitivity analyses examining the robustness of the UAV corridor planning framework to parameter uncertainty. The analyses focus on two critical components: the semi-quantitative risk assessment parameters (severity and likelihood scores, described in Chapter 4) and the energy consumption assumptions (described in Chapter 5). All sensitivity experiments were conducted using the Breda urban case study, with parameters systematically perturbed to evaluate their impact on corridor selection and performance metrics. A detailed description of the methods and results can be found in Appendix F.

6.5.1. Sensitivity analysis on semi-quantitative risk assessment parameters

To evaluate the robustness of the corridor optimization to risk parameter uncertainty, a series of perturbation experiments was conducted focusing on the urban case (Breda). The analysis targeted the area types most frequently traversed in the computed UAV corridors: meadows and open grasslands, pedestrian and cycling paths, regional roads, rivers/streams/canals, and tracks and rural access roads.

For severity analysis, scores across all three consequence domains (fatality risk, property damage risk, and societal impact) were systematically varied by ± 1 ordinal level within the defined scale. For likelihood analysis, scores for all five external risk factors (obstacles, interference, communication, navigational environment, and electrical environment) were similarly perturbed. Each sensitivity run altered only one parameter domain at a time per area type to isolate effects. Corridor routes were recomputed across α values of 0.0, 0.5, and 1.0.

At $\alpha = 0$, all routing metrics show exactly 0.0% change across all severity and likelihood perturbations, confirming that efficiency-based optimization is independent of risk parameters. For severity adjustments at $\alpha = 0.5$, waterway corridors (rivers, canals, and streams) show the strongest influence: changing their severity score causes average risk across all routes to shift by up to 22.13%, with turn frequency increasing 4.8-6.3%. At $\alpha = 1.0$, a single-level change in waterway severity causes average height changes across all routes to vary by 26.29%. This shift from horizontal (turns) to vertical (height) sensitivity reflects the removal of energy penalties under pure risk minimization, allowing routes to freely switch between corridor types at different altitudes.

Likelihood adjustments produce substantially smaller impacts than severity changes. Most metrics change by less than 5% even under balanced optimization. A strong clustering pattern emerges where Obstacles and Communication factors produce identical results, while Interference, Navigation, and Electrical factors form a second group with matching impacts. This clustering suggests the five theoretical dimensions may effectively represent only two underlying risk mechanisms.

6.5.2. Sensitivity analysis on energy parameters

Energy parameter sensitivity was assessed by systematically perturbing the three core parameters: horizontal energy consumption per meter, vertical energy consumption for ascent, and vertical energy consumption for descent. Parameters were varied by $\pm 20\%$ individually and simultaneously, resulting in seven scenarios. The full corridor optimization was recomputed for each scenario.

Energy parameter variations have minimal impact on route selection. At $\alpha = 1.0$, routing metrics, except for energy consumption, remain completely unchanged across all energy perturbations, validating that pure risk optimization is independent of energy parameters. Energy consumption responds proportionally: $\pm 20\%$ parameter changes yield approximately $\pm 20\%$ consumption changes.

At $\alpha = 0.$ and $\alpha = 0.5$, despite efficiency considerations, route changes remain negligible. Vertical energy perturbations produce virtually no impact: path length varies by at most 0.1%, risk by 0.3%, and turns by 0.5. Horizontal energy changes yield minimal routing adjustments with maximum length changes of 1.2%. This limited influence can be attributed to restricted route options, dominance of horizontal consumption (>98% of total energy), and limited altitude transition variability in the discrete network structure. These findings raise questions about the effectiveness of height-based energy optimisation within infrastructure network constraints.

6.6. Conclusion on experiments

The experimental evaluation in urban, rural and suburban contexts shows that the proposed UAV corridor planning framework successfully generates feasible routes that balance safety and operational efficiency. The results reveal several key findings that validate the practical applicability of the framework.

The α parameter effectively controls the risk-efficiency trade-off, producing distinctly different corridor network topologies. Efficiency-focused routing yields direct paths through higher-risk infrastructure, while risk-averse planning systematically routes through natural corridors like waterways and meadows. The Pareto analysis confirms that substantial risk reduction can be achieved with modest efficiency penalties at lower α values, with diminishing returns thereafter. Optimal trade-off points vary by context, occurring earlier in urban environments than in rural or suburban settings.

Geographic context fundamentally shapes the magnitude and nature of these trade-offs. Urban environments demonstrate moderate efficiency penalties due to abundant alternative corridors, while rural settings experience severe route adaptations with nearly doubled path complexity. The tested suburban context shows an interesting pattern where substantial detours are offset by the availability of long,

continuous waterway corridors that maintain directional consistency.

Infrastructure utilization patterns shift dramatically under risk-aware planning. Each geographic context exhibits characteristic substitutions: urban areas shift from roads to waterways, rural areas from regional infrastructure to agricultural corridors, and suburban areas abandon all high-risk corridors in favor of natural infrastructure. The multi-depot case reveals that risk considerations can fundamentally restructure service areas, transforming balanced geographic allocation into highly asymmetric patterns dominated by high-risk infrastructure barriers.

Battery constraints expose critical differences between urban and rural viability. Urban environments maintain high service coverage even under risk-averse routing, while rural areas face severe coverage penalties that worsen dramatically with risk considerations. This disparity underscores that operational range, not just risk-efficiency trade-offs, determines practical feasibility in dispersed geographic settings.

The sensitivity analyses reveal fundamental differences in how the UAV corridor network responds to parameter uncertainty. Risk parameters can trigger complete route reorganization, while energy parameters only adjust the cost of existing routes without changing their relative preference. Severity parameters show the highest sensitivity (up to 26.29%), particularly for heavily-used corridor types. The high sensitivity to severity scores for dominant corridor types emphasizes the importance of accurate risk assessment. Likelihood parameters demonstrate moderate sensitivity (<5%) with unexpected clustering suggesting redundancy in the five-factor framework. Energy parameters produce minimal routing changes (<2%), as the discrete network structure and dominance of horizontal flight limit optimization flexibility. The patterns validate most of the model's internal consistency while highlighting which parameters require careful calibration.

The framework demonstrates robust performance across all tested conditions, successfully generating viable corridor networks that adapt to local geographic constraints while maintaining consistent trade-off structures.

7

Discussion and Conclusion

The rapid growth of e-commerce and urban logistics has increased societal concerns about congestion, emissions and quality of life in densely populated areas. Unmanned aerial vehicles (UAVs) have emerged as a promising technological alternative, but their widespread deployment is still constrained by spatial, regulatory and operational factors, especially in a complex airspace environment such as that of the Netherlands. The aim of this thesis was to investigate whether fixed UAV corridors for freight operations are feasible within these constraints, while systematically dealing with complex trade-off's between safety risks and operational efficiency.

To address this question, this thesis developed and implemented a spatial corridor design framework based on regulatory constraints, operational feasibility and a risk-aware routing strategy. The model integrates airspace regulations, land use restrictions and infrastructure-based routing into a graphical representation of feasible corridors, enabling planning with trade-off's between energy consumption and risk.

7.1. Discussion on results

The experimental results reveal several critical insights that extend beyond the immediate findings and warrant deeper examination of their implications for UAV corridor design and implementation.

Balancing connectivity with infrastructure risk

The results highlight a fundamental trade-off between safety and network connectivity. While the framework designates most areas restricted for 'open' category UAVs as no-fly zones, it retains high-speed roads (>80 km/h), railways, and power lines with elevated risk scores to avoid network fragmentation. This pragmatic choice has clear implications: at full risk aversion (α = 1), some high-risk segments remain indispensable for maintaining connectivity and are therefore included in optimal routes. These findings suggest, unstead of blanket avoidance, implementing targeted mitigation at unavoidable crossings, such as increased flight altitude, speed limitations, temporal restrictions, or protective structures. This selective risk acceptance combined with localized safety measures aligns with EASA U-space regulations, which require defining mitigations to maintain acceptable risk levels [11].

Infrastructure-aligned facility placement

Current logistics infrastructure is poorly aligned with risk-aware UAV operations. Distribution centres along highways, optimal for ground transport, force UAVs to first utilize or even avoid these high-risk areas before accessing safer corridors like waterways. The suburban case study further demonstrates how high-risk infrastructure fragments operational space into distinct, risk-bounded domains rather than continuous service areas. Consequently, UAV networks cannot simply overlay traditional geographical regions but must instead operate within areas defined and constrained by risk profiles. These findings highlight a critical mismatch between the spatial logic of ground-based logistics and the requirements of safe, efficient UAV operations.

Operational constraints versus policy preferences

Battery constraint analysis shows that preferences for risk mitigation are inextricably linked to operational viability. As demonstrated in the national case study, risk-averse routing drastically reduces coverage, thereby jeopardising economic viability. These results call for a hierarchical optimisation approach where operational constraints define the boundaries for risk-efficiency trade-offs. Rather than treating α as a purely policy-driven parameter, its selection must first ensure minimum viable coverage given technological and geographic limitations. Within these operational limits, risk preferences can then be applied. This approach prevents adopting routing strategies that, while safer in theory, render the service economically unsustainable due to inadequate customer reach.

Risk assessment framework validity and parameter sensitivity

The differing sensitivities of severity and probability parameters validate the risk calculation framework while highlighting areas for improvement. The pronounced sensitivity to severity scores, particularly for heavily used corridors, confirms that impact estimates drive routing decisions more than probability estimates. This aligns with the framework's structure, where severity influences all three risk domains (fatality, property, societal), while probability affects a portion of only two domains. These patterns have practical implications. Validation efforts should prioritise accurate severity estimates for infrastructure types, as errors here can restructure entire networks. In contrast, the framework's relative insensitivity to probability suggests that precise probability estimates are less critical, enabling simplified assessments for these parameters.

Dimensional reduction in likelihood assessment

The probability sensitivity analysis uncovers a notable structural finding: although the framework defines five external risk factors, it effectively captures only two independent dimensions. Obstacles and communication cluster together, as do interference, navigation, and electrical factors, indicating systematic correlations in the assessment approach. This dimensional collapse likely reflects genuine factor correlations. For example, physical density affects both obstacle presence and communication quality, while dense technical infrastructure influences interference, navigational accuracy, and electrical disturbance. This suggests either that urban risk factors are inherently more interlinked than theoretical models assume, or that the SQRA framework oversimplifies these relationships.

Geographic disparities in feasibility

Urban environments show greater adaptability to risk-aware routing. For example, Breda achieves similar risk reductions at lower efficiency costs than rural Borsele (50% vs 60% path length increase at their respective optimal α values). This is largely due to differences in infrastructure density and diversity. Urban areas offer numerous alternative corridors, such as waterways, parks, and quiet streets, enabling efficient detours around high-risk zones. In contrast, rural areas face sparse infrastructure and concentrated barriers like railways and high-voltage lines, creating bottlenecks that force longer detours despite comparable risk reduction goals. These findings suggest that UAV corridor viability depends not only on delivery distance but also on infrastructure topology. Urban areas, despite their density, may thus offer more favourable conditions for risk-aware UAV operations due to their rich network of alternative routes. However, the current grid-based model may underestimate rural routing efficiency by forcing rectilinear paths along grid cells rather than allowing direct diagonal trajectories. A continuous-space model could better capture routing flexibility in rural areas, potentially narrowing the urban-rural disparity.

7.2. Limitations and further research

Despite the valuable contributions of this study, it is important to acknowledge its limitations. This section critically reviews these limitations and suggests options for future research to address them.

Energy model limitations and oversimplification

The sensitivity analysis reveals that energy parameters have negligible influence on routing decisions, with $\pm$20% variations producing less than 2% change in route selection. Although the framework optimizes risk versus efficiency, the simplified energy model effectively reduces this to a comparison of path length versus risk. With horizontal consumption (0.025 Wh/m) dominating (>98% of total energy)

and vertical costs (2.03/0.93 Wh per 30m) representing less than 2%, altitude decisions rarely affect routing.

Critical limitations include the absence of turn penalties and the binary altitude choice (30m/60m), which fails to capture realistic energy consumption dynamics. Routes with frequent sharp turns are treated identically to smooth paths of similar length, particularly impacting risk-averse routes that often involve circuitous trajectories. The infrastructure-aligned corridor design further constrains optimization by limiting altitude adjustments to discrete transition points between corridor types. Despite effectively demonstrating the feasibility of risk-aware corridor planning, operational deployment requires enhanced energy models incorporating turning costs, continuous altitude optimization, environmental effects, and vehicle-specific parameters. Future research should explore whether refined models can overcome these constraints, though results suggest infrastructure-bound routing may intrinsically limit three-dimensional optimization flexibility.

Limitations of grid-based network representation

A key limitation of the current model is its grid-based path generation. Using a 50×50 meter grid to manage computational complexity forces UAVs to follow rectilinear paths along grid edges rather than direct diagonal routes, systematically lengthening travel distances. This discretisation introduces biases. For instance, open agricultural areas, suitable for direct flights, are penalised relative to linear infrastructure corridors. Additionally, under efficiency-based optimisation (α = 0.0), route variations remain minimal despite $\pm$20% changes in energy parameters. This could also suggest that the grid constraints, rather than energy dynamics, often dictate routing outcomes.

Future research should explore alternatives to grid-based routing. Although the model allows adjustable grid resolutions, empirical analysis is needed to identify optimal balances between computational cost and routing accuracy. More advanced methods, such as pre-computed key crossing points or pruned visibility graphs, could enable direct path generation without imposing rectilinear constraints while maintaining computational feasibility. Despite these limitations, the grid-based approach does enable critical functionality for polygon-based areas. As validated in the Borsele case study (Figure 6.11), the grid successfully integrates large agricultural and open land areas into the corridor network, providing essential connectivity in rural environments where linear infrastructure is sparse.

Network connectivity through snapping tolerances

The network construction relies on snapping algorithms to ensure connectivity between different infrastructure types, such as grid-based polygons and adjacent linear corridors. Fixed snapping tolerances were applied to connect nodes within specified distances, creating network continuity where physical infrastructure may be spatially separated. This approach is reasonable for aerial operations where UAVs can easily traverse small gaps, unlike ground vehicles that face physical barriers such as curbs or bollards. The chosen tolerance values represent pragmatic assumptions rather than empirically optimized distances. The model's configurable snapping parameters allow future research to investigate optimal tolerances that balance network connectivity with operational realism, potentially varying by area type or infrastructure category.

Risk Assessment Framework Limitations

The SQRA framework facilitates clear risk communication and stakeholder engagement but has key operational limitations. First, it excludes dynamic hazards, focusing solely on static ground risks. For example, open grasslands are rated favourably despite potential for strong wind gusts and turbulence that can destabilise UAVs [27], while waterways may induce fog formation or wind channeling effects. More critically, intersection conflicts from climbing or descending flights at route transitions are omitted, despite representing major collision risks that require active management [8].

Second, likelihood analysis reveals only two independent risk dimensions rather than the five theoretical factors. Obstacles and communication cluster under physical density, while interference, navigation, and electrical factors cluster under technical infrastructure density. This suggests inherent correlations among urban risk factors or possible oversimplification in the current methodology. Future models could adopt a two-factor approach for simplicity. However, it is more likely that this clustering results from the framework's reliance on structured assumptions rather than empirical data. While this approach enables rapid deployment, it introduces uncertainties, as relative rankings may not accurately reflect

operational realities. In effect, the framework risks measuring similar underlying phenomena multiple times.

Future work should prioritise validation using incident databases or operational data to replace assumptions with observed likelihoods. Severity estimates could be improved through expert elicitation or detailed geographic data analysis, such as using high-resolution satellite imagery. Developing standardised data collection protocols in collaboration with UAV operators and aviation authorities would further enhance model credibility and regulatory acceptance. Despite these limitations, the framework provides a robust foundation for risk-integrated UAV planning. Its modular design supports iterative refinement as empirical evidence accumulates, facilitating the transition from theoretical assessments to evidence-based corridor design.

Data quality limitations

The framework's dependence on OpenStreetMap data introduces inherent uncertainties regarding infrastructure completeness and classification accuracy. While OSM provides comprehensive coverage, the crowd-sourced nature of the data may result in inconsistent tagging of socially sensitive locations, outdated infrastructure representations, or missing critical features that affect risk assessment. Additionally, the PostNL pickup point dataset may not reflect current operational reality, as network optimization frequently adjusts service point locations. The distribution center coordinates, obtained indirectly through Schuberg Philis rather than directly from PostNL, introduce additional uncertainty regarding their accuracy and currency. However, the framework's modular design allows users to input custom origin and destination points, mitigating these data quality concerns for operational deployment where accurate facility locations would be available.

Operational prerequisites and deployment readiness

This study intentionally excludes major urban centres within controlled airspace, such as Rotterdam and Amsterdam. As outlined in Appendix B.1, integrating operations in these areas would require extensive coordination with air traffic control, introducing complexities beyond this proof-of-concept. The current framework thus demonstrates fundamental corridor design principles in simpler airspace environments before tackling the layered challenges of CTR integration.

More importantly, the model assumes a mature UAV ecosystem that remains largely theoretical in the Netherlands. Full implementation would depend on continuous GNSS coverage, widespread 5G connectivity, pre-allocated timeslot systems, and comprehensive U-space services including strategic deconfliction, conformance monitoring, and emergency management. Future research should examine the technical and regulatory prerequisites for operations in complex airspace and the development of supporting infrastructure. Despite these limitations, the framework provides crucial groundwork. By identifying optimal corridor locations and quantifying risk trade-offs, it offers a spatial blueprint to guide infrastructure investments and regulatory development as the UAV ecosystem evolves.

7.3. Main findings

To explore whether fixed UAV corridor for freight operations are feasible within the spatial, operational and regulatory constraints of Dutch airspace, with a balance between ground risk and energy efficiency, six sub-questions have been drawn up. This chapter answers the main question and these six sub-questions separately.

Main Research Question

Fixed UAV corridors can be feasibly designed within the constraints of Dutch airspace while effectively balancing ground risk and energy efficiency. This study shows that by addressing three key operational constraints (competing feasibility requirements, limited airspace availability and European regulatory mandates), a practical corridor framework can be established. The proposed infrastructure-based corridor approach, combined with a Semi-Quantitative Risk Assessment (SQRA) methodology and α-weighted optimisation, enables systematic trade-offs between safety and design efficiency. Experimental validation in urban, rural and suburban contexts confirms the practical feasibility, with balance points around α = 0.10-0.20 achieving 70-80% risk reduction with only 15-20% efficiency loss.

Sub-question 1: Design Space Constraints

Three overarching constraints define the UAV operational design space in the Netherlands.

- Competitive feasibility requires UAV operations to be cost-effective relative to conventional delivery modes, achievable through fast delivery requirements, low route density scenarios, short-to-medium distances, and fixed-point delivery structures.
- Limited airspace availability stems from the Dutch airspace system's strict spatial limitations, including controlled airspace zones, civil and military usage areas, and designated no-fly zones over sensitive ground locations.
- European regulatory constraints mandate compliance with the U-space framework requiring airspace management services such as geo-awareness and flight authorization and a formal risk assessment, while delegating corridor spatial design responsibilities to individual Member States.

Sub-question 2: Spatial Corridor Structuring

UAV corridors can be effectively structured using fixed, infrastructure-aligned routing based on linear infrastructure and natural corridors. However, corridor placement cannot be arbitrary, as implementation in crowded urban environments poses additional challenges related to ground-level risks such as noise, privacy issues and third-party safety risks. Pursuing airspace safety through fixed corridors may therefore come at the expense of operational efficiency and requires a careful trade-off between safety risks and operational performance.

Despite these trade-offs, the fixed corridor approach proves particularly suitable for Very Low Level (VLL) airspace operations, where the constrained environment requires structured routing via predefined origin-destination pairs. The fixed corridor structure provides a direct solution to the three primary integration challenges identified in this study: enabling competitive feasibility through simplified operations and faster regulatory approval processes, enabling strategic avoidance of controlled airspace and no-fly zones to ensure spatial feasibility, and supporting regulatory compliance through corridor-level risk assessments and the implementation of U-space services, as required by the EASA.

Sub-question 3: Network Model Translation

The spatial corridor design was successfully translated into a computational network model by graphical representation using OSMnx-derived geospatial data. Distribution centres and pick-up points serve as origin-destination nodes, while linear infrastructure and natural corridors define the edges of the network. The graph construction process integrates legal no-fly zones, connects infrastructure segments through node procedures and integrates polygon-based areas through grid views. Additional design assumptions, including slot-based traffic management, GNSS positioning and 5G connectivity, ensure computational ease while maintaining essential safety and feasibility aspects.

Sub-question 4: Risk Assessment Integration

Third-party ground risks are systematically integrated through a semi-quantitative risk assessment (SQRA) tailored for corridor-level evaluation. The methodology includes hazard identification focusing on spatially defined external risk factors and systematic probability and severity scores on ordinal scales. Risk is aggregated across three consequence areas: fatalities, property damage and societal impact. The overall risk score for each area depends on adjustable domain weights (α_f = 0.4, α_p = 0.3, α_s= 0.3), which can be adjusted to changing policy priorities or stakeholder preferences through the transparent structure of the framework.

The SQRA approach balances analytical tractability with regulatory transparency, in line with EASA requirements for hazard identification and risk analysis under U-space regulations. Moreover, the methodology enables spatial quantification that supports route optimisation across different corridor sections. It provides area-specific risk scores that can be directly integrated into the path planning algorithm for systematic safety efficiency trade-offs.

Sub-question 5: Optimization Framework

Optimal corridor planning is achieved through a composite cost function that integrates third-party risk scores and energy consumption estimates using a weighting parameter α. The framework weighs area-

based risk assessments against energy costs that include both horizontal flight (25 Wh/km) and vertical transitions (2.03 Wh ascent, 0.93 Wh descent per 30m). The parameter α controls the relative emphasis between risk minimisation and energy efficiency, with α = 0 focusing entirely on energy optimisation, α = 1 prioritising risk reduction only, and intermediate values allowing balanced trade-offs between the two objectives.

Dijkstra's algorithm provides globally optimal route selection under the given cost structure and provides computational efficiency for large-scale network optimisation. The α parameter allows flexible policy control over the trade-off between safety and efficiency, allowing corridor designs to be adapted to different operational contexts, legal requirements or stakeholder priorities. This parametric approach makes the framework scalable for large-scale corridor planning applications while maintaining transparency in the trade-off decisions.

Sub-question 6: Model Performance Validation

The corridor design model successfully produces safe and efficient routes across all geographic contexts tested, although the balance between safety and efficiency varies slightly between environments. The framework demonstrates its core capability by consistently generating viable corridor networks that adapt to local constraints while maintaining the intended trade-off structures. The α parameter yields distinctly different network topologies, with efficiency-focused routing producing direct paths through higher-risk infrastructure, and risk-averse planning systematically avoiding high-risk areas via longer detours along natural corridors. This validates the fixed corridor concept's adaptability to different policy priorities. The framework's consistent optimisation principles across urban, rural, and suburban settings suggest broader applicability beyond the test cases, supporting strategic airspace planning at both local and national levels.

The framework's capability to integrate diverse infrastructure types is further validated through the successful implementation of polygon-based area integration. As demonstrated in Figure 6.11 for the Borsele case study, the model effectively generates grid structures within polygon areas and establishes seamless connections to adjacent linear infrastructure through boundary nodes. This validation confirms that the network construction methodology successfully handles both linear corridors (roads, railways, waterways) and area-based zones (agricultural lands, open grasslands), enabling comprehensive coverage even in rural environments where linear infrastructure is sparse. The ability to route through grid-based polygons while maintaining connectivity to traditional linear corridors demonstrates the framework's adaptability to heterogeneous infrastructure configurations.

The effectiveness of the model is particularly clear in urban settings, where dense infrastructure networks allow significant safety improvements with minimal efficiency losses at low parameter settings. However, significant safety gains can also be achieved in both urban and suburban contexts with only a small decrease in efficiency. This indicates that the framework allows feasible planning of routes that are both efficient and safe within the Netherlands. The tested suburban environment clearly shows the adaptability of the model, as it replaces risky corridors over railways with safer natural corridors such as waterways, thus preserving safety while maintaining topological efficiency as much as possible.

Sensitivity analyses confirm expected behavioural patterns: changes in energy parameters produce no effect under risk-averse planning (α = 1.0), while changes in severity parameters have minimal impact under efficiency-focused routing (α = 0.0). This validates the framework's internal logic and demonstrates its robust performance under extreme scenarios, indicating that the model functions as intended across varying assumptions and conditions.

7.4. Comparison with Literature

The results strongly validate the structured corridor approach advocated by Doole [8] and Zhang et al. [67]. The experimental evidence demonstrates that one-way, infrastructure-aligned corridors can successfully enhance safety through systematic risk reduction while maintaining operational viability. This confirms Doole's simulation findings that structured designs outperform unstructured alternatives in VLL airspace. However, this study extends beyond Doole's exclusive focus on road networks by demonstrating that waterways, railways, and natural corridors provide equally viable, and often superior, routing alternatives, particularly under risk-aware planning.

The observed efficiency penalties align with the theoretical predictions of Zhang et al. [67] and Bauranov and Rakas [4], who argued that structured airspace designs inherently trade efficiency for safety. However, the magnitude of these penalties is notably lower than might be expected from the literature's emphasis on significant detours. This suggests that infrastructure diversity in the Netherlands provides more routing flexibility than anticipated, particularly in urban environments.

The experimental results provide empirical validation for the competitive feasibility criteria outlined by Kennisinstituut voor Mobiliteitsbeleid [27]. The stark contrast between urban and rural coverage under battery constraints directly supports their emphasis on short-to-medium distances as a key feasibility requirement. Urban areas like Breda achieve near-complete coverage even under risk-averse routing, while rural Borsele faces severe limitations (17.4% coverage at 400 Wh), confirming that route density and distance fundamentally determine UAV viability relative to traditional van delivery.

The minimal impact of altitude variations on routing decisions contrasts with the findings of Ren and Cheng [48] and Michel et al. [35], who emphasized the importance of altitude optimization for both safety and energy efficiency. The experimental results show altitude changes remaining stable ($\pm$10%) across all risk configurations, suggesting that infrastructure-aligned routing may inherently constrain vertical optimization opportunities.

While the Metropolis study by Sunil et al. [63] examined airspace structuring in unconstrained environments, this research demonstrates how infrastructure diversity enables effective structuring even within highly constrained VLL airspace. The systematic substitution patterns observed, urban areas shifting from roads to waterways, rural areas utilizing agricultural corridors, represent a novel finding that extends beyond the literature's typical focus on single infrastructure types.

7.5. Policy implications

This section discusses the policy implications of the proposed framework, outlining how its technical capabilities can support risk-informed decision-making, regulatory development, and operational deployment strategies.

Adaptive risk management through parameter customization

A key strength of the framework is its adaptability to local contexts via customizable parameters. Authorities can use the α parameter to differentiate spatial risk tolerance, allowing higher risk acceptance in some zones while enforcing stricter standards elsewhere by adjusting α values accordingly. Additionally, domain-specific risk weights (α_f, α_p, α_s) enable prioritisation of local concerns, such as emphasising social impacts in privacy-sensitive areas or reducing fatality risks in dense urban centres. Furthermore, the simplicity of the framework's structure and risk formulation ensures that additional risk dimensions can be seamlessly integrated if required for specific policy applications. This transparent approach bridges technical modelling and policy implementation, providing interpretable trade-off mechanisms that replace opaque algorithmic decisions with clear and communicable strategies.

Systematic design with local refinement

The implementation highlights both strengths and areas for improvement. While the framework integrates critical infrastructure such as schools as elevated-risk zones, many sensitive ground elements remain unmapped. Event venues, sports parks, and similar gathering places require explicit inclusion by local authorities. The model's architecture supports such customisation, enabling municipalities to assign high-risk scores or designate these sites as no-fly zones. A two-stage implementation is recommended. First, generate network-wide corridors using carefully selected α parameters with operational constraints. Second, refine these networks with local knowledge, incorporating site-specific adjustments such as identifying unmapped sensitive locations, optimising infrastructure crossings via existing bridges or tunnels, or accounting for seasonal recreational area use.

Rather than excluding all high-risk infrastructures, targeted mitigation at critical crossings is preferable. The model naturally avoids high-risk areas, but where crossings are unavoidable, specific interventions can ensure safety. Rail crossings may require protective netting or mandatory altitude requirements, while roads can utilise existing overpasses. This pragmatic strategy preserves connectivity while aligning with EASA risk mitigation standards [11].

Phased deployment strategy

The U-space Concept of Operations [59] offers a structured pathway for implementation via its VLL airspace typology (Types X, Y, Z), which categorises areas based on UAV density, ground risk, manned aircraft activity, and public acceptance. This typology supports a phased corridor deployment strategy: starting in Type X rural areas with minimal airspace complexity and risk, then expanding to Type Y suburban areas, and finally to Type Z urban environments as U-space services and operational experience mature. Initial operations should adopt conservative α values (0.80–0.90) to prioritise safety and build public trust. As safety records strengthen, parameters can gradually shift toward more balanced values (0.10–0.20) to optimise efficiency. This phased approach enables regulatory refinement through operational data, maintains safety standards during expansion, and fosters public confidence by demonstrating success in lower-risk contexts.

Multi-use corridor development

Corridor infrastructure should be viewed as a versatile platform supporting applications beyond freight transport. This could maximise infrastructure returns, which is important for competitive feasibility and enable tailored operational parameters for diverse uses. Medical logistics, such as organ transport and pharmaceutical deliveries, could benefit from dedicated express lanes with adjusted risk thresholds reflecting their societal importance. Emergency response operations may require dynamic corridor activation, while agricultural monitoring and industrial inspections can utilise scheduled access during off-peak hours, reducing congestion and maintaining service quality. Each application can leverage the same infrastructure-based routing principles in a customised manner.

This multi-use framework transforms UAV corridors from single-purpose routes into versatile air highways. By accommodating varied operational needs under a unified system, governments can justify investments while fostering innovation across sectors. Flexibility to adapt risk parameters, planning protocols, and performance standards for each application is key to maintaining system integrity and safety.

7.6. Strategic recommendations

This section presents strategic recommendations derived from the study's findings, offering guidance for facility placement, U-space service development, and evidence-based expansion of UAV corridor networks.

Strategic facility placement and infrastructure optimization

Building on the findings in this study advises companies to reimagine their logistics infrastructure beyond traditional road-focused models. As highlighted in Section 7.1, there is a clear need to reconsider facility location strategies for UAV logistics. Businesses should evaluate relocating distribution centres away from highway intersections towards low-risk nodes, such as waterway junctions or the edges of green spaces, to reduce delivery risks while maintaining operational efficiency. Also, the UAV distribution networks service areas should not be viewed as geographical regions, but as risk-defined areas naturally bounded by high-risk infrastructure. Additionally, companies managing customer pick-up and delivery points can benefit from positioning these facilities at the periphery of residential areas or near low-risk corridors, rather than within dense commercial centres. Similarly, healthcare providers, such as pharmacies, could enhance accessibility and resilience by integrating dedicated UAV access points into their facilities. Organisations are encouraged to use this framework's data-driven insights to optimise facility placement and adapt their infrastructure for effective UAV integration.

U-space service development opportunities

Fixed corridor networks offer focused opportunities for U-space service development that are both technically feasible and economically viable. Instead of aiming for comprehensive urban coverage, service providers can prioritise corridor-specific solutions with immediate value. Key development areas include real-time weather monitoring for corridor microclimates, optimised 5G coverage leveraging existing infrastructure, pre-approved flight authorisations, dynamic capacity management during peak hours, and automated compliance monitoring with deviation alerts. These specialised services benefit from the economies of scale of fixed corridors, enabling cost-effective deployment and maintenance compared to citywide systems. A phased development strategy allows operators to expand service

areas incrementally, based on proven success and operational learning. This approach reduces initial capital investment while building expertise and public confidence, supporting scalable growth of U-space services alongside corridor network expansion.

Collaborative framework for evidence-based development

Phased deployment enables the collection of empirical risk data, supporting evidence-based policy development over rigid initial regulations. Standardised data collection procedures, developed collaboratively with operators, insurers, and aviation authorities, enhance model credibility and regulatory acceptance while facilitating the transition from theoretical assessments to practical corridor planning.

Robust stakeholder collaboration is essential to bridge theoretical models and operational reality. Joint efforts to establish standardised data protocols will validate and refine models, support evidence-based policymaking, and build a shared understanding across the ecosystem. Early deployments in low-risk environments can generate baseline performance data, while subsequent expansions provide comparative insights across operational contexts. This empirical foundation improves credibility and accelerates regulatory acceptance. Standardisation ensures consistency across operators and jurisdictions, enabling meaningful comparisons and fostering collective learning. Such collaboration transforms isolated operational experiences into shared knowledge that advances the entire sector.

Comprehensive performance benchmarking

The competitive advantages of UAV networks designed using this framework should be strictly validated by systematic benchmarking against traditional van-based systems on identical origin-destination pairs. Key performance indicators include journey times under varying traffic conditions, emissions per energy source, energy consumption per package, operational costs (including labour and infrastructure) and delivery reliability. Such a comprehensive analysis will determine where the corridor-based approach of this framework outperforms conventional logistics. This will guide strategic deployment and regulatory decisions. This evidence-based assessment will ensure that the application of this design concept delivers real value. The results will inform both private sector investment decisions and public sector regulatory priorities.

Technical capacity constraints and operational boundaries

Implementing UAV corridors requires careful consideration of technical limitations such as battery capacity, payload capacity, range, and communication reliability. Energy consumption analysis highlights critical operational boundaries for corridor planning: while current battery technology suffices for dense urban environments, rural deployments face constraints that necessitate higher-capacity batteries or strategically located charging infrastructure. Municipalities or private investors need clear operational guidelines that translate these technical constraints into actionable investment decisions, ensuring realistic planning while retaining flexibility for future expansion. This is particularly important given the pronounced differences in urban and rural feasibility identified by energy analyses. Continuous reassessment of these boundaries is essential as technology evolves. Improvements in battery capacity and communication reliability should be integrated into operational frameworks to prevent overcommitment and enable systematic growth. Aligning deployment strategies with technological progress ensures that early implementations remain feasible while paving the way for more ambitious future networks.

References

[1] A. Allouch et al. “Qualitative and Quantitative Risk Analysis and Safety Assessment of Unmanned Aerial Vehicles Missions over the Internet”. In: (Apr. 2019). DOI: 10.1109/ACCESS.2019.2911980.

[2] Amazon. *Amazon - Help & Customer Service*. 2025. URL: https://www.amazon.com/gp/help/customer/display.html?nodeId=T3jxhuvPfQ629BOIL4.

[3] C. Badea et al. *Limitations of Conflict Prevention and Resolution in Constrained Very Low-Level Urban Airspace*. Tech. rep. 2021.

[4] A. Bauranov and J. Rakas. “Designing airspace for urban air mobility: A review of concepts and approaches”. In: *Progress in Aerospace Sciences* 125 (Aug. 2021). ISSN: 03760421. DOI: 10.1016/j.paerosci.2021.100726.

[5] CBS. *Inwoners per gemeente*. 2025. URL: https://www.cbs.nl/nl-nl/visualisaties/dashboard-bevolking/regionaal/inwoners.

[6] L. Davies et al. “Review of Unmanned Aircraft System Technologies to Enable beyond Visual Line of Sight (BVLOS) Operations”. In: *2018 10th International Conference on Electrical Power Drive Systems, ICEPDS 2018 - Conference Proceedings*. Institute of Electrical and Electronics Engineers Inc., Dec. 2018. ISBN: 9781538647134. DOI: 10.1109/ICEPDS.2018.8571665.

[7] L. De Filippis, G. Guglieri, and F. Quagliotti. “A minimum risk approach for path planning of UAVs”. In: *Journal of Intelligent and Robotic Systems: Theory and Applications*. Vol. 61. 1-4. Jan. 2011, pp. 203–219. DOI: 10.1007/s10846-010-9493-9.

[8] M. M. Doole. “Urban airspace design for autonomous drone delivery”. In: (2022). DOI: 10.4233/uuid:1e5679a7-e9d1-40d0-8857-d3f0eb290510. URL: https://doi.org/10.4233/uuid:1e5679a7-e9d1-40d0-8857-d3f0eb290510.

[9] Dutch Drone Delta. *Living Lab UAV delivery Amsterdam*. Tech. rep. Apr. 2022.

[10] EASA. *Drones - regulatory framework background*. 2025. URL: https://www.easa.europa.eu/en/domains/civil-drones/drones-regulatory-framework-background.

[11] EASA. *Easy Access Rules for U-space*. Tech. rep. 2024. URL: https://www.easa.europa.eu/en/document-library/easy-access-rules/online-publications/easy-access-rules-u-space?page=6.

[12] EASA. *Met een drone vliegen in de buurt van mensen*. 2025. URL: https://www.easa.europa.eu/nl/light/topics/flying-drones-close-people.

[13] EASA. *Opinion No 01/2020*. Tech. rep. 2020. URL: https://www.easa.europa.eu/sites/default/files/dfu/ToR%20RMT.0230%20%E2%80%93%20Issue%202.pdf.

[14] European Commission. *Commission Implementing Regulation (EU) 2021/664 of 22 April 2021 on a regulatory framework for the U-space.* Tech. rep. 2021.

[15] European Commission. “Commission Implementing Regulation (EU) 2024/1109 of 10 April 2024 laying down rules for the application of Regulation (EU) 2018/1139 of the European Parliament and of the Council as regards competent authority requirements and administrative procedures for the certification, oversight and enforcement of the continuing airworthiness of certified unmanned aircraft systems, and amending Implementing Regulation (EU) 2023/203”. In: (2024).

[16] European Commission. *Implementing Regulation (EU) 2019/947 of the Commission of 24 May 2019 on the rules and procedures for the operation of unmanned aircraft.* Tech. rep. 2021.

[17] V. Garg et al. “Drones in last-mile delivery: A systematic review on Efficiency, Accessibility, and Sustainability”. In: *Transportation Research Part D: Transport and Environment* 123 (Oct. 2023). ISSN: 13619209. DOI: 10.1016/j.trd.2023.103831.

[18] A. Goodchild and J. Toy. “Delivery by drone: An evaluation of unmanned aerial vehicle technology in reducing CO2 emissions in the delivery service industry”. In: *Transportation Research Part D: Transport and Environment* 61 (June 2018), pp. 58–67. ISSN: 13619209. DOI: 10.1016/j.trd.2017.02.017.

[19] Government of the Netherlands. *Regeling beperking of verbod uitoefening burgerluchtverkeer in bepaalde gebieden 2018*. 2025. URL: https://wetten.overheid.nl/BWBR0040730/2025-02-20.

[20] M. H. Güngören et al. “Investigating the Impact of Drone Transport on the Stability of Monoclonal Antibodies for Inter-Hospital Transportation”. In: *Journal of Pharmaceutical Sciences* 113.7 (July 2024), pp. 1816–1822. ISSN: 15206017. DOI: 10.1016/j.xphs.2024.04.002.

[21] X. Hu et al. “Risk Assessment Model for UAV Cost-Effective Path Planning in Urban Environments”. In: *IEEE Access* 8 (2020), pp. 150162–150173. ISSN: 21693536. DOI: 10.1109/ACCESS.2020.3016118.

[22] S. H. Huang et al. “Solving the vehicle routing problem with drone for delivery services using an ant colony optimization algorithm”. In: *Advanced Engineering Informatics* 51 (Jan. 2022). ISSN: 14740346. DOI: 10.1016/j.aei.2022.101536.

[23] M. Jacewicz et al. “Quadrotor Model for Energy Consumption Analysis”. In: *Energies* 15.19 (Oct. 2022). ISSN: 19961073. DOI: 10.3390/en15197136.

[24] A. Jazairy et al. “Drones in last-mile delivery: a systematic literature review from a logistics management perspective”. In: *International Journal of Logistics Management* (2024). ISSN: 17586550. DOI: 10.1108/IJLM-04-2023-0149.

[25] D. Karve and F. Kapadia. *Multi-UAV Path Planning using Modified Dijkstra’s Algorithm*. Tech. rep. 28. Oct. 2020, pp. 975–8887. DOI: 10.5120/ijca2020920816.

[26] U. C. Kaya, A. Dogan, and M. Huber. *A Probabilistic Risk Assessment Framework for the Path Planning of Safe Task-Aware UAS Operations*. Tech. rep. July 2019. DOI: 10.2514/6.2019-2079.

[27] Kennisinstituut voor Mobiliteitsbeleid. *Drones in het personen-en goederenvervoer*. Tech. rep. 2017.

[28] R. Y. Kim. “The Impact of COVID-19 on Consumers: Preparing for Digital Sales”. In: *IEEE Engineering Management Review* 48.3 (July 2020), pp. 212–218. ISSN: 19374178. DOI: 10.1109/EMR.2020.2990115.

[29] Y. Kim and J. Bae. “Risk-Based UAV Corridor Capacity Analysis above a Populated Area”. In: *Drones* 6.9 (Sept. 2022). ISSN: 2504446X. DOI: 10.3390/drones6090221.

[30] A. Konert and T. Dunin. “A harmonized European drone market? – New EU rules on unmanned aircraft systems”. In: *Advances in Science, Technology and Engineering Systems* 5.3 (Feb. 2020), pp. 93–99. ISSN: 24156698. DOI: 10.25046/aj050312.

[31] B. López et al. “Path planning and collision risk management strategy for multi-UAV systems in 3D environments”. In: *Sensors* 21.13 (July 2021). ISSN: 14248220. DOI: 10.3390/s21134414.

[32] LVNL. *Downloadable Aeronautical Information Products*. Mar. 2025. URL: https://www.lvnl.nl/diensten/aip/downloads.

[33] S. Lyon-Hill et al. “Virginia Tech: Measuring the Effects of Drone Delivery in the United States”. In: (Sept. 2020). URL: http://hdl.handle.net/10919/100104.

[34] M. Majka. *Semi-Quantitative Risk Assessment: Bridging the Gap Between Qualitative and Quantitative Methods*. Tech. rep. Aug. 2024. URL: https://www.researchgate.net/publication/383239835.

[35] N. Michel et al. “Energy-Optimal Planning of Waypoint-Based UAV Missions – Does Minimum Distance Mean Minimum Energy?” In: (Oct. 2024). DOI: 10.1109/IROS58592.2024.10801888. URL: http://arxiv.org/abs/2410.17585.

[36] Ministerie van Infrastructuur en Waterstaat. *Aanzet voor een visie op U-space*. Tech. rep. Monitor Deloitte, 2021. URL: https://ec.europa.eu/defence-industry-space/civil-aerospace-industry/unmanned-aircraft_en.

[37] Ministerie van Infrastructuur en Waterstaat. *Actieplan-Programma Onbemande Luchtvaart 2023-2025*. Tech. rep. Apr. 2023.

[38] Ministerie van Infrastructuur en Waterstaat. *Toetsingskader U-spaceluchtruim Deel: verkenning en advies*. Tech. rep. 2023. URL: https://www.rijksoverheid.nl/documenten/rapporten/2023/02/28/toetsingskader-u-spaceluchtruim-verkenning-en-advies.

[39] Ministerie van Infrastructuur en Waterstaat. *Uitvoeringsagenda Luchtvaartnota*. Tech. rep. Nov. 2020.

[40] T. Mu et al. "Improved Network Reliability Optimization Model with Head Loss for Water Distribution System". In: *Water Resources Management* 35.7 (May 2021), pp. 2101–2114. ISSN: 15731650. DOI: 10.1007/s11269-021-02811-9.

[41] Nationaal Georegister. *Natura2000 WFS*. Jan. 2024. URL: https://nationaalgeoregister.nl/geonetwork/srv/dut/catalog.search#/metadata/32b1eb9e-c54f-4598-92d2-328eb77fa0d3.

[42] Odido. *Drone Delivery Services: klaar voor de toekomst met 5G.* 2025. URL: https://www.odido.nl/zakelijk/oplossingen/5g/drone-delivery-services.

[43] K. L. Palmerius et al. "End-to-end drone route planning in flexible airspace design". In: *Transportation Research Interdisciplinary Perspectives* 27 (Sept. 2024). ISSN: 25901982. DOI: 10.1016/j.trip.2024.101219.

[44] B. Pang et al. *Third Party Risk Modelling and Assessment for Safe UAV Path Planning in Metropolitan Environments*. Tech. rep. 2021. DOI: 10.48550/arXiv.2107.01834.

[45] Platform Nederlandse Luchtvaart. "Fact sheet Luchtruimclassificatie en -indeling". In: (2013).

[46] E. Politi et al. "A Survey of UAS Technologies to Enable Beyond Visual Line Of Sight (BVLOS) Operations". In: *International Conference on Vehicle Technology and Intelligent Transport Systems, VEHITS - Proceedings*. Vol. 2021-April. Science and Technology Publications, Lda, 2021, pp. 505–512. ISBN: 9789897585135. DOI: 10.5220/0010446905050512.

[47] S. Primatesta, G. Guglieri, and A. Rizzo. "A Risk-Aware Path Planning Strategy for UAVs in Urban Environments". In: *Journal of Intelligent and Robotic Systems: Theory and Applications* 95.2 (Aug. 2019), pp. 629–643. ISSN: 15730409. DOI: 10.1007/s10846-018-0924-3.

[48] X. Ren and C. Cheng. "Model of third-party risk index for unmanned aerial vehicle delivery in urban environment". In: *Sustainability (Switzerland)* 12.20 (Oct. 2020), pp. 1–15. ISSN: 20711050. DOI: 10.3390/su12208318.

[49] Rijksoverheid. *Waar mag ik niet vliegen met een drone?* 2025. URL: https://www.rijksoverheid.nl/onderwerpen/drone/vraag-en-antwoord/waar-mag-ik-niet-vliegen-met-een-drone.

[50] J. P. Rodrigue, C. Comtois, and B. Slack. *The geography of transport systems*. Taylor and Francis, Jan. 2016, pp. 1–440. ISBN: 9781317210108. DOI: 10.4324/9781315618159.

[51] T. A. Rodrigues et al. "Drone flight data reveal energy and greenhouse gas emissions savings for small package delivery". In: (Nov. 2021). DOI: 10.1016/j.patter.2022.100569. URL: http://arxiv.org/abs/2111.11463.

[52] T. A. Rodrigues et al. "Drone flight data reveal energy and greenhouse gas emissions savings for very small package delivery". In: *Patterns* 3.8 (Aug. 2022). ISSN: 26663899. DOI: 10.1016/j.patter.2022.100569.

[53] E. Rudnick-Cohen, J. W. Herrmann, and S. Azarm. "Risk-based path planning optimization methods for unmanned aerial vehicles over inhabited areas". In: *Journal of Computing and Information Science in Engineering* 16.2 (June 2016). ISSN: 15309827. DOI: 10.1115/1.4033235.

[54] A. T. Sage et al. *Testing the delivery of human organ transportation with drones in the real world Last-mile transportation of human donor lungs in a densely populated urban environment has been made pos-sible with drones*. Tech. rep. 2022. DOI: 10.1126/scirobotics.adf5798.

[55] J. R. Scalea et al. "An initial investigation of unmanned aircraft systems (UAS) and real-time organ status measurement for transporting human organs". In: *IEEE Journal of Translational Engineering in Health and Medicine* 6 (2018). ISSN: 21682372. DOI: 10.1109/JTEHM.2018.2875704.

[56] SDU UAS Center. *Sustainable Development Goals at SDU UAS Center*. 2025. URL: https://www.sdu.dk/en/forskning/sduuascenter/sdg.

[57] SEO. *Maatschappelijke Effecten van Drones*. Tech. rep. Decisio, 2021. URL: https://www.seo.nl/publicaties/maatschappelijke-effecten-van-drones/.

[58] Sesar. *European Drones Outlook Study Unlocking the value for Europe*. Tech. rep. Nov. 2016.

[59] Sesar. *U-space Concept of Operations (CONOPS)*. Tech. rep. 2023. URL: https://www.sesarju.eu/node/4544.

[60] P. Silveira, A. P. Teixeira, and C. Guedes Soares. "Ais based shipping routes using the dijkstra algorithm". In: *TransNav* 13.3 (Sept. 2019), pp. 565–571. ISSN: 20836481. DOI: 10.12716/1001.13.03.11.

[61] Single European Sky ATM Research 3 Joint Undertaking. "U-space : blueprint". In: (2017). DOI: doi/10.2829/335092.

[62] J. K. Stolaroff et al. "Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery". In: *Nature Communications* 9.1 (Dec. 2018). ISSN: 20411723. DOI: 10.1038/s41467-017-02411-5.

[63] E. Sunil et al. *Metropolis: Relating Airspace Structure and Capacity for Extreme Traffic Densities*. Tech. rep. 2015. URL: https://www.researchgate.net/publication/279845394.

[64] H. Tang et al. "UAV path planning based on third-party risk modeling". In: *Scientific Reports* 13.1 (Dec. 2023). ISSN: 20452322. DOI: 10.1038/s41598-023-49396-4.

[65] S. Verlinde. *Promising but challenging urban freight transport solutions: freight flow consolidation and off-hour deliveries*. Tech. rep. 2015.

[66] M. Viu-Roig and E. J. Alvarez-Palau. *The impact of E-Commerce-related last-mile logistics on cities: A systematic literature review*. Aug. 2020. DOI: 10.3390/su12166492.

[67] H. Zhang et al. "Research on Public Air Route Network Planning of Urban Low-Altitude Logistics Unmanned Aerial Vehicles". In: *Sustainability (Switzerland)* 15.15 (Aug. 2023). ISSN: 20711050. DOI: 10.3390/su151512021.

[68] H. Zhang et al. "UAV safe route planning based on PSO-BAS algorithm". In: *Journal of Systems Engineering and Electronics* 33.5 (Oct. 2022), pp. 1151–1160. ISSN: 16711793. DOI: 10.23919/JSEE.2022.000111.

[69] K. Zhou et al. "A risk-based unmanned aerial vehicle path planning scheme for complex air–ground environments". In: *Risk Analysis* (2024). ISSN: 15396924. DOI: 10.1111/risa.17685.

[70] Zipline. *Zipline Drone Delivery & Logistics*. 2025. URL: https://www.flyzipline.com/.

A

Airspace and infrastructure

This appendix provides a comprehensive overview of the airspace constraints and infrastructure elements relevant to the UAV corridor planning framework. Although the main study primarily focused on regions characterised by limited airspace complexity, this section describes the broader set of airspace categories and considerations relevant to the practical implementation and future extensions of the model. The data sources and cleaning procedures applied to derive the geospatial data of the areas presented here are described in Appendix B.1.

A.1. Airspace categories and UAV operational constraints

The following sections outline all airspace categories in Very Low Level (VLL) airspace as designated by Dutch authorities, explaining their implications for UAV corridor planning and how they were addressed in this study [38, 45, 38, 19].

Aviation categories

- CTR zones: Control zones around major airports (Schiphol, Rotterdam, Beek, Eelde, Lelystad) and military airfields where air traffic control manages all movements. While 'open' category flights are prohibited, 'specific' category operations may be authorized. UAV operations are feasible under strict conditions, as these zones primarily handle take-off and landing at higher altitudes, leaving much of the VLL airspace unused. Fixed and secured UAV corridors can operate here if air traffic control agreements are established. For instance, Rotterdam city center lies within CTR Rotterdam, but coordinated UAV corridors remain feasible with proper authorization and fixed routing.
- Regional Airports: Areas surrounding smaller airports with limited air traffic control. Similar to CTRs, these zones require 'specific' category privileges. UAV operations require special privileges and careful routing to avoid aircraft flight paths. Despite broad restriction zones (often 3000m radius), significant VLL space remains available under proper coordination.

Military categories

- General Training Areas (EHD): These are small zones (approximately 1.2 km radius) above military installations reserved for military exercises and training. Treated as strict no-fly zones due to safety and security concerns.
- Low-Flying Areas (GLV) and Routes: Areas where military aircraft can fly below standard minimum altitudes. UAV operations in OPEN A1/A2 categories are permitted up to 30 meters; A3 is prohibited. This study assumes corridors can be established below 30 meters here.
- Permanently Reserved Airspace (EHR): Military training areas requiring explicit permission for any flights. Examples include explosive ordnance disposal training grounds. Treated as no-fly zones.
- Temporarily Reserved Airspace (EHTRA): Temporary military training areas within a temporary

restricted airspace. It is a defined piece of airspace in which military flight activities take place at specific, announced times that may endanger civilian flights. All flights are prohibited unless permission is obtained from the military air traffic service provider. These zones are used intermittently and are often extensive. Therefore, similar treatment as CTRs is assumed, with the prerequisite of coordination with military air traffic control.

- Temporarily Segregated Airspace (EHTSA): This is a temporary segregated area, requiring reservation of airspace (EHTSA). It is intended for exclusive use by specific users for a predetermined time period. Again, all flights are prohibited unless permission is obtained from the military air traffic service provider. These areas are generally small and easier to avoid. However, for large zones such as EHTSA De Peel, covering populated areas like Deurne, Asten, and Venray, similar special arrangements to those in CTRs are required when delivery routes are necessary.

Security categories

- Restricted Zones (EHR): An area with restrictions for security reasons, typically a geographic zone around a secure facility or location with vital processes. Flights in the OPEN category are not allowed, and in the SPEC category special permission is required to enter the restricted area. These zones surround secure facilities or vital infrastructure. UAV operations in the OPEN category are prohibited, and SPEC category flights require special permission. As such, they are treated as no-fly zone.
- Prohibited Airspace (EHP): A defined airspace above the land areas or territorial waters of the Netherlands, within which the flight of aircraft is completely prohibited. This applies to areas such as Drakenstein Castle, royal palaces, government buildings, and Scheveningen. All flights are prohibited at all times under any circumstances. These zones should be treated as absolute no-fly areas.

Special operation categories

- Air Ambulance Landing Sites: Areas around landing sites reserved for trauma helicopters and rescue helicopters. Flights in the 'open' category are not allowed here. Treated as strict no-fly zones due to safety and emergency response priorities.
- Fire-fighting Training Areas:Training areas for helicopters specializing in external firefighting. OPEN A1/2 are allowed up to a max height of 30m above the ground; A3 not allowed. Although few in number and often located above natural corridors, these areas are excluded from routing to avoid operational interference. This may pose challenges in specific routes, such as between Leerdam and Houten.

Other categories

- High-Risk Areas: These zones typically cover hazardous industrial sites with a 500-meter safety radius and are treated as no-fly zones. Flights in the 'open' category are not allowed, and in the 'specific' category special permission is required to enter the restricted area. So, exceptions may be considered for key logistics hubs such as the ports of Rotterdam and Amsterdam.
- Protected Nature Reserves: Areas such as the Wadden Sea and Natura 2000 sites are fully excluded from UAV operations to prevent environmental disturbance. Additional protected areas not included in initial datasets were manually sourced and added from the national georegister.

A.2. Area types included in corridor model

The previous section detailed airspace categories relevant to UAV planning. This section outlines which ground infrastructure and sensitive areas are included in the corridor model for routing and risk assessment.

A.2.1. Linear and natural infrastructure

The following linear infrastructure and natural corridors were included based on spatial continuity, suitability for low-risk UAV operations, and availability in open geospatial datasets (corresponding OSM tags are listed in Table B). Policymakers can adjust these selections to align with local planning preferences.

1. Transportation infrastructure
    - Highways and major roads (motorways, trunks, primaries)
    - Regional/local roads (secondary, tertiary, residential)
    - Living streets
    - Tracks and rural access roads
    - Pedestrian/cycling paths (footways, pedestrian zones, cycle paths, bridleways)
    - Railways (heavy rail, light rail, trams)
2. Energy and industrial infrastructure
    - Power lines
    - Power generation plants (e.g., solar, gas, waste-to-energy)
    - Communication towers and masts
    - High infrastructures (silos, chimneys, water towers, wind turbines, windmills)
3. Land use types
    - Industrial zones
    - Commercial zones (offices, business parks)
    - Retail and shopping areas
    - Residential areas
    - Recreational grounds (sports fields, playgrounds)
    - Agricultural land (farmland, orchard)
4. Water and natural corridors
    - Rivers, streams, and canals
    - Lakes, reservoirs, and ponds
    - Wetlands
    - Forests, woodlands, and scrubland
    - Meadows and open grass areas
5. Other relevant areas
    - Public parks
    - Cemeteries

A.3. Manually selected sensitive areas

Certain public and sensitive facilities are not strict no-fly zones but are treated as high-risk areas to be avoided in corridor routing. This list is preliminary and should be expanded by policymakers using this framework.

- Schools, universities and kindergartens
- Hospitals and healthcare centres
- Prisons
- Religious institutions (churches, mosques, temples)
- Cultural heritage sites and landmarks

B

Model Creation

This appendix provides detailed documentation of the data collection, preprocessing, and model construction processes used in this study. It covers the technical implementation of the UAV corridor network generation, including all data sources, cleaning procedures, and graph construction algorithms.

B.1. Data collection and analysis

The model development required integration of multiple geospatial datasets and extensive preprocessing to ensure data quality and compatibility. All data collection, cleaning, and analysis procedures were implemented using Python notebooks, which are publicly available in the accompanying GitHub repository at `https://github.com/catomartens/thesis_cf_martens`. The repository includes comprehensive documentation and a README file with step-by-step instructions for reproducing the analysis. The following subsections detail each data source and its preprocessing pipeline.

B.1.1. UAS geozone data for airspace constraints

The current dataset for UAS geozones is retrieved from the official website of the Dutch government [49]. The dataset, titled *"Alle statische zonering en andere gebieden in ED269 format - 22juli2024"*, is provided as a JSON file. Since this dataset lacks complete Natura 2000 area coverage, it was supplemented with a separate Natura 2000 geospatial dataset retrieved from Nationaal Georegister [41]. Both datasets were imported and processed into a unified GeoPandas dataframe. The original JSON dataset contains geozones with attributes including identifier, country, name, type, restriction, reason, zoneAuthority details, applicability, message, and geometry specifications (upper/lower limits, vertical references, and horizontal projections). The notebook get_data.ipynb guides the user through each of the steps.

After processing, the unified dataframe includes the following columns:

- **name**: Descriptive name of the geozone
- **description**: Additional information about the zone (renamed from 'message')
- **air_type**: Specific airspace classification (e.g., CTR Zones, Protected Nature Reserves, Military Training Areas)
- **category**: Broader categorization (Aviation, Military, Environment, Security, Industry, Special Operations)
- **geometry**: Spatial representation of the geozone
- **no_fly**: Boolean indicating whether the zone is a no-fly area for drones

Zones with lower altitude limits exceeding 120 meters were filtered out as they fall outside the Very Low Level (VLL) airspace relevant for drone operations. Circle geometries were converted to polygons using their radius values for consistent spatial analysis.

B.1.2. Ground based data

The ground-based infrastructure data was retrieved using the Python package OSMnx through a structured multi-step process in `get_data.ipynb`. The data collection was organized into seven distinct categories, each targeting specific types of infrastructure and land use relevant to UAV corridor planning:

- Transportation infrastructure: Roads, railways, and paths
- Energy and industrial infrastructure: Power lines, plants, and communication towers
- Land use areas: Residential, commercial, industrial, and agricultural zones
- Natural areas: Forests, meadows, and wetlands
- Water features: Rivers, canals, lakes, and ponds
- Public facilities: Schools, hospitals, religious sites, and cultural locations
- PostNL locations: Post offices, post boxes, and distribution points

For each category, specific OpenStreetMap tags were used to query the relevant features, as shown in Table B.1. The extraction process followed these steps:

1. **Boundary definition**: Study areas were defined by specifying one or more municipalities (e.g., 'Alphen aan den Rijn, Zuid-Holland, Netherlands'). Multiple boundaries could be combined to create larger study regions.
2. **Category-specific queries**: For each infrastructure category, the OSMnx `features_from_place()` function was called with the appropriate tag dictionary to retrieve all matching features within the boundaries.
3. **Data standardization**: Each retrieved dataset was processed to retain only the essential columns: `geometry`, `name`, `id`, `description`, `area_type`, and `category`. Missing columns were added with null values to ensure consistency across all datasets.
4. **Area type classification**: A comprehensive classification function (`determine_area_type_from_row()`) was applied to map specific OSM tags to meaningful area types. This function implements detailed logic to categorize infrastructure based on its attributes.
5. **Data consolidation**: All category-specific datasets were concatenated into a single GeoPandas DataFrame, with duplicate entries removed based on geometry, name, id, and description.
6. **Export**: The final consolidated dataset was exported as a GeoJSON file for use in the corridor network construction phase.

Since not all socially sensitive locations are well-represented in OpenStreetMap (OSM) as precise polygons, certain modelling assumptions were necessary. Specifically, for features available only as point geometries, a circular buffer of 50 meters was applied to approximate a conservative exclusion radius. This approach provides a safety margin intended to capture both privacy considerations and societal acceptance thresholds, particularly in densely built or human-centric environments.

The extraction process also revealed data availability limitations. Aviation infrastructure (aerodromes, helipads, runways) was notably absent in several study areas, resulting in failed queries for this category. Additionally, some infrastructure types may be underrepresented in OSM, requiring supplementary data sources for comprehensive coverage.

The final GeoPandas DataFrame contains the following fields for each area or object:

- geometry: The spatial representation of the feature (Point, LineString, Polygon, or MultiPolygon)
- name: The name of the feature, serving as a description
- id: The unique OpenStreetMap identifier
- description: Additional descriptive information
- area_type: The classified type based on the mapping function

- category: The high-level category (Transportation, Energy/Industrial, Land use, Nature, Water, Public facility, or PostNL)

Table B.1: Mapped area types and their corresponding OpenStreetMap tags

Area Type	Representative OSM Tags
Motorways and major roads	`highway=motorway, trunk, primary`
Regional roads	`highway=secondary, tertiary`
Tracks and rural access roads	`highway=track, unclassified, service`
Living and residential streets	`highway=living_street, residential`
Pedestrian and cycling paths	`highway=footway, cycleway, path, bridleway`
Railways	`railway=rail, light_rail`
Bridges	`bridge=yes`
Power lines	`power=line, power=tower`
Power plants	`power=plant`, `man_made=waste_water_plant`
Communication towers	`man_made=communications_tower, mast`
Wind turbines	`man_made=windmill, wind_turbine`
High infrastructures	`man_made=silo, chimney, tank, water_tower`
Industrial zones	`landuse=industrial`
Commercial zones	`landuse=commercial`
Retail zones	`landuse=retail`
Residential areas	`landuse=residential`
Recreational zones	`landuse=recreation_ground`
Agricultural lands	`landuse=farmland, orchard`
Forests and woodlands	`landuse=forest, wood`
Meadows and open grass	`landuse=grass, grassland, meadow, scrub`
Rivers, canals, streams	`waterway=river, stream, canal`
Lakes and ponds	`water=lake, pond`
Water reservoirs	`water=reservoir`
Wetlands	`natural=wetland`, `wetland=marsh, bog, fen`
Schools and universities	`amenity=school, kindergarten, university`
Hospitals	`amenity=hospital`
Prisons	`amenity=prison`
Religious sites	`amenity=place_of_worship`
Cultural sites	`historic=*, tourism=attraction`
Cemeteries	`landuse=cemetery`
Parks	`leisure=park`

B.1.3. Distribution center data processing

The distribution center dataset was obtained via Schuberg Philis and contains information about PostNL's logistics facilities. The original CSV file included various depot types and operational statuses. The data was filtered to include only active facilities (depotLive = True) specifically designated as parcel sorting centers (PostNL Sorteercentrum Pakketten). The processed dataset retains the following attributes:

- depotAbbreviation: Short code identifying each distribution center
- depotCity: Municipality where the facility is located
- id: Unique identifier for each depot (renamed from depotId)
- geometry: Point geometry created from GPS coordinates
- type: Classification as 'distribution' center

The cleaned data was converted to a GeoDataFrame with WGS84 (EPSG:4326) coordinate reference system and exported as a GeoJSON file for integration with the routing network analysis.

B.2. UAV corridor network construction

This section presents the methodology for constructing a graph-based representation of feasible UAV corridors connecting PostNL distribution centers to delivery points. The network integrates multiple spatial data sources into a unified routing graph that respects both operational requirements and airspace restrictions.

B.2.1. Data preprocessing pipeline

The graph creation process requires three sequential preprocessing steps to prepare the spatial data. First, geospatial data of all areas and PostNL locations are collected (from `get_data.ipynb`), as described in Section B.1.2. This initial dataset contains all linear infrastructure (roads, railways), natural corridors (waterways, green spaces), and PostNL delivery points for the study area.

After this, risk scores are calculated and assigned to each area type based on a predefined risk assessment framework (`assign_risk.ipynb`). The risk calculation uses weighted factors for: fatality risk (α_f = 0.4), property damage risk (α_p = 0.3) and societal disruption risk (α_s = 0.3). Each area type receives a normalized risk score between 0 and 1, with PostNL points assigned a risk value of 0. The risk scores incorporate five external risk factors (L1-L5) and severity ratings for each domain.

Finally, the no-fly zones need to be integrated in the data set, which is done in `assign_no_fly_zones.ipynb`. No-fly zones relevant to the study area are extracted from the national dataset and added to the spatial data. These zones are filtered based on geographic intersection with the study area boundaries and assigned a 'no_fly_zone' risk classification. The final preprocessed dataset combines all spatial features with their associated risk values, ready for graph construction in the routing analysis.

B.2.2. Network graph construction

To construct the network capturing all feasible paths between distribution centers and PostNL pickup points, an eight-step procedure is followed (including data loading and validation). Each step systematically builds upon the previous one to create a comprehensive routing graph. The full implementation can be found in the accompanying GitHub repository, with the notebook `get_graph.ipynb` guiding users through the entire workflow.

Step 0: Data Loading and Preparation

The process begins by loading the preprocessed GeoJSON file using the `read_geojson` function. This function:

- Reads the combined spatial data file containing all geometries with assigned risk values
- Converts the coordinate system to EPSG:28992 (Dutch RD New) for accurate distance calculations
- Separates the data into five distinct GeoDataFrames:
  - `lines_gdf`: LineString geometries representing linear infrastructure
  - `polygons_gdf`: Polygon geometries representing open areas
  - `post_nl_gdf`: Point geometries for PostNL pickup locations
  - `distribution`: Distribution center locations filtered by depot city
  - `no_fly_zones_gdf`: Polygon geometries marked as no-fly zones
- Ensures risk values are numeric (float) and fills missing values with 0
- Sets height and risk to 0 for all PostNL-related points

The depot parameter accepts a list of depot city names to accommodate distribution centers that may serve multiple municipalities.

Step 1: Initializing the Graph with PostNL Nodes

The graph structure is initialized using `create_graph_with_nodes`, which processes the `post_nl_gdf` containing both distribution centers and pickup points. The function:

- Creates an empty NetworkX graph
- Adds distribution centers (area_type = 'distribution') as nodes with:
  - Node ID from the 'id' column (or DataFrame index if missing)
  - `geometry`: Point coordinate
  - `ntype`: 'distribution'
  - `risk`: 0
- Adds PostNL pickup points (area_type = 'postnl point') as nodes with:
  - Node ID from the DataFrame index
  - `geometry`: Point coordinate
  - `ntype`: 'postnl'
  - `risk`: 0
- Validates that all geometries are Point types (non-Point geometries are logged and skipped)
- Reports summary statistics of nodes added vs. skipped

Step 2: Adding Linear Infrastructure as Edges
In the second step, all relevant linear infrastructure from OpenStreetMap (OSM) is integrated into the graph as edges. These consist of road segments, railways, paths and other eligible `LineString` or `MultiLineString` geometries in the GeoDataFrame (`lines_gdf`). The aim is to transform these infrastructure elements into a coherent edge network representing all feasible UAV corridors in the study area. The function `add_linesstrings_to_graph` executes this step in four phases:

1. **Geometry processing and intersection detection:** All `MultiLineStrings` are first exploded into individual `LineStrings` to simplify processing. A spatial index (`STRtree`) is then used to identify pairwise intersections between all line segments. Intersections are merged with a snap tolerance of 5 metres using `cKDTree`, based on the assumption that UAVs can reasonably treat endpoints within this distance as spatially identical. This tolerance (`snap_tolerance = 5`) helps resolve minor misalignments in OSM data. Coordinates are rounded to 3 decimal places for consistency.
2. **Networking and edge segmentation:** Each `LineString` is split at the identified intersection points using Shapely's `split` function, resulting in smaller segments that can be added individually as graph edges. Nodes are created at each segment endpoint and labelled as `line_intersection` (for snapped intersection points) or `waypoint` (if newly introduced). Each valid segment is added to the graph as an edge with the following attributes:
   - `geometry`: the corresponding `LineString`
   - `length`: the physical length of the segment
   - `etype`: the infrastructure type (from area_type)
   - `risk`: the pre-assigned risk value from the GeoDataFrame
   - `height`: the designated flight altitude for this segment
   - `name`: optional name attribute if present
3. **Dealing with segments without intersections:** `LineStrings` that do not intersect other lines are added directly as single edges. These are also rounded to 3 decimal places to ensure spatial consistency and are given identical attributes as above.
4. **Quality control:** The function tracks and reports various statistics including number of lines processed, intersections found, nodes added, edges created, and segments skipped due to invalidity or excessive shortness (< 0.001 units).

Step 3: Integrating Polygon-Based Areas into the Corridor Network
In the third step, all polygon-based areas suitable for UAV routing are integrated into the graph. These include open areas with land use and other spatially contiguous regions, as defined in the GeoDataFrame (`polygons_gdf`). Unlike linear infrastructure, these areas are modelled using grid-based internal structures and sampled boundaries to enable safe, efficient routing across and around these zones. The `add_grid_polygons_to_graph` function performs this step in eight sub-steps, grouped into three main stages:

1. **Grid generation and polygon filling:** For each polygon not classified as a no-fly zone, a uniform grid with 50-metre cells (`grid_size = 50`) is generated. Cells that intersect the polygon are retained and a node is placed at the centroid of the intersected area. These nodes are labelled as `grid_point`. Edges are created between neighbouring grid points, provided they do not intersect the polygon boundary. All nodes and edges inherit the `risk`, `height`, and `etype` attributes of the source polygon. A maximum of 10,000 cells per polygon is enforced to maintain computational efficiency.
2. **Boundary sampling and connection:** Each polygon boundary is sampled every 50 metres (`boundary_sample_distance = 50`) to form a set of `polygon_boundary` nodes. These are connected to each other to represent the perimeter of the polygon. Grid points within the polygon are then connected to the nearest boundary node if the distance is within 35 metres (`snap_tolerance = 35`). These connections preserve accessibility between internal and perimeter nodes. The nearest boundary connections are efficiently computed using `cKDTree` spatial indexing.
3. **Integration with existing network:** To connect the polygon-based subgraphs to the existing linestring-based network:
   (a) All pairwise intersections between existing edges are detected using a spatial index (`STRtree`). Intersections are snapped using a 2-metre tolerance to avoid near-duplicate nodes.
   (b) New nodes are introduced at these snapped intersections, labelled as `poly_intersection`, and existing edges are split accordingly by removing the original edge and adding two new segments.
   (c) Finally, all `line_intersection` and `waypoint` nodes are connected to the nearest `polygon_boundary` node, provided the distance is less than 15 metres (`max_connection_distance = 15`). These connector edges are labelled as `etype = connector`, with a fixed `height = 30`.

Step 4: Connecting Distribution and PostNL Nodes to the Network
In the fourth step, all postnl and distribution nodes are connected to the corridor network by linking them to the nearest edge. This step ensures that all origin and destination nodes are reachable within the graph structure. The function `connect_to_nearest_edge` handles this process as follows:

1. For each node of type `postnl` or `distribution`, the nearest edge in the graph is identified based on Euclidean distance. Only edges represented by valid `LineString` geometries are considered. The projection point on the edge is calculated using Shapely's `project` and `interpolate` functions.
2. If the nearest edge is within the maximum connection distance (set as `max_connection_distance = 10` meters), a new `waypoint` node is inserted at the closest projection point on the edge. The original edge is then split into two new edges, each connecting one endpoint of the original edge to the new waypoint node.
3. Finally, a connector edge is added between the original `postnl` or `distribution` node and the newly inserted waypoint. This edge is assigned `etype = postnl_connector` and a fixed `height = 0`, indicating its sole purpose is spatial linkage rather than an operational corridor segment.

Nodes beyond the maximum allowed distance (10 meters) are skipped and reported to prevent unrealistic routing assumptions. This step guarantees that all service locations are embedded within the navigable corridor graph.

Step 5: Enforcing No-Fly Zone Restrictions

In the fifth construction step, all edges that fall within, intersect, touch, or cross designated no-fly zones are removed from the graph. This ensures full regulatory compliance and enforces strict avoidance of restricted airspace in the resulting UAV corridor network. The function `remove_no_fly_zones_from_graph` performs this operation in three stages:

1. All valid edges in the graph are extracted into a GeoDataFrame with CRS EPSG:28992. Only edges with LineString geometries are retained, and their spatial properties (length, geometry, and risk) are preserved.
2. The function `get_lines_inside_no_fly_zones` checks each edge for spatial overlap with the no-fly zone polygons using a spatial index (`STRtree`). An edge is flagged for removal if it is fully contained in (`within`), intersects, crosses, or touches any no-fly zone geometry.
3. All flagged edges are removed from the graph based on exact geometry matches using Well-Known Binary (WKB) representation to ensure precision.

Step 6: Graph Serialization

After the no-fly zone enforcement, the graph is saved as a pickle file with suffix `_raw.pkl` for efficient storage and potential debugging of the unvalidated network structure.

Step 7: Visualization (Optional)

The graph can be visualized using the `plot_graph_on_folium` function to create interactive web-based maps. This visualization overlays the network on OpenStreetMap tiles and can optionally display waypoint nodes. This step is commented out by default to avoid dependency issues.

Step 8: Validation and Height Completion

As a final step, the graph undergoes validation and attribute completion using two functions:

1. `diagnose_graph`**:** Performs comprehensive validation including:
    - Graph connectivity analysis (reports if fully connected or number of components)
    - Node type (ntype) distribution counting
    - Edge type (etype) distribution counting
    - Missing attribute detection for nodes (geometry, ntype) and edges (geometry, length, risk, height, etype)
2. `fill_missing_edge_heights_from_csv`**:** If height values are missing:
    - Loads a predefined mapping from `all_edge_types_with_properties.csv`
    - Creates a lookup table mapping area_type to default height values
    - Fills missing edge heights based on their etype attribute
    - Reports the number of edges patched and any unmapped edge types

The validated and complete graph is then saved as the final pickle file (without the _raw suffix) for use in routing analysis.

B.3. Model assumptions, user-defined parameters, and sources of uncertainty

The corridor construction model is built on a series of structural assumptions, configurable parameters, and known sources of uncertainty. These are discussed in this section to provide transparency and clarify the flexibility and limitations of the current setup.

Modelling Assumptions

The model relies on the following assumptions to ensure tractability, consistency, and regulatory compliance:

- Only `Point` geometries are allowed for PostNL and distribution nodes; non-point geometries are excluded from the graph.
- No-fly zones are treated as strict exclusions: all edges intersecting these areas are removed entirely.
- PostNL and distribution nodes are assigned a risk value of 0, assuming that unreachable or high-risk locations will be filtered out via edge removal.
- All `MultiLineString` geometries are exploded into individual `LineString` segments prior to graph insertion.
- Intersections between lines are snapped together using a 5-meter tolerance.
- Polygons are filled using a uniform 50-meter grid to balance resolution and computational cost.
- Polygon boundaries are sampled at 50-meter intervals to allow perimeter traversal.
- Grid points are only connected to boundary nodes if within a 35-meter threshold.
- Intersections between line and polygon edges are snapped within 2 meters to prevent duplicate nodes.
- Connectors between polygon boundaries and intersecting lines are only created if the distance is less than 15 meters.
- Connector edges between PostNL nodes and the main graph are only added if the node is within 10 meters of an edge.
- Missing height values are imputed using a lookup table based on the edge's `etype`.
- Coordinates are rounded to 3 decimal places for consistency across all operations.
- The graph uses an undirected structure (NetworkX Graph, not DiGraph).
- All spatial operations use EPSG:28992 (Dutch RD New) coordinate system.

User-Defined Parameters

Within the structure of these assumptions, several key parameters have been manually set. These parameters are summarized in Table **??** and were used for the default model runs presented in this study. Each parameter can be easily adjusted in future research to explore alternative corridor configurations or sensitivities.

Table B.2: User-defined parameters in the UAV corridor graph model

Parameter	Value	Description
`grid_size`	50 m	Cell size used to create grid points within polygons
`boundary_sample_distance`	50 m	Distance between sampled points along the polygon boundary
`snap_tolerance`	5 m	Tolerance for merging nearby intersection points from line geometries
`snap_tolerance`	35 m	Maximum distance to connect grid points to nearby boundary points
`snap_tolerance`	2 m	Snapping threshold for connecting line and polygon edges at intersections
`max_connection_distance`	10 m	Maximum distance allowed to connect PostNL or distribution nodes to existing edges
`max_connection_distance`	15 m	Maximum distance allowed to connect line/polygon nodes to polygon boundaries

Sources of Uncertainty

The model also contains several unavoidable sources of uncertainty, related to both data quality and methodological choices:

- Accuracy of OSM geometries: roads, paths, and polygons may be misaligned or inconsistently digitized.
- No-fly zone data completeness: The national dataset may not include all temporary or recently updated restrictions.
- Grid discretisation error: the use of 50-meter grid cells may overlook narrow features or fine spatial patterns.
- Maximum grid cells limit: The 10,000 cell cap per polygon may underrepresent very large areas.
- Snapping tolerances: merging points within a fixed radius introduces small geometric approximations.
- Edge splitting precision: Shapely's split function may introduce small gaps or overlaps at intersection points.
- Fixed height assignments: assumed altitudes may not fully align with regulatory or operational practices.
- Connection thresholds: rigid distance limits (e.g., 10 meters) may exclude otherwise feasible corridor connections.
- Uniform boundary sampling: regular perimeter spacing may underrepresent sharp corners or complex geometries.
- Uncertainty in risk and height attributes: both rely on external datasets and simplified mappings per area type.

C

Routing Algorithm Implementation

This appendix provides an overview of how the routing algorithm is implemented in the `notebook_find_paths` notebook which is publicly available in the accompanying GitHub repository at `https://github.com/catomartens/thesis_cf_martens`. This notebook is used to find optimal UAV corridors between origins and destinations in a Graph created based on geospatial data.

C.1. Routing workflow

The path finding process follows a structured workflow implemented across several key steps:

Graph preparation

The routing begins by loading the preprocessed graph (saved as a pickle file, and created with `get_graph.ipynb`) that contains the complete corridor network with all nodes, edges, and their attributes. The graph includes distribution centers, PostNL pickup points, and the connecting infrastructure network with assigned risk values and operational parameters.

Distribution center connectivity

In some cases, distribution centers may not be optimally connected to the main network due to the conservative 10-meter connection threshold used during graph construction. The notebook includes a manual connection function that allows these distribution points to be connected to suitable edges within a larger radius (up to 1000 meters) when necessary. This ensures all distribution centers can serve as valid origins for route planning.

Cost function application

Before executing the routing algorithm, the weighted cost function is applied to all edges in the network. This function combines third-party risk and energy consumption based on the selected α parameter. The energy model accounts for both horizontal travel distance and vertical transitions between the 30-meter and 60-meter flight levels.

C.1.1. Path computation

For each PostNL delivery point, the algorithm evaluates potential routes from all available distribution centers using Dijkstra's shortest path algorithm. The distribution center yielding the lowest total cost path is selected as the optimal origin for that specific delivery point. This ensures each PostNL location is served by its most suitable distribution center considering the risk-efficiency trade-off.

C.2. Performance metrics

The routing implementation calculates comprehensive performance metrics for each computed path:

- **Path length**: Total horizontal distance traveled in meters

- **Risk exposure**: Cumulative third-party risk along the route
- **Energy consumption**: Total energy required including horizontal and vertical components
- **Route complexity**: Number of turns and altitude changes
- **Infrastructure usage**: Types of corridors utilized (roads, railways, waterways, etc.)

These metrics are collected in a structured DataFrame that enables systematic analysis of routing outcomes across different α values.

C.3. Multi-scenario analysis

The notebook supports evaluation across multiple scenarios by varying the risk-efficiency trade-off parameter α:

- $\alpha = 0$: Pure efficiency focus, minimizing energy consumption
- $\alpha = 1$: Pure safety focus, minimizing third-party risk exposure
- Intermediate values: Balanced approaches exploring the trade-off spectrum

For each α value, the algorithm recomputes all distribution-to-PostNL routes, enabling the construction of Pareto frontiers that visualize the relationship between safety and efficiency objectives.

C.4. Visualization capabilities

The routing results can be visualized through several complementary approaches:

C.4.1. Corridor usage maps

The notebook generates maps showing which infrastructure segments are utilized for different α values. These visualizations use color coding to distinguish between different corridor types (motorways, railways, parks, etc.) and highlight how route preferences shift as safety priorities change.

C.4.2. Distribution area coverage

Service area maps illustrate which PostNL points are served by each distribution center, revealing how the optimal assignment changes based on the risk-efficiency balance. This helps identify potential capacity imbalances or service area overlaps.

C.4.3. Battery range analysis

The notebook includes functionality to assess service coverage under different battery capacity constraints. By analyzing the energy requirements of computed routes, it determines which delivery points remain accessible for various battery configurations, supporting practical deployment planning.

C.5. Output generation

The routing analysis produces several outputs for further analysis:

- Summary statistics tables comparing key metrics across different α values
- Detailed path geometries for geographic analysis and corridor planning
- Infrastructure usage frequency data to identify critical corridor segments
- LaTeX-formatted tables for direct inclusion in reports and publications

These outputs enable comprehensive evaluation of the corridor network's performance under different operational priorities and support evidence-based decision-making for UAV corridor implementation.

D

Risk Analysis

This appendix presents a comprehensive risk analysis framework for UAV operations in diverse operational environments. The analysis employs a structured approach to assess and quantify risks associated with UAV flights across different area types, providing a foundation for informed decision-making in U-space operations. The method is onderbouwt en discussed also in 4.

This appendix presents a comprehensive risk analysis framework for UAV operations in various operational environments. The analysis employs a structured approach to assess and quantify risks associated with UAV flights across different area types. The underlying methodology and theoretical foundations are thoroughly substantiated and discussed in Section **??**. This appendix focuses on the practical application of that methodology, presenting the detailed attribution of severity and likelihood scores for each area type and risk factor. These scores form the basis for calculating overall risk ratings that guide operational planning.

D.1. Severity

Severity is assessed in terms of the potential impact on humans and the surrounding environment. The methodology of Allouch et al. [1] introduces a four-level qualitative categorization of severity: catastrophic, critical, marginal, and negligible.

- Catastrophic: A catastrophic event involves serious injury or loss of human life and must be addressed with the highest priority.
- Critical: A critical event causes significant damage to third-party assets or infrastructure, but does not pose a direct threat to human safety.
- Marginal: A marginal event results in damage to the drone system itself without affecting external entities.
- Negligible: A negligible event has minimal operational impact, potentially causing minor degradation in system performance without leading to mission failure.

To ensure consistency with existing UAV risk literature, this study adopts a severity structure inspired by common components in third-party risk (TPR) models. Pang et al. [44] assign risk scores based on three components: fatality risk, property damage risk, and societal impact (including noise and privacy). Zhou et al. [69] similarly construct a detailed site-based model incorporating fatality risk, property risk, and noise impact. Tang et al. [64] apply a comparable structure, defining risk through obstacle risk, fatality risk, and property loss. Hu et al. [21] include risks to people, vehicles, and manned aircraft in urban UAV operations, while Zhang et al. [68] compute risk per grid cell based on crash probability, UAV impact characteristics, and population density.

Building on these approaches, this thesis standardizes severity assessment across three domains: (1) fatality risk, (2) risk of property damage, and (3) societal impacts. Each domain is linked to a proxy

variable that enables comparison of severity across area types. These proxies and their implementation are detailed in the sections below.

D.1.1. Fatality severity

Fatality severity refers to the risk that a UAV crash causes direct injury or fatality to either pedestrians, cyclists or individuals in vehicles. In this study, fatality severity is assessed by considering the likely physical consequences of a UAV impact in a given area, specifically: are people physically present in the area, and can they be directly struck by a crashing UAV? The categorisation of severity levels fatality risk per area type is detailed in Table D.1.

For example:

- Wind turbines are considered Negligible. No people are usually present in these areas under normal operating conditions. UAVs flying over or near turbines do not pose a significant risk of affecting people, as maintenance areas are limited and human presence is rare.
- Residential areas are considered Marginal. Although people live here, most are indoors and are not directly exposed to drones flying overhead. UAVs flying over rooftops are unlikely to hit people.
- Commercial zones are assessed as critical to property severity. Pedestrian and vehicle traffic is frequent in these areas, such as business districts. People often move between buildings or along sidewalks, so there is a risk of direct impact from UAVs, although the density is still lower than, for example, schools.
- Living and residential streets and schools are classified as Catastrophic. These locations have constant and concentrated pedestrian activity, including children, bicyclists and other vulnerable road users. A UAV crash in these areas is likely to hit someone directly, making the potential consequences severe.

Table D.1: Standardized fatality severity scores per area type

Severity Level	Interpretation	Applicable Area Types
1 – Negligible	Virtually no human presence under normal conditions which can be directly hit by the UAV.	Power lines, Power plants, Communication towers, High infrastructures, Agricultural lands, Forests and woodlands, Meadows and open grass, Rivers, canals and streams, Water reservoirs, Wetlands
2 – Marginal	Occasional human presence which can be directly hit by the UAV; limited activity	Tracks and rural access roads, Industrial zones, Residential areas, Lakes and ponds, Hospitals, Prisons, Religious sites
3 – Critical	Frequent but less dense human presence; regular pedestrian or vehicle traffic which can be directly hit by the UAV.	Regional roads, Living and residential streets, Pedestrian and cycling paths, Commercial zones, Schools and universities, Cemeteries
4 – Catastrophic	Constant and dense human presence throughout the day, which can be directly hit by the UAV.	Motorways and major roads, Railways, Retail zones, Recreational zones, Cultural sites, Parks

D.1.2. Property severity

Property severity reflects the level of non-human damage that may result from a UAV crash, specifically the potential for impact on infrastructure, buildings, or essential services. This includes both the financial cost of property damage and the operational consequences that such damage may cause. The categorisation of severity levels of property risk per area type is detailed in Table D.2.

For example:

- Open meadows are considered Negligible, as a UAV crash in these areas is unlikely to cause any third-party property damage, only to itself.
- Tracks and rural access roads are considered Marginal. While some property is present, damage would typically be minor and isolated, such as to a fence.
- Communications towers are assessed as critical to property severity because they offer high-value infrastructure essential for mobile networks, broadcasting or emergency communications. A UAV crash can damage sensitive equipment and lead to local service interruptions, with significant financial and operational consequences.
- Power plants and power lines are classified as Catastrophic due to their high concentration of critical infrastructure. Damage may lead to widespread power outages, cascading technical failures, and indirect safety or security consequences.

Table D.2: Standardized property severity scores per area type

Severity Level	Interpretation	Applicable Area Types
1 – Negligible	Little or no third-party property is present; any damage would have minimal external impact.	Forests and woodlands, Meadows and open grass, Meadows and open grass, Lakes and ponds, Wetlands
2 – Marginal	Some properties of low value or sparsely scattered are present; damage is unlikely to cause broader disruption.	Tracks and rural access roads, Pedestrian and cycling paths, Recreational zones, Agricultural lands, Water reservoirs, Cemeteries, Parks
3 – Critical	Significant presence of third-party assets; damage may result in local operational, financial or logistical consequences.	Motorways and major roads, Regional roads, Railways, Communication towers, High infrastructures, Commercial zones, Hospitals, Schools and universities, Prisons, Religious sites, Cultural sites
4 – Catastrophic	High concentration of sensitive or valuable property; damage may result in broader disruption, costs or cascading effects with indirect security or safety implications.	Living and residential streets, Power lines, Power plants, Industrial zones, Retail zones, Retail zones

D.1.3. Societal severity

Societal severity in this study refers to the risk of public disturbance or lack of acceptance associated with UAV operations in a given area. This includes concerns related to privacy, noise, and social discomfort caused by the visible or audible presence of drones. Unlike fatality severity, which considers physical harm from a crash, societal severity focuses on how individuals perceive and experience UAV activity, even in the absence of direct risk. For example, drones flying over private gardens or residential areas may raise strong objections due to privacy concerns or audible disruption, particularly during quiet hours. Conversely, a drone passing unnoticed over remote woodland is unlikely to provoke public concern, even if people are present. The severity assessment is thus informed by population presence, but interpreted through the lens of social sensitivity, not physical exposure. The categorisation of severity levels of societal risk per area type is detailed in Table D.3. *Note: Environmental consequences niet meegenomen*

For example:

- Rivers, canals, and streams are considered Negligible. UAVs flying over open water are rarely seen or heard, resulting in minimal public disturbance or awareness.
- Railways are considered Marginal. Although trains may carry passengers, the surrounding space is not typically occupied by the public.
- Cultural sites are considered Critical. These are public spaces with regular foot traffic and a

moderate level of social or symbolic sensitivity. Drone presence may be noticed and perceived as intrusive or inappropriate.

- Residential areas are considered Catastrophic. Drones flying over private homes or gardens are highly likely to trigger objections due to privacy concerns, noise, and the perception of surveillance. These areas are socially sensitive, and UAV presence is often considered unacceptable by residents.

Table D.3: Standardized societal severity scores per area type

Severity Level	Interpretation	Applicable Area Types
1 – Negligible	Very low or no population presence; the drone is unlikely to be seen or heard and disturbance is minimal or non-existent.	Motorways and major roads, Tracks and rural access roads, Power lines, Communication towers, Agricultural lands, Forests and woodlands, Meadows and open grass, Rivers, canals and streams, Lakes and ponds, Water reservoirs, Wetlands, Natura2000 areas
2 – Marginal	Low sensitivity; people may notice drones but are unlikely to have major concerns.	Regional roads, Railways, Power plants, High infrastructures, Industrial zones
3 – Critical	Moderate to high social sensitivity; drone activity is likely to be noticed and possibly cause complaints or discomfort.	Pedestrian and cycling paths, Commercial zones, Retail zones, Recreational zones, Cultural sites, Parks, Prisons
4 – Catastrophic	High potential for public disruption and objection; dense residential areas or sensitive social areas where the presence of drones is very intrusive.	Living and residential streets, Residential areas, Schools and universities, Hospitals, Religious sites, Cemeteries

D.2. Likelihood

The probability of a risk is defined as the likelihood of the risk occurring. Based on the analysis of Allouch et al. [1], likelihood is described through five ordinal categories: frequent, probable, occasional, remote, and improbable.

- Frequent: A frequent event occurs repeatedly or has been observed often.
- Probable: A probable event is expected to occur regularly, though not as often as frequent ones.
- Occasional: An occasional event occurs sporadically. A remote event is unlikely but possible,
- Remote: Unlikely to occur but possible or has occurred rarely
- Improbable: Event is extremely rare or has never been observed.

Each area is systematically assessed for all relevant external risk factors. For each factor, a probability score is assigned based on the expected frequency of occurrence under local spatial conditions. This results in a structured set of factor-specific probability scores per area. For each external risk factor, a standardization is assigned to the probability of the risk profile in an area. described in the sections below:

D.2.1. External risk 1: Obstacles

Obstacle risk acan be defined as the presence of both fixed obstacles (e.g., trees, power lines, buildings, bridges) and dynamic obstacles (e.g., birds, vehicles, cyclists) that may interfere with UAV operations [1]. In this study, the likelihood score reflects how often such obstacles are present along the UAV's route and how dense or unavoidable they are within a given area type. Importantly, this metric reflects the structural and operational complexity of the environment, despite the flight altitude of the UAV.

To standardize obstacle likelihood assessment, this study assigns each area type to one of five likelihood categories, ranging from Improbable (1) to Frequent (5). While these scores are not derived from empirical obstacle measurements, they are based on approximations on observable spatial characteristics such as infrastructure density, land use, and known activity patterns. To enhance interpretability and support future standardization, each score is linked to an indicative obstacle density range (in objects per km^2), as shown in Table D.4.

For example:

- Lakes and ponds are considered *Improbable*. UAVs flying over open water face negligible risk of obstacle collision, and the estimated obstacle density is below 5 objects/km^2.
- Rivers, canals, and streams are considered *Remote*. While largely unobstructed, occasional boats or bridges may introduce minor risks (estimated density: 5–20 objects/km^2).
- Cemeteries are considered *Occasional* due to scattered trees, monuments, and fencing (20–50 objects/km^2).
- Bridges are considered *Probable* as they often include structural elements, traffic, and power lines (50–150 objects/km^2).
- Industrial zones are considered *Frequent*, given their dense clusters of fixed infrastructure and dynamic activity (over 150 objects/km^2).

Table D.4: Expert-based likelihood scores for obstacle risk per area type

Score	Interpretation	Applicable area types	Indicative obstacle density (objects/km^2)
1 – Improbable	Open and flat terrain; no vertical structures. Obstacle collisions are highly unlikely.	Lakes and ponds, Water reservoirs	< 5
2 – Remote	Sparse vertical structures; natural elements like trees may occur occasionally.	Tracks and rural access roads, Power lines, Agricultural lands, Meadows and open grass, Rivers, canals and streams, Wetlands, Natura2000 areas	5–20
3 – Occasional	Obstacles are present but scattered. Some collision potential remains.	Motorways and major roads, Wind turbines, Cemeteries	20–50
4 – Probable	Many fixed or moving obstacles; moderate infrastructure and traffic.	Regional roads, Pedestrian and cycling paths, Bridges, Power plants, Communication towers, High infrastructures, Recreational zones, Forests and woodlands, Cultural sites	50–150
5 – Frequent	High density of tall infrastructure and dynamic obstacles; dense urban fabric.	Living and residential streets, Railways, Industrial zones, Commercial zones, Retail zones, Residential areas, Schools and universities, Hospitals, Care homes, Prisons, Religious sites, Parks	> 150

These indicative density thresholds are not based on direct measurement but offer a conceptual framework for further research. Future studies could use drone-based obstacle detection data, satellite imagery, or LIDAR-based scans to refine and validate these likelihood assignments on a quantitative basis.

D.2.2. External risk 2: Interference

Interference includes risks arising from electromagnetic interference (EMI), intentional or unintentional human interference, such as eavesdropping or jamming, and general communications interference [1]. In this study, the scope is narrowed to focus specifically on unintentional signal interference resulting from high densities of overlapping wireless signals. Interference refers to degradation or loss of UAV communication quality due to competition with Wi-Fi networks, cellular signals, or other consumer RF devices.

The likelihood score reflects how saturated an area is in terms of wireless signal emissions, based on land use, infrastructure density, and human activity. These scores are assigned through approximations and spatial reasoning rather than direct RF measurements. To aid future standardization and interpretability, each category is linked to an indicative RF signal density (measured conceptually as concurrent emitting sources per km^2). The proposed standardization across area types is summarized in Table D.5.

- Forests and woodlands are scored as *Improbable*, since these areas lack RF infrastructure and human presence. Indicative density: < 10 emitters/km^2.
- Agricultural lands are considered *Remote*. Although lightly populated, sporadic RF sources may be present (e.g., farm equipment or isolated routers). Density estimate: 10–50 emitters/km^2.
- Pedestrian and cycling paths are scored as *Occasional*, with moderate use and nearby buildings introducing Wi-Fi or mobile interference. Density: 50–200 emitters/km^2.
- Residential areas are scored as *Probable*, due to a high concentration of consumer electronics and wireless networks. Density estimate: 200–1000 emitters/km^2.
- Communication towers and industrial zones are scored as *Frequent*, due to RF saturation from multiple antennas, mobile base stations, and high-volume device deployments. Density: > 1000 emitters/km^2.

Although these emitter densities are indicative and not empirically verified, they provide a structured basis for spatial comparison. Future work could validate these scores through RF mapping campaigns, spectrum analysers on UAV platforms, or integration of crowd-sourced signal coverage datasets.

D.2.3. External risk 3: Communication

Communication risk refers to the likelihood of network congestion, delays, jitter, or unavailability in the communication link between the UAV and external control or coordination systems [1]. This study assesses communication reliability based on the availability and expected performance of mobile network infrastructure (mainly 5G) in different area types in the Netherlands. The underlying assumption is that UAVs operating within U-space require continuous and reliable mobile network connectivity for real-time flight management, navigation updates and coordination.

The likelihood score reflects the expected quality of mobile connectivity per area, expressed as a function of 5G availability, fallback to 4G, and network saturation. Scores are assigned based on approximations about mobile coverage patterns across the Dutch landscape. Each category is linked to an indicative spatial coverage profile, which is defined as the estimated percentage of surface area within the area type where reliable 5G or 4G connectivity is available, as shown in Table D.6. Given the Netherlands' well-developed telecom infrastructure, the occurrence of scores 4 and 5 is considered highly unlikely under normal conditions but included for methodological completeness.

For example:

- Living and residential streets are scored as *Improbable* (1), with near-complete 5G coverage and high network reliability.
- Recreational zones are scored as *Remote* (2); these areas typically have stable 4G and partial 5G coverage, with minor potential signal degradation.
- Agricultural lands are scored as *Occasional* (3), as these rural zones may lack continuous 5G access and rely on less consistent 4G connectivity.

Table D.5: Standardized likelihood scores for Interference risk per area type

Score	Interpretation	Applicable Area Types	RF Emitter Density (emitters/km^2)
1 – Improbable	Remote or natural areas with no RF activity or significant human presence; negligible signal interference.	Tracks and rural access roads, Wind turbines, Agricultural lands, Forests and woodlands, Meadows and open grass, Rivers, canals and streams, Lakes and ponds, Water reservoirs, Wetlands, Religious sites, Cemeteries, Natura2000 areas	< 10
2 – Remote	Areas with low population density or light infrastructure; low likelihood of spectrum congestion or RF overlap.	Motorways and major roads, Pedestrian and cycling paths, Railways, Bridges, Power lines, High infrastructures, Recreational zones, Care homes, Cultural sites, Parks	10–50
3 – Occasional	Moderate infrastructure or crowds; possible interference from Wi-Fi, mobile networks, or RF-based equipment.	Regional roads, Power plants, Prisons	50–200
4 – Probable	High-density public or commercial areas with consistent wireless activity.	Living and residential streets, Commercial zones, Retail zones, Residential areas, Schools and universities, Hospitals	200–1000
5 – Frequent	Areas saturated with RF emissions, such as near mobile base stations, broadcast towers, or dense IoT deployments.	Communication towers, Industrial zones	> 1000

While this study uses expert-based likelihood scores for communication risk, future research could validate and refine these estimates using empirical data. Promising approaches include spatial analysis of mobile signal strength from coverage maps, UAV-based network logging during test flights, or integration of crowd-sourced signal quality datasets (e.g., OpenSignal or CellMapper). These methods would allow area-type-specific calibration of network reliability, transforming the current likelihood estimates into data-driven metrics.

D.2.4. External risk 4: Navigational Environment

Navigational environment risk refers to the potential for UAV navigation errors due to factors such as GPS signal loss or inaccuracy, GPS spoofing, ADS-B errors, attitude estimation faults, or waypoint following mistakes [1]. While this may be due to several factors, this study focuses specifically on urban canyon effects: the degradation of GNSS performance due to dense built environments that obstruct direct satellite lines of sight and introduce multipath reflectivity. This operational focus follows from a key design assumption: UAVs in U-space depend on continuous GNSS positioning for flight coordination and tracking. The risk of navigation degradation is thus defined as a function of structural density and sky cover caused by the built environment.

Instead of modelling technical errors, this study uses approximations to assign likelihood scores based on the expected severity of the urban canyon effect by area type. Scores reflect the degree of structural containment and are linked to the expected building density of each area type, expressed as the estimated number of buildings per square kilometre. The proposed standardization is shown in Table **??**.

Table D.6: Standardized likelihood scores for Communication risk per area type

Score	Interpretation	Applicable Area Types	Indicative mobile coverage
1 – Improbable	Full 5G coverage with high reliability and capacity; communication loss is highly unlikely.	Motorways and major roads, Regional roads, Living and residential streets, Bridges, Communication towers, High infrastructures, Industrial zones, Commercial zones, Retail zones, Residential areas, Wetlands, Schools and universities, Hospitals, Care homes, Prisons, Cultural sites	5G ≥ 95%
2 – Remote	Stable 4G or partial 5G; minor signal degradation may occur, but with low operational impact.	Pedestrian and cycling paths, Railways, Power plants, Recreational zones, Meadows and open grass, Rivers, canals and streams, Lakes and ponds, Religious sites, Cemeteries, Parks	5G 40–95%, 4G fallback
3 – Occasional	Irregular mobile coverage; only stable 4G or weak 5G; potential delays or dropouts.	Tracks and rural access roads, Power lines, Wind turbines, Agricultural lands, Forests and woodlands, Water reservoirs, Natura2000 areas	5G < 40%, 4G partial
4 – Probable	Weak or inconsistent 4G/5G coverage; expected instability and data latency.	–	4G unreliable, 5G largely absent
5 – Frequent	No mobile network coverage or frequent signal loss; UAV communication severely impaired.	–	No 4G or 5G coverage

For example:

- Motorways and major roads are considered *Improbable*. These are fully open linear environments with an unobstructed sky view, resulting in minimal GNSS degradation.
- Pedestrian and cycling paths are considered *Remote*. These routes typically pass through open or semi-open areas with some vegetation or low structures, where GNSS signal loss is rare.
- Regional roads are considered *Occasional*. These areas may have intermittent buildings, trees, or other obstructions that cause temporary interference with GNSS reception.
- Retail zones are considered *Probable*. These semi-enclosed commercial areas often feature built-up environments that restrict sky visibility and may induce multipath effects.
- Industrial zones are considered *Frequent*. These are dense, structurally complex environments with metallic structures, machinery, and narrow corridors. GNSS signal loss and multipath effects are common.

These likelihood scores are based on expert judgement using building density as a proxy for urban GNSS interference. Future research could validate and refine these estimates using spatial data on building footprints and densities derived from OpenStreetMap, 3D BAG datasets, or LiDAR-based urban models. By calculating actual building counts or surface area per square kilometre per area type, more precise and data-driven risk assessments can be established.

Table D.7: Standardized likelihood scores for Navigational Environment risk per area type

Score	Interpretation	Applicable Area Types	Indicative Building Density (buildings/km^2)
1 – Improbable	Fully open areas with no built structures or isolated infrastructure.	Motorways and major roads Railways, Bridges, Power lines, Wind turbines, Agricultural lands, Meadows and open grass, Rivers, canals and streams, Lakes and ponds, Water reservoirs, Wetlands, Natura2000 areas	< 10
2 – Remote	Rural areas with scattered structures or linear infrastructure.	Tracks and rural access roads, Pedestrian and cycling paths, Communication towers, Cemeteries	10–50
3 – Occasional	Areas with moderate and dispersed buildings.	Regional roads, Power plants, High infrastructures, Recreational zones, Forests and woodlands, Care homes, Religious sites, Cultural sites, Parks	50–150
4 – Probable	Built-up areas with regular building rows and commercial structures.	Commercial zones, Retail zones, Schools and universities, Hospitals, Prisons	150–300
5 – Frequent	High-density built environments; enclosed urban spaces.	Living and residential streets, Industrial zones, Residential areas	> 300

D.2.5. External risk 5: Electrical Environment

Electrical environment risk refers to the likelihood that a UAV experiences electromagnetic interference from its surrounding environment. Following the categorization by Allouch et al. [1], this includes two key components:

- High Intensity Radio Transmission Areas (HIRTAs): geographic areas near military or civilian RF emitters (e.g., radar stations, long-range communication towers, public address systems) where powerful electromagnetic fields may interfere with UAV electronics;
- Electrostatic phenomena: naturally occurring or operationally induced buildup of static charge on the surface of the UAV, especially near large conductive structures such as high-voltage cables or industrial equipment

While this study does not measure electromagnetic field strength directly, an approximation is proposed based on the expected number of electromagnetic interference sources by area type. This includes infrastructure such as power lines, substations, RF communication towers, industrial facilities or radar installations. The more such sources are present, the greater the likelihood of sustained electromagnetic exposure. Table D.8 presents the standardized scoring based on this proxy.

For example:

- Forests and woodlands are considered *Improbable*. These areas are free of infrastructure and electronic installations and typically contain less than 2 electromagnetic interference sources per km^2.
- Agricultural lands are considered *Remote*. These areas may contain scattered equipment or isolated buildings of businesses, with an estimated 2-5 potential EM sources per km^2.
- Motorways and major roads are considered *Occasional*. Although mostly open, they are often near infrastructure such as lighting systems, digital signage or communication towers, resulting in about 5-10 EM sources per km^2.

- Railways are considered *Probable*, as electrified train lines, signal systems, and adjacent technical infrastructure can lead to 10–20 electromagnetic sources per km^2.
- Power lines and industrial zones are considered *Frequent*. These areas are directly associated with dense clusters of high-voltage facilities, substations or heavy machinery, with typically more than 20 electromagnetic interference sources per km^2.

Table D.8: Standardized likelihood scores for Electrical Environment risk per area type

Score	Interpretation	Applicable Area Types	Indicative EM Disturbance Sources (per km^2)
1 – Improbable	Natural or remote areas with no known EM infrastructure; only low-power electronics.	Forests and woodlands, Meadows and open grass, Rivers, canals and streams, Lakes and ponds, Wetlands, Cemeteries, Natura2000 areas	< 2
2 – Remote	Rural or recreational areas with occasional technical structures or isolated antennas.	Regional roads, Tracks and rural access roads, Pedestrian and cycling paths, Recreational zones, Agricultural lands, Water reservoirs, Care homes, Religious sites, Cultural sites, Parks	2–5
3 – Occasional	Mixed-use areas with moderate exposure to infrastructure like base stations, lighting systems, or substations.	Motorways and major roads, Living and residential streets, Bridges, Wind turbines, High infrastructures, Commercial zones, Retail zones, Residential areas, Schools and universities, Prisons	5–10
4 – Probable	Infrastructure-heavy zones near known RF emitters, signal systems, or electrical compounds.	Railways, Communication towers, Hospitals	10–20
5 – Frequent	Areas directly adjacent to multiple high-power installations or transmission infrastructure.	Power lines, Power plants, Industrial zones	> 20

Future research could validate electrical environment risk estimates by measuring actual electromagnetic field strength (in V/m) using mobile RF spectrum analyzers or UAV-based logging. Alternatively, publicly available transmitter datasets and electromagnetic zoning maps could be used to quantify the number and type of RF sources per km². Simulation approaches using 3D environmental models and known transmitter power could further support risk estimation in complex environments.

D.3. Risk scores

The risk scores presented in Table D.9 represent the risk scores calculated using the methodology described in Section 4.3.1. The resulting risk matrix provides a hierarchical ranking of area types, ordered from highest to lowest overall risk.

Table D.9: Risk assessment scores for all area types

Area Type	S_f	S_p	S_s	L_1	L_2	L_3	L_4	L_5	**Risk**
Schools and universities	4.0	3.0	4.0	5.0	4.0	1.0	4.0	3.0	0.873
Living and residential streets	3.0	4.0	4.0	5.0	4.0	1.0	5.0	3.0	0.864
Retail zones	4.0	4.0	3.0	5.0	4.0	1.0	4.0	3.0	0.838
Residential areas	2.0	4.0	4.0	5.0	4.0	1.0	5.0	3.0	0.747
Hospitals	2.0	3.0	4.0	5.0	4.0	1.0	4.0	4.0	0.678
Commercial zones	3.0	3.0	3.0	5.0	4.0	1.0	4.0	3.0	0.663
Industrial zones	2.0	4.0	2.0	5.0	5.0	1.0	5.0	5.0	0.632
Parks	4.0	2.0	3.0	5.0	2.0	2.0	3.0	2.0	0.607
Cultural sites	4.0	3.0	3.0	4.0	2.0	1.0	3.0	2.0	0.586
Recreational zones	4.0	2.0	3.0	4.0	2.0	2.0	3.0	2.0	0.574
Railways	4.0	3.0	2.0	5.0	2.0	2.0	1.0	4.0	0.561
Religious sites	2.0	3.0	4.0	5.0	1.0	2.0	3.0	2.0	0.556
Prisons	2.0	3.0	3.0	5.0	3.0	1.0	4.0	3.0	0.529
Cemeteries	3.0	2.0	4.0	3.0	1.0	2.0	2.0	1.0	0.482
Pedestrian and cycling paths	3.0	2.0	3.0	4.0	2.0	2.0	2.0	2.0	0.463
Regional roads	3.0	3.0	2.0	4.0	3.0	1.0	3.0	2.0	0.440
Power plants	1.0	4.0	2.0	4.0	3.0	2.0	3.0	5.0	0.409
Motorways and major roads	4.0	3.0	1.0	3.0	2.0	1.0	1.0	3.0	0.312
High infrastructures	1.0	3.0	2.0	4.0	2.0	1.0	3.0	3.0	0.272
Communication towers	1.0	3.0	1.0	4.0	5.0	1.0	2.0	4.0	0.226
Power lines	1.0	4.0	1.0	2.0	2.0	3.0	1.0	5.0	0.222
Tracks and rural access roads	2.0	2.0	1.0	2.0	1.0	3.0	2.0	2.0	0.144
Agricultural lands	1.0	2.0	1.0	2.0	1.0	3.0	1.0	2.0	0.066
Forests and woodlands	1.0	1.0	1.0	4.0	1.0	3.0	3.0	1.0	0.062
Water reservoirs	1.0	2.0	1.0	1.0	1.0	3.0	1.0	2.0	0.051
Lakes and ponds	2.0	1.0	1.0	1.0	1.0	2.0	1.0	1.0	0.039
Rivers, canals and streams	1.0	1.0	1.0	2.0	1.0	2.0	1.0	1.0	0.010
Meadows and open grass	1.0	1.0	1.0	2.0	1.0	2.0	1.0	1.0	0.010
Wetlands	1.0	1.0	1.0	2.0	1.0	1.0	1.0	1.0	0.000

E

Results

This appendix provides comprehensive experimental results and detailed visualizations that supplement the findings presented in Chapter **??**. While the main chapter focuses on key insights and comparative analyses across the three study areas, this appendix presents the complete dataset including all intermediate α values, detailed infrastructure utilization patterns, and extended metrics.

E.1. Hard Constraints in UAV Routing

As described in Section 6.2, the implementation of regulatory safety requirements creates substantial operational constraints for UAV routing in urban environments. This section visualizes the spatial impact of these hard constraints across the three study areas.

Figures E.1, E.2, and E.3 illustrate the extent of no-fly zones (shown in red) when strict regulatory compliance is enforced. These restricted areas encompass residential zones, retail zones, recreational areas, cultural sites, and religious sites. These are all locations where the presence of uninvolved persons prohibits standard UAV operations. The analysis reveals significant operational limitations across all study regions:

- Breda: 78.46% of the operational area falls within hard constraint zones
- Borsele: 95.83% of the area is classified as restricted
- Alphen aan den Rijn and Waddinxveen: 81.93% of the combined area is subject to no-fly restrictions

These percentages directly translate to delivery accessibility constraints, with only 21.54%, 4.17%, and 18.07% of PostNL pickup points remaining accessible in Breda, Borsele, and Alphen aan den Rijn/Waddinxveen, respectively. Such severe limitations underscore the operational challenge of implementing UAV delivery services under current regulatory frameworks and highlight the necessity of the risk-based corridor design approach developed in this study.

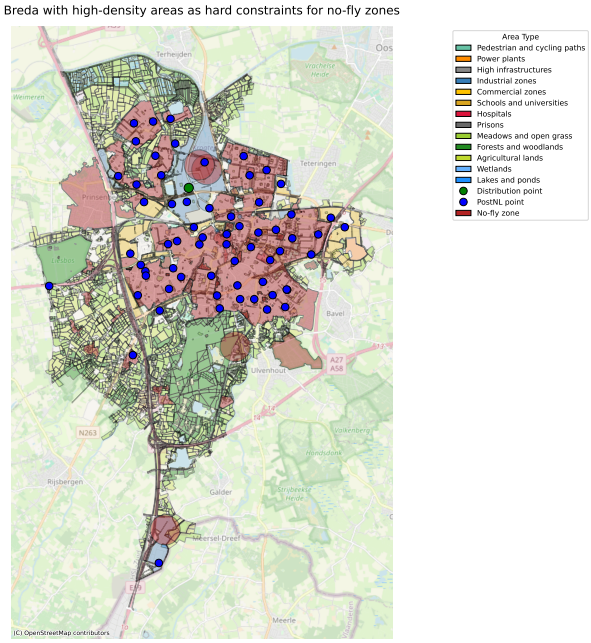

Figure E.1: Hard constraints in Breda municipality.

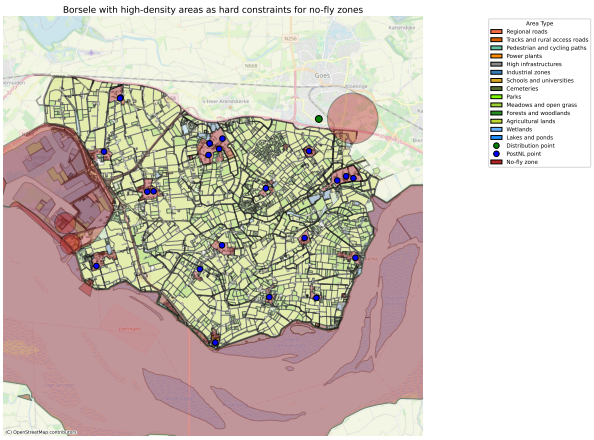

Figure E.2: Hard constraints in Borsele municipality.

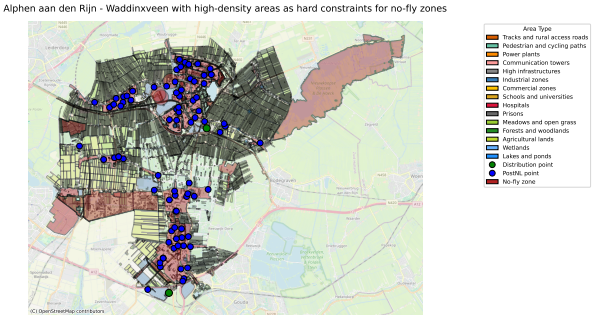

Figure E.3: Hard constraints in Alphen aan den Rijn and Waddinxveen.

E.2. Urban use case: Breda

As described in Section 6.3, the urban context of Breda demonstrates distinctive characteristics in the trade-off between operational efficiency and risk mitigation. This section provides detailed results from the experimental evaluation across all tested α values.

Table E.1 presents comprehensive metrics for each α value, including minimum, maximum, and mean values for all key performance indicators. The data reveals a clear progression from efficiency-oriented routing ($\alpha = 0$) to risk-averse planning ($\alpha = 1$), with notable inflection points occurring at $\alpha = 0.1$ and $\alpha = 0.3$.

Table E.1: Summary statistics for UAV corridor paths in Breda, showing mean, minimum, and maximum values across key metrics as α increases.

α	Risk			Energy (Wh)			Length (m)			Height changes			Turns		
	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
0.00	1372.6	234.5	6045.8	113.6	33.6	367.7	4172.7	862.7	14388.2	5.86	5.00	7.00	123.5	34	225
0.10	393.9	188.5	753.1	130.4	33.8	407.3	4840.1	872.2	15975.4	6.02	5.00	9.00	112.1	31	265
0.20	362.3	182.7	677.8	135.8	33.8	409.4	5042.6	872.2	16058.1	6.21	5.00	9.00	116.8	31	267
0.30	339.2	178.8	671.9	143.6	33.8	409.4	5352.1	872.2	16058.1	6.21	5.00	9.00	123.0	31	267
0.40	333.5	172.1	604.5	146.4	34.8	409.5	5466.4	913.0	16061.6	6.24	5.00	9.00	125.5	31	272
0.50	329.0	172.1	597.5	150.2	34.8	410.1	5609.3	913.0	16085.0	6.27	5.00	11.00	129.1	31	276
0.60	327.1	172.1	597.5	152.2	34.8	410.6	5690.3	913.0	16104.7	6.27	5.00	11.00	133.6	31	278
0.70	324.6	172.0	570.1	157.2	34.8	412.1	5885.3	913.0	16164.9	6.43	5.00	11.00	140.1	31	297
0.80	323.3	172.0	570.1	160.5	34.8	438.7	6016.3	913.0	17228.9	6.43	5.00	11.00	146.8	31	355
0.90	322.1	172.0	555.5	167.9	34.8	441.2	6315.2	913.0	17331.7	6.43	5.00	11.00	157.1	31	367
1.00	322.1	172.0	555.5	169.7	34.8	449.0	6385.2	913.0	17641.1	6.46	5.00	11.00	161.4	31	420

The evolution of infrastructure utilization patterns across different α values is illustrated in Figures E.4, E.5, and E.6. These visualizations reveal several key insights:

- Initial transition ($\alpha = 0.0$ to $\alpha = 0.1$): A dramatic shift occurs from pedestrian/cycling paths and regional roads toward waterway infrastructure and open green spaces. This transition achieves a 71% reduction in average risk while increasing energy consumption by only 14%, as shown in Table E.1.
- Stabilization phase ($\alpha = 0.1$ to $\alpha = 0.3$): The infrastructure mix remains relatively stable, with waterways maintaining dominance.
- Refinement phase ($\alpha = 0.3$ to $\alpha = 1.0$): Step changes further reduce the use of high-risk corridors, but the return in risk reduction clearly decreases (by merely 5.0%). The final configurations show an almost complete elimination of pedestrian paths, regional roads and the use of residential streets.

Notably, certain infrastructure types demonstrate consistent utilization across all α values. Tracks and rural roads maintain steady usage levels, suggesting their balanced risk-efficiency profiles make them suitable for both operational paradigms. This consistency indicates that some corridor types naturally align with both safety and efficiency objectives in urban environments.

The turn frequency metrics reveal an interesting pattern: initial risk considerations ($\alpha = 0$ to $\alpha = 0.1$) actually reduce average turn frequency from 123.5 to 112.1, suggesting that preliminary risk avoidance may eliminate marginally efficient but structurally complex paths. Only at higher α values does turn frequency increase, reaching 161.4 at $\alpha = 1.0$ as routes become increasingly circuitous to avoid residual high-risk areas.

E.3. Rural Use Case: Borsele

As described in Section 6.3, the rural context of Borsele presents unique challenges for UAV corridor design due to its sparse infrastructure network and limited routing alternatives. This section provides detailed experimental results across all tested α values, revealing how rural environments necessitate more severe trade-offs between efficiency and risk mitigation.

Table E.2 presents comprehensive metrics revealing how rural environments impose unique operational challenges for risk-averse routing. Borsele's efficiency penalties are substantial; path length increase

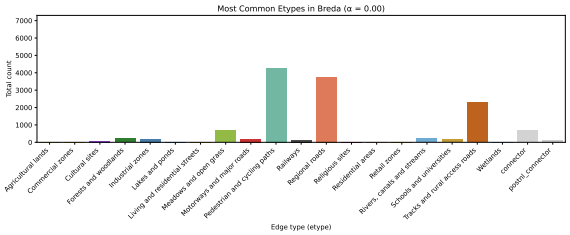

(a) $\alpha = 0.00$

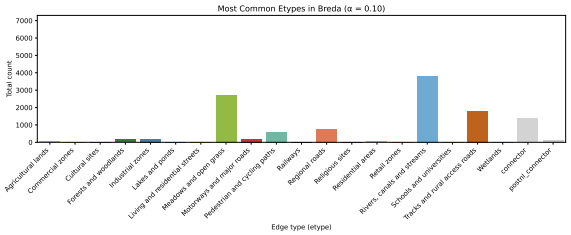

(b) $\alpha = 0.10$

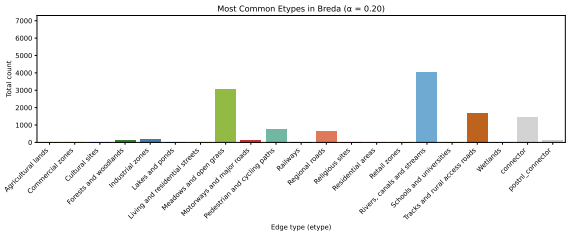

(c) $\alpha = 0.20$

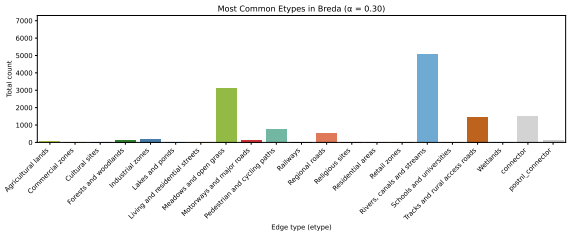

(d) $\alpha = 0.30$

Figure E.4: Corridor segment usage by edge type in Breda for $\alpha = 0.00 - 0.30$

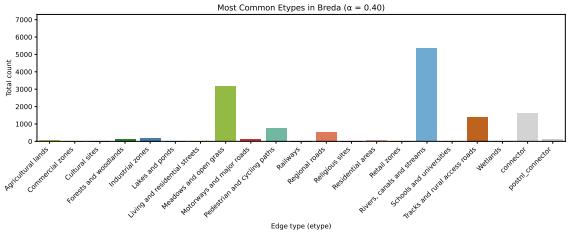

(a) $\alpha = 0.40$

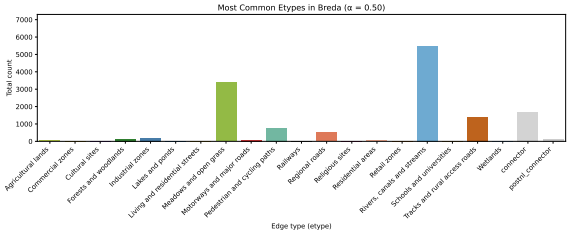

(b) $\alpha = 0.50$

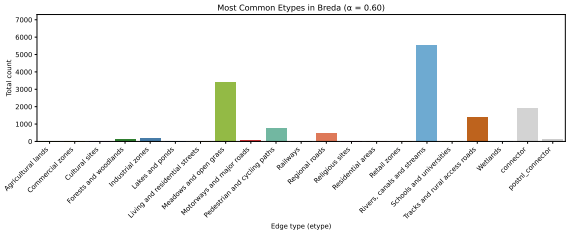

(c) $\alpha = 0.60$

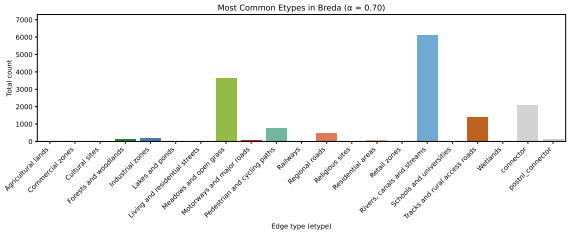

(d) $\alpha = 0.70$

Figure E.5: Corridor segment usage by edge type in Breda for $\alpha = 0.40 - 0.70$

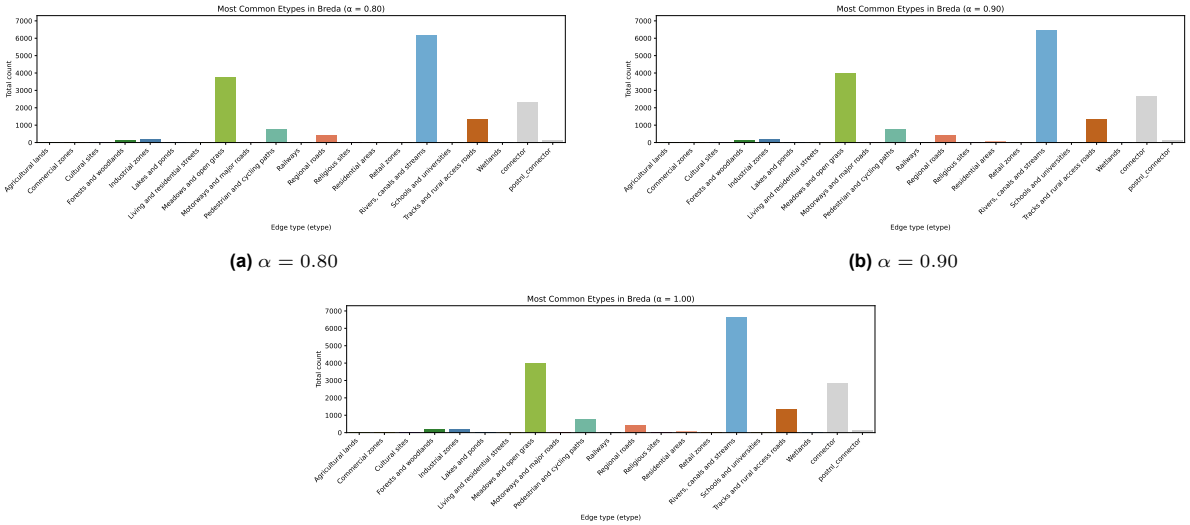

(a) $\alpha = 0.80$

(b) $\alpha = 0.90$

(c) $\alpha = 1.00$

Figure E.6: Corridor segment usage by edge type in Breda for $\alpha = 0.80 - 1.00$

by 61% and energy consumption by 60% when transitioning from efficiency-focused ($\alpha = 0.0$) to risk-averse planning ($\alpha = 1$). The most striking characteristic is the dramatic increase in routing complexity, with turn frequency nearly doubling (92% increase) from 158.1 to 304.1 turns.

Table E.2: Summary statistics for UAV corridor paths in Borsele, showing mean, minimum, and maximum values across key metrics as α increases.

α	Risk			Energy (Wh)			Length (m)			Height changes			Turns		
	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
0.00	1872.7	483.9	3702.2	222.7	57.1	381.4	8748.8	2210.5	15099.4	4.90	3.00	7.00	158.1	43	270
0.10	564.3	205.6	892.2	243.7	59.8	421.4	9559.1	2317.8	16373.6	5.38	3.00	11.00	209.1	43	345
0.20	430.7	188.3	708.0	266.3	62.7	469.1	10477.8	2434.4	18610.0	5.00	3.00	9.00	228.1	45	377
0.30	385.1	171.9	626.0	280.5	69.4	472.3	11053.4	2701.7	18735.3	4.62	3.00	7.00	242.4	57	377
0.40	377.5	164.6	622.4	284.9	73.9	478.6	11231.0	2879.8	18989.9	4.62	3.00	7.00	248.1	58	384
0.50	358.3	149.4	593.3	301.1	88.6	508.3	11851.5	2983.5	20176.6	5.00	3.00	9.00	260.1	71	401
0.60	341.6	147.3	589.0	320.9	90.8	532.7	12645.0	3069.6	21151.0	5.10	3.00	9.00	271.5	73	409
0.70	340.8	147.3	585.7	322.3	90.8	533.1	12701.1	3069.6	21166.7	5.10	3.00	9.00	273.0	73	411
0.80	337.5	147.3	582.9	332.7	90.8	535.5	13109.5	3069.6	21263.8	5.29	3.00	9.00	280.8	73	419
0.90	332.4	146.8	547.6	354.2	94.1	570.5	13971.9	3203.4	22662.3	5.29	3.00	9.00	299.8	75	490
1.00	332.2	146.8	547.4	357.5	94.1	572.5	14103.3	3203.4	22745.8	5.29	3.00	9.00	304.1	75	491

The evolution of infrastructure utilization patterns illustrated in Figures E.7, E.8, and E.9 reveals several distinctive characteristics of rural UAV routing:

- Dominant infrastructure shift ($\alpha = 0.0$ to $\alpha = 0.1$): The most striking transformation occurs in the abandonment of linear infrastructure corridors. Efficiency-based routing heavily utilizes tracks and rural roads, pedestrian/cycling paths, and regional roads. Risk-aware planning dramatically shifts toward agricultural lands and open grasslands, which become the predominant corridor types for all $\alpha > 0$.
- Turn complexity explosion: Unlike the urban context, Borsele shows a nearly 92% increase in turn frequency between $\alpha = 0$ (158.1 turns) and $\alpha = 1$ (304.1 turns). This dramatic increase reflects the sparse rural network's inability to provide direct low-risk alternatives, forcing UAVs to navigate circuitous paths through agricultural areas.

A critical observation from the detailed metrics is the initial risk reduction efficiency. Between $\alpha = 0$ and $\alpha = 0.1$, average risk plummets from 1872.7 to 564.3, a 70% reduction. Energy consumption increases by only 9.4%. However, further risk reductions come at increasingly steep costs. The transition from $\alpha = 0.1$ to $\alpha = 1.0$ achieves only an additional 41% risk reduction while imposing a 47% energy penalty.

Height change metrics remain relatively stable across all α values (4.90 to 5.29 average changes), in-

dicating that vertical complexity in rural environments is primarily determined by terrain characteristics rather than risk considerations. This stability contrasts with the dramatic increase in horizontal complexity. This suggests that rural risk mitigation works mainly in the horizontal plane through path extension rather than height adjustments.

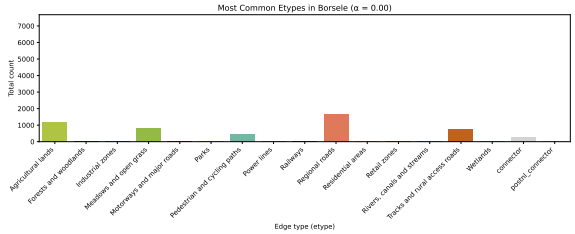

(a) $\alpha = 0.00$

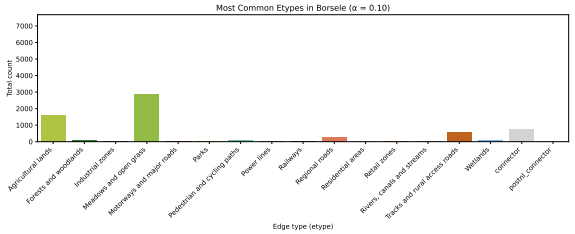

(b) $\alpha = 0.10$

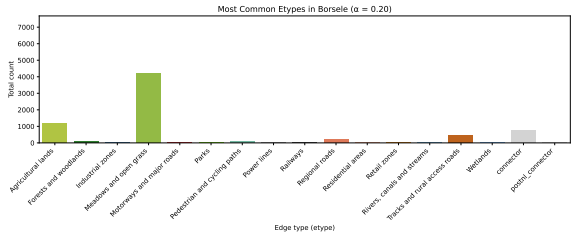

(c) $\alpha = 0.20$

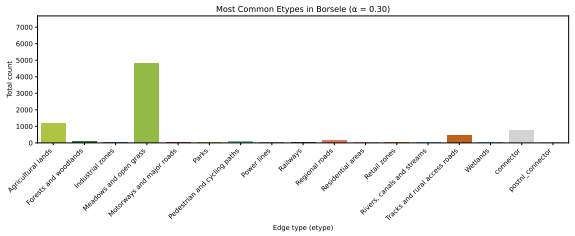

(d) $\alpha = 0.30$

Figure E.7: Corridor segment usage by edge type in Borsele for $\alpha = 0.00 - 0.30$

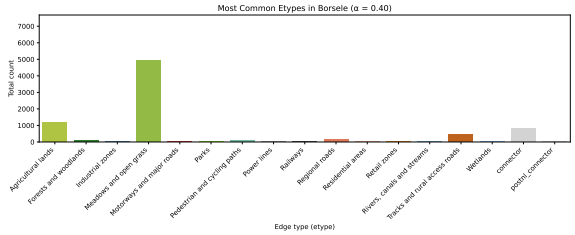

(a) $\alpha = 0.40$

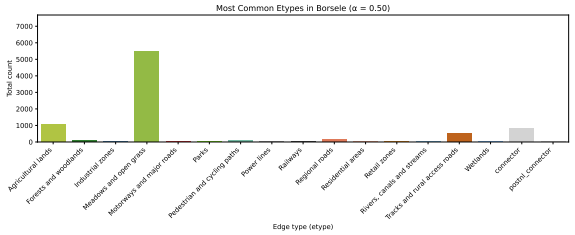

(b) $\alpha = 0.50$

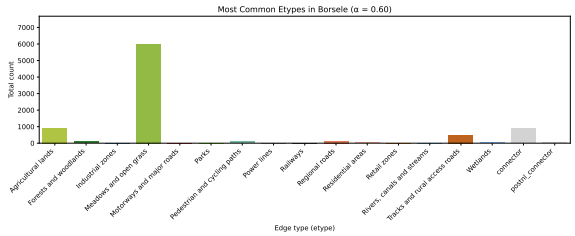

(c) $\alpha = 0.60$

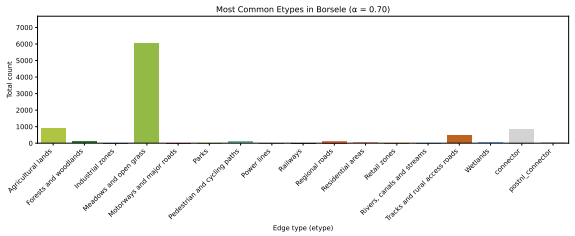

(d) $\alpha = 0.70$

Figure E.8: Corridor segment usage by edge type in Borsele for $\alpha = 0.40 - 0.70$

E.4. Suburban Use Case: Alphen aan den Rijn and Waddinxveen

As described in Section 6.3, the suburban context of Alphen aan den Rijn and Waddinxveen presents a unique hybrid environment that combines urban areas with natural infrastructure's in between. This section provides detailed experimental results demonstrating how the dual distribution center configuration and suburban area influence routing decisions across all tested α values.

Table E.3 reveals the most severe efficiency penalties among all study areas, with energy consumption increasing by 86% and path length by 88% when transitioning from efficiency-focused to risk-averse planning. Despite these substantial increases, the topological complexity grows by only 22%, suggesting that the suburban environment offers spatially simple risk mitigation pathways.

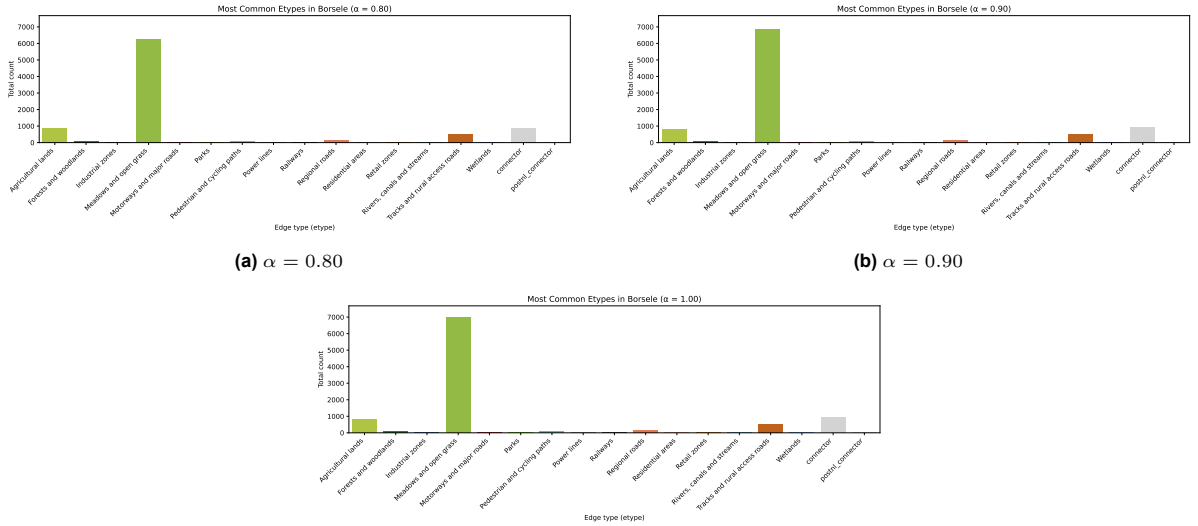

(a) $\alpha = 0.80$ **(b)** $\alpha = 0.90$

(c) $\alpha = 1.00$

Figure E.9: Corridor segment usage by edge type in Borsele for $\alpha = 0.80 - 1.00$

Table E.3: Summary statistics for UAV corridor paths in Alphen–Waddinxveen, showing mean, minimum, and maximum values across key metrics as α increases.

α	Risk			Energy (Wh)			Length (m)			Height changes			Turns		
	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
0.00	1507.5	129.9	3092.5	117.0	17.2	271.4	4524.2	615.0	10780.2	4.48	3.00	9.00	102.2	17	196
0.10	352.4	41.3	931.7	129.3	19.1	287.0	5004.5	688.5	11406.6	4.43	3.00	9.00	109.3	22	209
0.20	290.4	41.3	708.9	140.2	19.1	295.6	5434.7	688.5	11587.6	4.51	3.00	9.00	102.0	22	220
0.30	262.7	38.3	608.3	149.6	20.1	361.0	5799.4	729.3	14365.8	4.48	3.00	9.00	100.6	25	251
0.40	247.2	38.3	512.7	157.1	20.1	379.9	6088.9	729.3	15083.7	4.51	3.00	9.00	103.3	25	274
0.50	240.5	38.3	425.5	162.6	20.1	379.5	6302.4	729.3	15106.6	4.53	3.00	9.00	102.7	25	275
0.60	239.8	38.3	425.4	163.5	20.1	379.7	6335.6	729.3	15112.0	4.58	3.00	9.00	103.4	25	276
0.70	233.9	32.3	425.4	174.1	29.7	379.7	6765.4	1112.6	15112.0	4.48	3.00	9.00	106.0	33	276
0.80	231.5	32.3	425.4	181.6	29.7	381.6	7062.2	1112.6	15119.4	4.48	3.00	9.00	108.9	33	279
0.90	228.5	32.3	409.4	200.3	29.7	456.6	7807.4	1112.6	18189.9	4.61	3.00	9.00	114.5	33	259
1.00	227.3	32.3	408.3	218.2	29.7	467.0	8524.6	1112.6	18604.3	4.56	3.00	9.00	124.6	33	269

The infrastructure utilization patterns illustrated in Figures E.10, E.11, and E.12 demonstrate a distinctive substitution pattern unique to suburban environments:

- Waterway dominance ($\alpha = 0.0$ to $\alpha = 0.1$): The most dramatic transformation occurs in the wholesale adoption of waterway infrastructure. Rivers, canals, and streams, which play a minimal role in efficiency-based routing, become the dominant corridor type for all $\alpha \geq 0.1$. This shift achieves a remarkable 77% risk reduction (from 1507.5 to 352.4) while increasing energy consumption by only 10.5%.
- Infrastructure abandonment pattern: Unlike the gradual elimination seen in other contexts, Alphen-Waddinxveen exhibits rapid and near-complete abandonment of high-risk corridors. Railways, motorways, major roads, and pedestrian/cycling paths, which collectively dominate at $\alpha = 0$, virtually disappear by $\alpha = 0.1$, but are not completely absent at higher values.
- Natural corridor emergence ($\alpha = 0.1$ to $\alpha = 1.0$): As risk aversion increases, the algorithm increasingly leverages meadows, open grasslands, and forested areas to complement the waterway network. This diversification of natural infrastructure could enable continued risk reduction without the dramatic path complexity increases observed in rural environments.

Despite the 88% increase in path length, average turn frequency increases by only 22% (102.2 to 124.6). This remarkable efficiency stems from the linear nature of waterway corridors, which provide continuous low-risk paths without requiring frequent directional changes. The stability of height change metrics (4.48 to 4.56 average changes) across all α values confirms that the suburban topology, dominated by waterways at or near ground level, inherently minimizes vertical complexity. This characteristic makes

waterway-based routing particularly suitable for energy-efficient risk mitigation, as UAVs can maintain consistent altitude while following these natural corridors.

A critical observation unique to this case study is the synergy between geographic features and risk mitigation strategies. The extensive waterway network—comprising the Oude Rijn, Gouwe, and numerous connecting canals—provides a pre-existing low-risk corridor system that requires minimal adaptation for UAV operations. This natural infrastructure advantage explains why Alphen-Waddinxveen achieves the highest risk reduction (85%) among all study areas while maintaining operational feasibility.

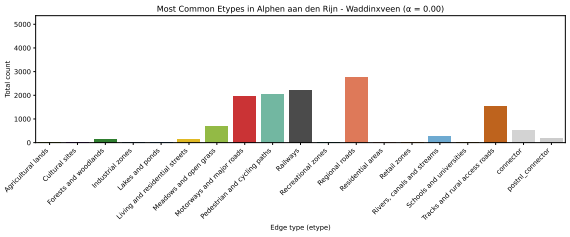

(a) $\alpha = 0.00$

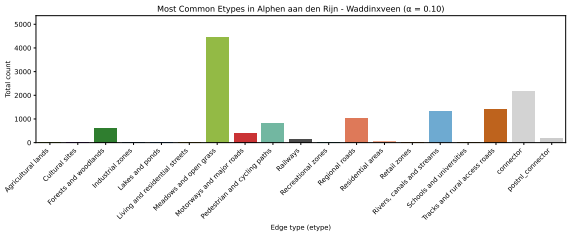

(b) $\alpha = 0.10$

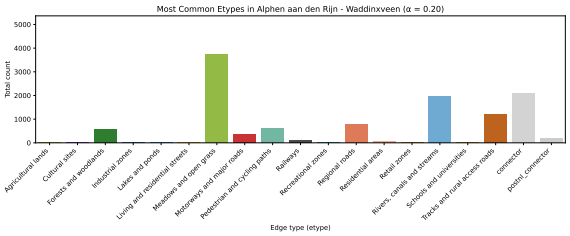

(c) $\alpha = 0.20$

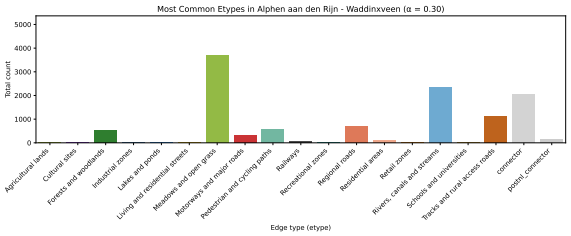

(d) $\alpha = 0.30$

Figure E.10: Corridor segment usage by edge type in Alphen aan den Rijn and Waddinxveen for $\alpha = 0.00 - 0.30$

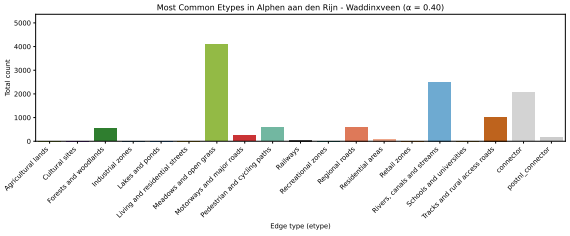

(a) $\alpha = 0.40$

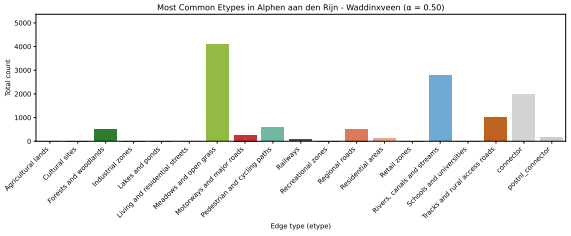

(b) $\alpha = 0.50$

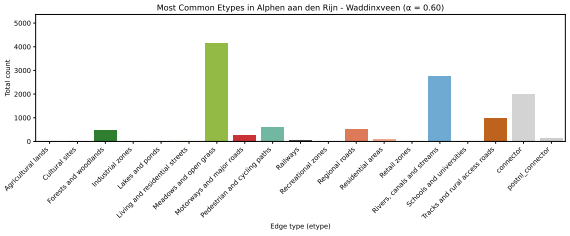

(c) $\alpha = 0.60$

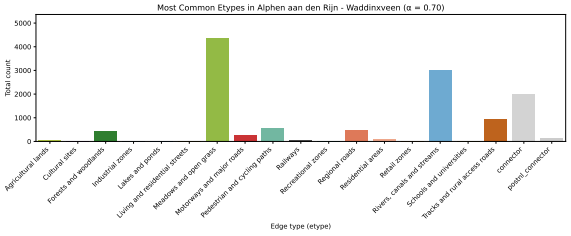

(d) $\alpha = 0.70$

Figure E.11: Corridor segment usage by edge type in Alphen aan den Rijn and Waddinxveen for $\alpha = 0.40 - 0.70$

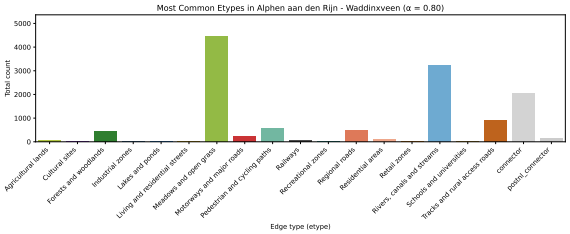

(a) $\alpha = 0.80$

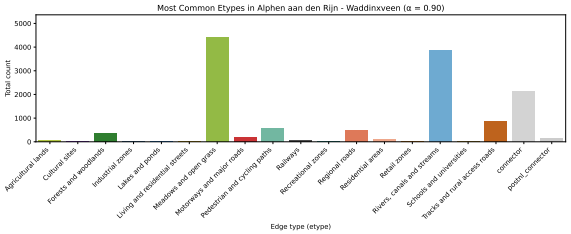

(b) $\alpha = 0.90$

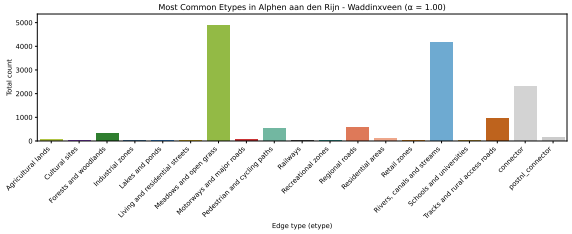

(c) $\alpha = 1.00$

Figure E.12: Corridor segment usage by edge type in Alphen aan den Rijn and Waddinxveen for $\alpha = 0.80 - 1.00$

F

Sensitivity Analysis

The proposed framework incorporates several assumptions, particularly in the semi-quantitative risk assessment parameters (severity and likelihood scores, described in Chapter 4) and energy consumption parameters (described in Chapter 5). This section presents detailed sensitivity analyses to evaluate the model's sensitivity to these assumptions.

F.1. Sensitivity analysis on severity

As discussed in Section 4.3.1, severity scores were assigned to different area types across three consequence domains: fatality risk, property damage risk, and societal impact. These severity levels, defined using ordinal categories ranging from Negligible to Catastrophic, were mapped to numerical values (1-4) to enable quantitative risk aggregation. Since these severity assumptions fundamentally determine which areas are classified as high-risk versus low-risk, and thereby directly influence corridor selection and the resulting safety-efficiency trade-offs, it is crucial to assess the sensitivity of the model to these input parameters.

To evaluate the robustness of the corridor optimization to severity score uncertainty, a series of perturbation experiments was conducted. The analysis focused on the urban case (Breda) and targeted the area types most frequently traversed in the computed UAV corridors: meadows and open grasslands, pedestrian and cycling paths, regional roads, rivers/streams/canals, and tracks and rural access roads.

For each area type, the severity scores across all three consequence domains were systematically varied by ± 1 ordinal level within the defined scale. For example, an initial severity classification of Critical (score = 3) was adjusted to both Catastrophic (+1, score = 4) and Marginal (−1, score = 2). Boundary conditions were enforced to prevent scores from exceeding the minimum (Negligible, score = 1) or maximum (Catastrophic, score = 4) levels.

To isolate the effects of individual consequence domains, each sensitivity run altered only one severity domain at a time per area type. This approach avoids confounding effects and enables evaluation of whether the model is more sensitive to changes in fatality risk assumptions versus property damage or societal impact assumptions. For each configuration, corridor routes were recomputed across α values of 0.0, 0.5, and 1.0.

The sensitivity analysis reports signed averages rather than absolute values of the percentage changes resulting from +1 and -1 perturbations. This approach provides insight into model behaviour beyond simple magnitude: when positive and negative perturbations produce symmetric responses, the average approaches zero, confirming linear model behaviour in that parameter range.

F.1.1. Results severity score sensitivity

The reported sensitivity metrics represent the average absolute change across both positive (+1) and negative (-1) severity perturbations for individual area types, showing their impact on network-wide

performance metrics.

First, the model's internal consistency is validated in Table F.1 through the $\alpha = 0$ results, where all routing metrics show exactly 0.0% change across all severity perturbations. This confirms that efficiency-based optimization is independent of risk parameters. The small risk variations (0-3.13%) observed reflect post-hoc risk calculations on unchanged routes.

At $\alpha = 0.5$ in Table F.2, adjusting severity scores for specific area types produces significant network-wide effects. Waterway corridors (rivers, canals, and streams) show the strongest influence: changing their severity score causes average risk across all routes to shift by up to 22.13%, with turn frequency increasing 4.8-6.3%. Meadows and open grasslands produce the second-largest network impact. This outsized influence reflects these corridor types' dominance in risk-aware networks from $\alpha = 0.1$ onwards (Figure E.4). The greater increase in turns compared to path length indicates that routes become more circuitous while still attempting to minimize distance penalties, which is logical for $\alpha = 0.5$.

At $\alpha = 1.0$ in Table F.3, the network's sensitivity to waterway severity adjustments manifests differently. A single-level change in waterway severity causes average height changes across all routes to vary by 26.29%, while turn frequency shows less dramatic changes. This shift from horizontal (turns) to vertical (height) sensitivity reflects a fundamental difference in optimization constraints. Height changes incur energy penalties (2.03 Wh for ascent, 0.93 Wh for descent) but do not affect risk scores, as areas maintain their risk profile regardless of flight altitude. At $\alpha = 0.5$, this energy penalty still influences routing decisions, leading to preference including more turns as dicussed above. However, at $\alpha = 1.0$, energy costs become irrelevant in the optimization. This removes the penalty for height changes, allowing routes to freely switch between corridor types at different altitudes when risk profiles change. The 26.29% height variation thus reflects newfound vertical flexibility under pure risk minimization.

The progressive increase in sensitivity from $\alpha = 0$ to 0.5 to 1.0 validates the weighted objective function. Across both $\alpha = 0.5$ and $\alpha = 1.0$, adjusting societal impact parameters produces the largest network-wide changes, followed by fatality and property damage adjustments. This hierarchy aligns with the risk calculation framework where societal impact applies directly while fatality and property scores are moderated by likelihood multiplication.

Table F.1: Severity score sensitivity analysis for α = 0.0 (Energy-only optimization)

Area Type	Metric	Fatality	Property	Societal
Meadows and open grass	Risk Change (%)	0.49	0.29	1.08
	Length Change (%)	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00
Pedestrian and cycling paths	Risk Change (%)	0.00	0.00	-0.00
	Length Change (%)	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00
Regional roads	Risk Change (%)	0.00	0.00	0.00
	Length Change (%)	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00
Rivers, canals and streams	Risk Change (%)	0.22	0.13	0.48
	Length Change (%)	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00
Tracks and rural access roads	Risk Change (%)	0.00	0.00	3.13
	Length Change (%)	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00

Table F.2: Severity score sensitivity analysis for α = 0.5 (Balanced optimization)

Area Type	Metric	Fatality	Property	Societal
Meadows and open grass	Risk Change (%)	6.37	4.39	10.70
	Length Change (%)	0.81	0.28	1.17
	Energy Change (%)	0.80	0.31	1.12
	Turn Change (%)	1.85	1.25	2.70
	Height Change (%)	0.77	0.77	0.26
Pedestrian and cycling paths	Risk Change (%)	-0.16	-0.09	-0.34
	Length Change (%)	-0.37	-0.03	-0.42
	Energy Change (%)	-0.31	-0.01	-0.36
	Turn Change (%)	-0.58	-0.37	-0.62
	Height Change (%)	0.77	0.51	0.77
Regional roads	Risk Change (%)	-0.63	-0.42	-0.79
	Length Change (%)	0.44	0.59	0.47
	Energy Change (%)	0.46	0.59	0.48
	Turn Change (%)	1.01	0.82	0.89
	Height Change (%)	0.77	0.77	0.77
Rivers, canals and streams	Risk Change (%)	13.21	8.36	22.13
	Length Change (%)	-1.59	-0.26	-2.78
	Energy Change (%)	-1.48	-0.24	-2.56
	Turn Change (%)	4.83	5.47	6.32
	Height Change (%)	0.51	0.51	1.53
Tracks and rural access roads	Risk Change (%)	-1.24	-0.65	5.91
	Length Change (%)	-0.90	-0.35	-0.41
	Energy Change (%)	0.56	1.04	1.07
	Turn Change (%)	0.72	0.74	0.19
	Height Change (%)	1.02	0.51	0.77

Table F.3: Severity score sensitivity analysis for α = 1.0 (Risk-averse optimization)

Area Type	Metric	Fatality	Property	Societal
Meadows and open grass	Risk Change (%)	6.97	4.76	11.45
	Length Change (%)	-2.93	-2.45	-2.66
	Energy Change (%)	-2.77	-2.33	-2.49
	Turn Change (%)	-4.08	-3.73	-3.03
	Height Change (%)	-0.74	-0.74	-0.25
Pedestrian and cycling paths	Risk Change (%)	-0.30	-0.08	-0.56
	Length Change (%)	-0.26	-0.33	-0.43
	Energy Change (%)	-0.25	-0.30	-0.39
	Turn Change (%)	0.69	-0.21	0.49
	Height Change (%)	-0.25	0.25	0.00
Regional roads	Risk Change (%)	-0.45	-0.14	-0.66
	Length Change (%)	-1.07	-0.15	-1.15
	Energy Change (%)	-0.85	-0.01	-0.90
	Turn Change (%)	-0.52	0.13	-0.57
	Height Change (%)	0.74	0.74	0.74
Rivers, canals and streams	Risk Change (%)	13.75	9.03	22.80
	Length Change (%)	-4.67	-4.26	-6.15
	Energy Change (%)	-4.12	-3.74	-4.39
	Turn Change (%)	-2.55	-3.33	-1.17
	Height Change (%)	6.39	6.39	26.29
Tracks and rural access roads	Risk Change (%)	-1.45	-0.65	5.99
	Length Change (%)	-0.46	0.00	-1.43
	Energy Change (%)	0.90	1.22	0.01
	Turn Change (%)	1.15	1.11	-0.76
	Height Change (%)	2.70	0.00	2.95

F.2. Sensitivity analysis on likelihood

Likelihood scores capture the probability of UAV-related incidents occurring due to five external risk factors: obstacles (L1), interference (L2), communication (L3), navigational environment (L4), and electrical environment (L5). These assessments, ranging from *improbable* (1) to *frequent* (5), were assigned to each area type based on environmental characteristics (Appendix D). Since likelihood scores directly multiply with severity scores in the risk calculation framework, uncertainties in these probability assessments can significantly impact corridor selection.

The analysis maintained focus on the urban case (Breda) and targeted the same five frequently traversed area types as in F.1. For each area type, the likelihood scores for all five external risk factors were systematically perturbed by ± 1 ordinal level within the five-point scale. For instance, a baseline likelihood assessment of *probable* (score = 4) was adjusted to both *frequent* (+1, score = 5) and *occasional* (−1, score = 3). Boundary constraints were enforced to maintain the ordinal structure, preventing scores from exceeding the scale limits of *improbable* (1) to *frequent* (5).

To isolate the influence of individual risk factors, each sensitivity experiment modified only one likelihood factor at a time per area type. This approach enables identification of which external risk factors most strongly influence corridor selection. Similar to the severity analysis, corridors were recomputed for each perturbed configuration across α values of 0.0, 0.5, and 1.0.

F.2.1. Results likelihood score sensitivity

Again, the reported sensitivity metrics represent the average absolute change across both positive (+1) and negative (-1) likelihood perturbations.

Similar to the severity analysis, efficiency-based routing ($\alpha = 0$) demonstrates complete independence from likelihood parameters, with exactly 0% change in all routing metrics across all external factors and area types, validating the model's internal consistency.

Tables F.5 and F.6 reveal a strong clustering pattern in the sensitivity results. In most cases, Obstacles

(L1) and Communication (L3) produce identical results, while Interference (L2), Navigation (L4), and Electrical (L5) form a second group with matching impacts. However, this clustering is not perfect: notable exceptions occur for regional roads, where Communication shows distinct behavior (e.g., +0.38% risk change versus -0.36% for other factors at $\alpha = 0.5$). This predominant but imperfect clustering suggests that the likelihood scoring methodology assigns similar scores to factors within each group across most area types, likely reflecting underlying correlations in how these risks manifest spatially.

The analysis confirms that adjusting likelihood scores for heavily-used corridor types produces the largest network-wide impacts. Rivers, canals, and streams show the highest sensitivity, particularly to the Interference/Navigation/Electrical group, with risk changes up to 3.93% and turn frequency variations reaching 10.96%. However, these impacts remain substantially smaller than those from severity adjustments. Most metrics change by less than 5% even under balanced optimization. This reduced sensitivity aligns with the risk calculation framework, where likelihood scores only moderate fatality and property components, unlike severity scores which directly influence all three risk domains.

Notably, no single likelihood factor consistently produces the largest network impact. The Obstacles/Communication group drives large turn frequency changes when meadows and waterways are adjusted, while the Interference/Navigation/Electrical group shows stronger effects on risk scores for other areas. This context-dependent pattern validates the multi-factor approach, though the observed clustering suggests the five theoretical dimensions may effectively represent only two underlying risk mechanisms.

Table F.4: Likelihood factor sensitivity analysis for α = 0.0 (Energy-based optimization)

Area Type	Metric	Obst.	Inter.	Comm.	Nav.	Elec.
Meadows and open grass	Risk Change (%)	0.00	0.11	0.00	0.11	0.11
	Length Change (%)	0.00	0.00	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00	0.00	0.00
Pedestrian/cycling paths	Risk Change (%)	0.00	0.00	0.00	0.00	0.00
	Length Change (%)	0.00	0.00	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00	0.00	0.00
Regional roads	Risk Change (%)	0.00	0.00	1.26	0.00	0.00
	Length Change (%)	0.00	0.00	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00	0.00	0.00
Rivers, canals, streams	Risk Change (%)	0.00	0.05	0.00	0.05	0.05
	Length Change (%)	0.00	0.00	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00	0.00	0.00
Tracks and rural roads	Risk Change (%)	0.00	0.64	0.00	0.00	0.00
	Length Change (%)	0.00	0.00	0.00	0.00	0.00
	Energy Change (%)	0.00	0.00	0.00	0.00	0.00
	Turn Change (%)	0.00	0.00	0.00	0.00	0.00
	Height Change (%)	0.00	0.00	0.00	0.00	0.00

Table F.5: Likelihood factor sensitivity analysis for α = 0.5 (Balanced optimization)

Area Type	Metric	Obst.	Inter.	Comm.	Nav.	Elec.
Meadows and open grass	Risk Change (%)	-2.03	1.78	-2.03	1.78	1.78
	Length Change (%)	1.38	0.32	1.38	0.32	0.32
	Energy Change (%)	1.33	0.30	1.33	0.30	0.30
	Turn Change (%)	6.75	0.00	6.75	0.00	0.00
	Height Change (%)	1.53	0.00	1.53	0.00	0.00
Pedestrian/cycling paths	Risk Change (%)	-0.04	-0.04	-0.04	-0.04	-0.04
	Length Change (%)	0.08	0.08	0.08	0.08	0.08
	Energy Change (%)	0.04	0.04	0.04	0.04	0.04
	Turn Change (%)	-0.53	0.09	-0.53	0.09	0.09
	Height Change (%)	-1.02	0.00	-1.02	0.00	0.00
Regional roads	Risk Change (%)	-0.36	-0.36	0.38	-0.36	-0.36
	Length Change (%)	0.69	0.69	0.57	0.69	0.69
	Energy Change (%)	0.68	0.68	0.57	0.68	0.68
	Turn Change (%)	0.00	1.75	0.00	1.75	1.75
	Height Change (%)	0.00	1.02	0.00	1.02	1.02
Rivers, canals, streams	Risk Change (%)	-2.26	3.48	-2.26	3.48	3.48
	Length Change (%)	1.46	-0.05	1.46	-0.05	-0.05
	Energy Change (%)	1.42	-0.03	1.42	-0.03	-0.03
	Turn Change (%)	0.00	5.62	0.00	5.62	5.62
	Height Change (%)	0.00	1.53	0.00	1.53	1.53
Tracks and rural roads	Risk Change (%)	-0.17	1.48	-0.17	-0.17	-0.17
	Length Change (%)	0.31	0.06	0.31	0.31	0.31
	Energy Change (%)	0.29	0.09	0.29	0.29	0.29
	Turn Change (%)	-0.12	0.68	-0.12	0.68	0.68
	Height Change (%)	-0.51	0.00	-0.51	0.00	0.00

Table F.6: Likelihood factor sensitivity analysis for α = 1.0 (Risk-based optimization)

Area Type	Metric	Obst.	Inter.	Comm.	Nav.	Elec.
Meadows and open grass	Risk Change (%)	-2.25	2.00	-2.25	2.00	2.00
	Length Change (%)	2.36	-1.34	2.36	-1.34	-1.34
	Energy Change (%)	2.25	-1.29	2.25	-1.29	-1.29
	Turn Change (%)	-5.25	2.55	-5.25	2.55	2.55
	Height Change (%)	-1.47	1.47	-1.47	1.47	1.47
Pedestrian/cycling paths	Risk Change (%)	-0.03	-0.03	-0.03	-0.03	-0.03
	Length Change (%)	-0.28	-0.28	-0.28	-0.28	-0.28
	Energy Change (%)	-0.26	-0.26	-0.26	-0.26	-0.26
	Turn Change (%)	-0.44	0.09	-0.44	0.09	0.09
	Height Change (%)	0.00	0.49	0.00	0.49	0.49
Regional roads	Risk Change (%)	-0.05	-0.05	0.54	-0.05	-0.05
	Length Change (%)	-0.18	-0.18	0.00	-0.18	-0.18
	Energy Change (%)	-0.14	-0.14	0.03	-0.14	-0.14
	Turn Change (%)	-0.11	0.25	-0.11	0.25	0.25
	Height Change (%)	0.00	0.98	0.00	0.98	0.98
Rivers, canals, streams	Risk Change (%)	-2.06	3.93	-2.06	3.93	3.93
	Length Change (%)	-1.24	-1.83	-1.24	-1.83	-1.83
	Energy Change (%)	-1.18	-1.72	-1.18	-1.72	-1.72
	Turn Change (%)	-10.96	0.00	-10.96	0.00	0.00
	Height Change (%)	-0.49	0.00	-0.49	0.00	0.00
Tracks and rural roads	Risk Change (%)	-0.26	1.28	-0.26	-0.26	-0.26
	Length Change (%)	-0.48	-0.50	-0.48	-0.48	-0.48
	Energy Change (%)	0.74	0.73	0.74	0.74	0.74
	Turn Change (%)	-0.27	0.04	-0.27	0.04	0.04
	Height Change (%)	-0.49	0.00	-0.49	0.00	0.00

F.3. Sensitivity analysis on energy parameters

The proposed framework jointly optimizes for safety and efficiency, trading off third party ground risk against estimated energy consumption. While the previous section focused on the model's sensitivity

to uncertainty in risk parameters, it is equally important to assess the robustness of the corridor design with respect to assumptions in the energy consumption model.

The energy estimates used in this study were derived from a combination of literature values and simplified physical assumptions, and as such, are subject to uncertainty. To evaluate how these uncertainties might affect the corridor outcomes, a sensitivity analysis was performed by systematically perturbing the three core parameters of the energy model: horizontal energy consumption per meter, vertical energy consumption for ascent, and vertical energy consumption for descent.

The following parameter scenarios were evaluated:

- Base: $(0.025, 2.03, 0.93)$
- Higher horizontal (+20%): $(0.030, 2.03, 0.93)$
- Lower horizontal ($-$20%): $(0.020, 2.03, 0.93)$
- Higher vertical (+20%): $(0.025, 2.436, 1.116)$
- Lower vertical ($-$20%): $(0.025, 1.624, 0.744)$
- All higher (+20%): $(0.030, 2.436, 1.116)$
- All lower ($-$20%): $(0.020, 1.624, 0.744)$

For each scenario, the full corridor optimization was recomputed for the urban test case (Breda) under $\alpha = 0.0$, $\alpha = 0.5$, and $\alpha = 1.0$.

F.3.1. Results of energy parameter sensitivity

Tables F.7, F.8, and F.9 demonstrate that energy parameter variations have minimal impact on route selection, in stark contrast to risk parameter sensitivity.

At $\alpha = 1.0$, routing metrics, except for the energy consumption, remain completely unchanged across all energy perturbations, validating that pure risk optimization is independent of energy parameters. Energy consumption responds proportionally to parameter changes: horizontal energy adjustments of $\pm$20% yield approximately $\pm$18.8% change in total consumption, with the remaining difference accounted for by the vertical component. When all parameters change by $\pm$20%, total consumption changes by exactly $\pm$20%, confirming the model's internal consistency.

Additionally at $\alpha = 0$ and $\alpha = 0.5$ energy consumption responds proportionally to parameter changes. However, other notable findings emerge.

First, vertical energy perturbations at $\alpha = 0$ and $\alpha = 0.5$ produce almost no change in route planning. Maximum changes are minimal: length increases by up to 0.1%, risk by up to 0.3%, turn count by up to 0.5 turns, and altitude changes by only 0.06, all effectively negligible. Horizontal energy changes also result in very limited planning adaptations, with the largest change in length being +1.2%, risk decreasing by 0.5%, turns increasing by 3, and altitude changes again negligible.

At both $\alpha = 1$ and $\alpha = 0$, simultaneous $\pm 20\%$ changes in all parameters yield exactly $\pm 20\%$ changes in total energy consumption. For $\alpha = 0.5$, the effect is slightly dampened due to the mixed objective function. This outcome illustrates that horizontal energy dominates total consumption, while a 20% change in vertical energy costs has minimal impact.

Despite efficiency-based optimization at $\alpha = 0$, route changes remain negligible. Vertical energy perturbations produce virtually no impact: path length varies by at most 0.1%, risk by 0.3%, and turns by 0.5. Horizontal energy changes, though more influential, still yield minimal routing adjustments with maximum length changes of 1.2% and risk changes of 0.5%.

The limited influence of energy parameters on routing decisions can be attributed to een combinatie van deze drie dingen:

- Restricted route options: If only two or three feasible routes exist from A to B, and all use similar corridor types, a 20% increase in horizontal energy makes all options approximately 20% more expensive, offering little incentive to switch routes.

- Dominance of horizontal consumption: For a typical 10km route with 2 altitude changes, horizontal energy (250 Wh) vastly exceeds vertical energy (4 Wh, <2% of total), making vertical cost variations negligible.
- Hoogte verschillen zijn te weinig variaties in:

The limited influence of energy parameters on routing decisions can be attributed to a combination of three factors:

- **Restricted route options**: If only two or three feasible routes exist from A to B, and all use similar corridor types, a 20% increase in horizontal energy makes all options approximately 20% more expensive, offering little incentive to switch routes.
- **Dominance of horizontal consumption**: For a typical 10km route with 2 altitude changes, horizontal energy (250 Wh) vastly exceeds vertical energy (4 Wh, <2% of total), making vertical cost variations negligible.
- **Limited altitude transition variability**: If most alternative routes have similar numbers of altitude transitions, these are dictated by the infrastructure types that must be traversed. Unlike continuous optimization where altitude could be gradually adjusted, the binary choice between 30m and 60m means routes cannot fine-tune vertical profiles for energy savings.

Table F.7: Energy parameter sensitivity analysis for α = 0.0 (Energy-based optimization)

Scenario	Length (m)	Risk	Energy (Wh)	Turns	Alt. Changes
Base	4,173	1,373	113.6	123.5	5.86
Horizontal +20%	4,172 (-0.02%)	1,372 (-0.04%)	134.5 (+18.4%)	123.4	5.86
Horizontal -20%	4,174 (+0.03%)	1,377 (+0.3%)	92.8 (-18.3%)	124.0	5.83
Vertical +20%	4,174 (+0.03%)	1,377 (+0.3%)	115.5 (+1.7%)	124.0	5.83
Vertical -20%	4,172 (-0.02%)	1,372 (-0.04%)	111.8 (-1.6%)	123.4	5.86
All +20%	4,173 (0.0%)	1,373 (0.0%)	136.3 (+20.0%)	123.5	5.86
All -20%	4,173 (0.0%)	1,373 (0.0%)	90.9 (-20.0%)	123.5	5.86

Table F.8: Energy parameter sensitivity analysis for α = 0.5 (Balanced optimization)

Scenario	Length (m)	Risk	Energy (Wh)	Turns	Alt. Changes
Base	5,609	329.0	150.1	129.1	6.27
Horizontal +20%	5,561 (-0.9%)	330.4 (+0.4%)	176.7 (+17.7%)	127.1	6.21
Horizontal -20%	5,677 (+1.2%)	327.4 (-0.5%)	123.5 (-17.7%)	132.2	6.27
Vertical +20%	5,603 (-0.1%)	329.3 (+0.1%)	151.9 (+1.2%)	129.3	6.21
Vertical -20%	5,609 (0.0%)	329.0 (0.0%)	148.2 (-1.3%)	129.1	6.27
All +20%	5,558 (-0.9%)	330.6 (+0.5%)	178.4 (+18.9%)	126.9	6.17
All -20%	5,677 (+1.2%)	327.4 (-0.5%)	121.5 (-19.0%)	132.2	6.27

Table F.9: Energy parameter sensitivity analysis for α = 1.0 (Risk-averse optimization)

Scenario	Length (m)	Risk	Energy (Wh)	Turns	Alt. Changes
Base	6,385	322.0	169.7	161.4	6.46
Horizontal +20%	6,385 (0.0%)	322.0 (0.0%)	201.6 (+18.8%)	161.4	6.46
Horizontal -20%	6,385 (0.0%)	322.0 (0.0%)	137.8 (-18.8%)	161.4	6.46
Vertical +20%	6,385 (0.0%)	322.0 (0.0%)	171.7 (+1.2%)	161.4	6.46
Vertical -20%	6,385 (0.0%)	322.0 (0.0%)	167.7 (-1.2%)	161.4	6.46
All +20%	6,385 (0.0%)	322.0 (0.0%)	203.6 (+20.0%)	161.4	6.46
All -20%	6,385 (0.0%)	322.0 (0.0%)	135.8 (-20.0%)	161.4	6.46

F.4. Conclusions from sensitivity analysis

The sensitivity analyses reveal fundamental differences in how the UAV corridor network responds to parameter uncertainty. Risk parameters (severity and likelihood) can trigger complete route reorganization by changing corridor attractiveness, while energy parameters only adjust the cost of existing routes without changing their relative preference.

Key findings include:

Differential parameter sensitivity: Severity parameters show the highest sensitivity (up to 26.29% for height changes), particularly for heavily-used corridor types in risk-aware planning. Likelihood parameters demonstrate moderate sensitivity (<5%) with unexpected clustering of five theoretical factors into two effective groups. Energy parameters produce minimal routing changes (<2%), as the discrete network structure and dominance of horizontal flight limit optimization flexibility.

Model validation: The analyses confirm proper model behaviour through parameter independence at extreme α values (risk parameters at α=0, energy parameters at α=1) and progressive sensitivity increases as objectives blend. The signed average approach reveals that most parameters exhibit symmetric responses, validating linear model behaviour within tested ranges.

Critical parameters for implementation: The high sensitivity to severity scores for dominant corridor types (waterways, meadows) emphasizes the importance of accurate risk assessment for all areas. Misspecification could lead to either overuse of incorrectly classified low-risk areas or underutilization of safe corridors incorrectly deemed high-risk.

Model limitations identified: The analysis reveals that turning manoeuvrers, while increasing path complexity and presumably energy consumption, are not penalized in the current energy model. This explains the why the preference for height changes over additional turns is only at $\alpha = 1.0$ and not at $\alpha = 0.5$ and suggests that true efficiency penalties may be underestimated for complex routes.

These findings suggest that future implementations should prioritize validation of severity scores, as path planning is highly dependent on these parameters. Likelihood scores are somewhat less critical but appear to have unintentionally clustered into groups. This clustering may result from risk factors being too similar across area types, indicating a need for more precise empirical measurement. Appendix D already outlines possible methodologies for such empirical validation.

To improve the energy model, which currently has limited influence on routing decisions, incorporating turn penalties would provide more realistic efficiency assessments, particularly important for balanced optimization scenarios where both risk and energy matter. Additionally, implementing a more complex altitude model is crucial. The current binary choice between flight levels means altitude decisions rarely affect route selection, effectively reducing the optimization to a comparison between path length and risk without meaningful energy trade-offs. A continuous or multi-level altitude framework would enable more nuanced energy optimization and better reflect real-world UAV operations.