An Agent-Based Safety Analysis of the Third Party Risk of UAS operations in the urban environment

Master of Science Thesis

Bastiaan Zwanenburg

TU Delft

An Agent-Based Safety Analysis of the Third Party Risk of UAS operations in the urban environment

Master of Science Thesis

by

Bastiaan Zwanenburg

To obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday February 18th, 2021.

Student number:	4393678	
Project duration:	May 2020 – February 2021	
Thesis committee:	Prof. dr. ir. J.M. Hoekstra	TU Delft, Chairman
	Dr. O. A. Sharpanskykh	TU Delft, Supervisor
	Dr. B. F. Santos	TU Delft, Examiner

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Preface

Dear reader,

This thesis reports on nine months of research carried out for obtaining a Master of Science degree in Aerospace Engineering at Delft University of Technology. This work focuses on understanding the safety of drone operations in cities. Despite all the hype concerning drones flying through cities, my work often felt as pioneering, as little is yet understood about the safety of these operations. It was a remarkable experience to dedicate myself fully to this research, for it has been a unique experience to spend nine months on a project like this.

I have learned a lot from the past months. In this educational process, it was my supervisor, Dr. Alexei Sharpans'kykh, who guided the way. Our meetings never failed to provide me with fundamental insights and spark my enthusiasm for this research in general. This is especially remarkable as Alexei achieved this by only asking a few well-directed questions. Having had the opportunity to learn about the agent-based modelling paradigm is something I will carry for a lifetime. This tool shaped my way of thinking on a wide range of problems, whether it is a COVID-pandemic, teamwork, or the safety of drones. Furthermore, I am grateful for having had the opportunity to serve as a Teaching Assistant in the Agent-Based Modeling course. It was very rewarding to pass on the lessons that I learned, to a new generation of students. Also, while Alexei is supervising more students than two hands can count, his availability and knowledge of the ins- and outs- of my project, always gave me the feeling like I was the only student under his supervision.

While this thesis reports on nine months of research, it also concludes an educational journey which has been 20 years in the making. This journey is defined not merely by the main road which had to be followed, but even more so by a forest of side-roads that distracted me from this main road. Even though it has been a bumpy ride, I am grateful for every step of the way. I owe much debt to my family for getting at where I am right now. It is my late grandfather who infected me with his enthusiasm about Delft's university, and its broader student community. However, most important were my parents, who taught me to navigate between the side roads and the main roads, as without this navigation it is conceivable that I would have gotten lost in the forest called "high school", let alone finish university. Also, they didn't miss a single beat of any of the 24 years of my life. Concluding this thesis, I am grateful for what I have achieved, I am even more grateful for who I have become, and I am most grateful for everyone who helped me getting here.

Bastiaan Zwanenburg
Delft, December 2020

Contents

List of Figures

List of Tables

List of Abbreviations

ABMS	Agent-Based Modeling and Simulation
AGL	Above Ground Level
AIS	Abbreviated Injury Scale
ANN	Average Nearest Neighbor Index
ANSP	Air Navigation Service Provider
ARC	Air Risk Class
ATC	Air Traffic Control
ATM	Air Traffic Management
BBN	Bayesian Belief Network
BC	Blunt Criterion
BVLOS	Beyond Visual Line of Sight
c&dphase	climb&descent-phase
CAA	UK Civil Aviation Authority
CBS	Conflict Based Search
CD & R	Conflict-detection and resolution
CNS	Communication, Navigation and Surveillance
ConOps	Concept of Operations
CT	Constraint Tree
DAA	Detect and Avoid
DP	Delivery Point
EASA	European Aviation Safety Agency
FAA	Federal Aviation Authority
FIMS	Flight Information Management System
FTA	Fault Tree Analysis
GRC	Ground Risk Class
HFRM	High Fidelity Risk Model
ID	Independence Detection
JARUS	Joint Authorities for Rulemaking on Unmanned Systems
KNMI	Dutch Meteorological Institute
MAPF	Multi-Agent Path Finding
MTBF	Mean Time Between Failure
OSM	Open Street Map
OSO	Operational Safety Objectives
PAV	Passenger Aerial Vehicle
PDF	Probability Density Function
POF	Probability of Fatality
RBPF	Risk Based Path Finding
SAIL	Specific Assurance and Integrity Level
SORA	Specific Operations Risk Assessment
TCL	Technical Capability Levels
TMPR	Tactical Mitigation Performance Requirement
TPR	Third Party Risk
UA	Unmanned Aircraft
UAS	Unmanned Aircraft System
USS	UAS Service Supplier
UTM	Unmanned Aircraft Systems Traffic Management
VLOS	Visual Line of Sight

Introduction

Delivering packages using drones in urban areas used to be a futuristic concept. However, several companies have conducted flight-tests that indicate that this is technologically feasible [51, 97, 132]. Furthermore, economic studies show that it is an attractive proposition with significant demand [27, 85]. One hurdle remains, which is to gain approval of regulators. These regulators indicate that they demand a better understanding of the risk associated with these operations, before they would lend their approval [32, 110].

In the past years, several researchers modeled operations of unmanned aircraft systems (UAS) in urban areas accurately. However, these models where mainly focused on analyzing airspace structuring and traffic management. Also, much research was done on analyzing the risk associated with crashes of drones. The objective of this research is to blend these fields, by creating an accurate model of urban UAS operations, and combining this with state-of-the-art methods for computing the risk associated with these operations.

This model is created by means of an Agent-Based Safety Analysis. In this model, the cities of Delft, New York and Paris are modeled by combining different data sources. The concept of operations regards the delivery of small packages using drones. Ultimately, this model is used to draw conclusions on two things. Firstly, the results in different environments are compared to draw conclusions on what causes the model to behave differently in different environments. Secondly, the methodology behind the risk computation is scrutinized.

This thesis report is organized as follows: In Part I, the scientific paper is presented. Part II contains the relevant Literature Study that supports the research. Finally, Part III presents supplementary work. This supplementary work includes an elaboration on the model description in chapter 1. Secondly, chapter 2 presents results of additional experiments, as well as several sensitivity studies. In chapter 3, the statistical methods used to obtain our results, are presented. Lastly, chapter 4 proposes fourteen recommendations for future research.

I

Scientific Paper

An Agent-Based Safety Analysis of the Third Party Risk of UAS operations in the urban environment

Bastiaan J.V. Zwanenburg,*

Delft University of Technology, Delft, The Netherlands

Abstract

There are many diverse concepts for autonomous drone operations. One of these concepts is the delivery of packages in urban areas. Transportation companies are already testing vehicles capable of performing such operations, and many studies indicate that this concept is both technically feasible and financially attractive. However, how much risk do civilians living in cities (also called third-party-risk, or TPR) face following these drone operations? Moreover, what measures can be employed to mitigate this risk? Regulators demand answers to these questions before allowing autonomous Unmanned Aircraft Systems (UAS) operations to take off. Recent research has focused on developing methods of calculating the TPR, or on creating models that accurately resemble drone operations in urban cities. The model proposed in this work bridges the gap between these categories using an agent-based safety risk analysis. In this analysis, we study the TPR of UAS package delivery operations in Delft, New York and Paris. This model leads to two main contributions. The first considers observations regarding the influence of the environment on the TPR. In particular, our model suggests that if the low-risk areas are clustered in a city, this leads to higher TPR. The second contribution is derived from the interaction of the risk computation with our model of the environment. Global- and local sensitivity analyses led to a few interesting observations. For example, our model suggests that it is more important to understand the vehicle's failure rate in the cruise-phase, than in the takeoff- and landing-phase. It is also suggested that it is more important to understand how buildings protect people from impact than to understand the effects of an impact directly on a human. Another finding is that modelling the impact speed with the terminal speed, as is common in literature, leads to a TPR that is 15% - 26% higher than when the impact speed is modelled based on the drone's dynamics.

1 Introduction

Vehicles flying through the world's largest cities used to be a concept belonging to science-fiction movies but is quickly becoming a reality. Technology enabling these operations is developing quickly, with companies like Amazon, Google, and Uber having conducted flight-tests with vehicles capable of transporting packages or even passengers through cities [1, 2, 3]. Recently, Wing *(a subsidiary of Alphabet, Google's parent company)* and Amazon have received FAA approval for commercial package deliveries with crewless vehicles, as long as there is still a human pilot in control [4, 5, 6, 7].

It is not without reason that so many organizations show interest in this field. Locascio et al. estimated that delivering packages using UASs could reduce financial costs by a factor of 7 compared to trucks [8]. Similarly, Doole et al. estimated a factor 5 cost reduction compared to e-bikes [9]. Doole et al. also forecasted that a city of Paris would have about 440 million package deliveries by drones per year in 2035 [9]. However, these estimations hinge on vehicles being allowed to fly autonomously over urban areas, as urban areas are where the bulk of the packages are to be delivered. Such operations are not allowed yet, and regulators in the EU and US alike demand a better understanding of the risk associated with these operations before they would grant their approval [10, 11].

What precisely is meant by "risk"? Kaplan et al. define risk as the combined answer to three questions [12]: (i) What can go wrong? (ii) How likely is it? And (iii) what are the consequences? In the case of urban UAS operations, we are especially interested in the consequences to **third parties**, which are not directly related to the operations at all. Third Party Risk (TPR) is so important because these people are most at risk, and because regulators have indicated that they are especially interested in third-party risk [10, 11].

An understanding of how the environment influences the risks can help further research into mitigating these risks. However, this understanding can only be gathered from a risk model integrated with several realistic environments. Despite significant research activity, no such research was found. Instead, the research found

*Msc Student, Air Transport and Operations, Faculty of Aerospace Engineering, Delft University of Technology

falls in one of three categories. The first category is research focused on a specific aspect of the risk computation (such as the failure rate), the second category focuses only on a small environment (i.e. a 100x100 meter patch of a city, or a single road), and the third category of research focuses on complex urban UAS operations. However, research in this third category focuses not on safety, but on efficiency or UAS Traffic Management (UTM) considerations.

This research sets out to fill this gap by developing a detailed risk model of drone package delivery operations in several urban metropoles. We do so by combining state-of-the-art methods for calculating the Third-Party Risk, with a realistic concept of operations of small packages deliveries using UASs. We have modelled the cities of Paris, New York, and Delft to evaluate these operations. By developing this model, we aim to contribute to the scientific field in two ways. Firstly, our model makes it possible to **identify the most essential elements of this risk computation**. The second contribution is through the comparison of different environments. Using three cities, it becomes possible to evaluate **the impact of individual environmental parameters on safety, and the interaction between these parameters**. Furthermore, these three cities are of different sizes, continents, layouts, and were founded millennia apart. Therefore, it is feasible that if we find results that generalize across these three cities, that they will generalize to many more cities. We have released this model as an open-source tool, allowing everyone to repeat and build upon our research easily.

This research aimed to compare risk values for different model variants and scenarios. Because real-world data is missing for several input parameters, such as the failure rate of drones, the absolute values of the risk may not be accurate. While this poses no limitation on the possibility to compare different model variants and parameters, it is impossible to evaluate whether or not the level of risk is acceptable.

We followed the steps of an Agent-Based Safety Risk Analysis [13], which is a process that assesses the safety of a proposed concept of operations. This is a suitable approach because it analyzes the system risks that emerge from drones' local behaviour in the environment. One of the steps is to model the agent-based safety risk model, which models the risk following from the concept of operations. This model is developed using Agent-Based Modeling and Simulation (ABMS). This modelling paradigm is appropriate for this research, as ABMS makes it possible to capture global phenomena (system-wide risk) following from local autonomous behaviours (movements of individual UASs) [14]. Also, ABMS is flexible and modular, which is especially advantageous in a developing field like UAS safety operations, because it allows for the addition of new components in the future [15]. Furthermore, ABMS has previously been used to study urban UAS operations [16], and it has been used as a means of safety-analysis [17]. Therefore, both by virtue of its characteristics and its successful application in similar research, it is chosen as the modelling approach for this research.

We present a model of UAS package deliveries in Delft, New York and Paris. We modelled these cities by combining data on the population density and buildings from different sources, and we modelled the TPR of these operations using the most advanced methods found in literature.

This paper is structured as follows. Firstly, section 2 elaborates on related work. An overview of the methodology is presented in section 3. Section 4 discusses the case that is studied. The model is outlined in section 5. After that, section 6 elaborates on the approach for verification and validation. Furthermore, section 7 describes the experimental approach, and section 8 considers both the setup and results of the experiments. In section 9, the model's limitations and the results are discussed. Lastly, the conclusions and recommendations for future work are presented in section 10.

2 Related Work

It is not only companies and regulators that have worked intensively on urban UAS operations, but academic research on this topic has also gained momentum. This section presents an overview of the most relevant academic research. Section 2.1 presents research on models of UAS operations in urban areas, and section 2.2 discusses work on the safety of such operations. Finally, section 2.3 synthesises the related work.

2.1 Modelling UAS Operations in Urban Areas

Researchers have extensively discussed envisioned operations of UASs in urban areas over the past years. For example, NASA has proposed four concepts of operations for future UAS operations in conjunction with the FAA. For each of the concepts, NASA specifies associated Technical Capability Levels (TCL). The fourth, and most advanced, of these concepts, would allow autonomous beyond-visual-line-of-sight (BVLOS) operations in high traffic density situations, flying over areas with high population densities [18, 19]. The feasibility of this concept of operations was proven with flight-tests in 2019 [20]. As of 2020, the FAA is working on integrating this Concept of Operations into existing regulations [7, 21]. Furthermore, Rios et al. argue that it is feasible that drones will be allowed to freely roam urban skies, as long as they comply with constraints set out by

Unmanned Aircraft Systems Traffic Management (UTM) [11]. These constraints could include no-fly zones, or zones that the drone must not leave. This condition's relevance is that drones can freely start missions without having to wait for UTM approval.

Also, much work has been done on modelling Urban UAS operations. In 2015, Hoekstra et al. conducted a research project, Metropolis, in which they modelled high-density traffic in the urban airspace. This research was pioneering in how it modelled the urban airspace, and it presents important findings on how structures such as layers and tubes can be employed to structure flights through the urban airspace. However, Third-Party-Risk (TPR) was not the focus of this research and therefore modelled rather simplistically. For example, paths were not optimized for risk, and the effect of sheltering by buildings, trees, or vehicles was largely neglected [22]. Salleh et al. built upon this research, comparing the efficiency of several methods of structuring the urban airspace, using a high-fidelity model of Singapore [23, 24]. Pongsakornsathien et al. developed an advanced representation of the city of Melbourne to study several CNS-related scenarios (Communication, Navigation and Surveillance) [25].

We thus conclude that despite research focusing on the development of accurate representations of Urban UAS operations, the goals of these research project are mostly related to UTM (Urban Traffic Management) related topics, such as airspace structuring and CNS.

2.2 UAS Safety Analysis

The field of UAS Safety has seen a rapid increase in research activity over the past years. Searching Scopus gives 2,406 papers published between 1990 and 2020. Of these, 50% have been published in 2018 or later. This section summarizes the state of the art of this research field.

In literature, several classifications of risk exist. Firstly, one can distinguish between collective risk (CR), individual risk (IR) and societal risk (SR). IR is the risk to one exposed person, CR is obtained by multiplying IR by the people exposed to an event, and SR is the CR summed over all possible scenarios [26]. Secondly, risk can also be categorized based on whether people belong to first-, second- or third-parties. First party-risk encompasses people directly involved in the operations, second party-risk concerns people who are not involved in the operations but do derive benefit from it, and third-party-risk (TPR) describes people who have nothing to do with the operations at all [27]. Regulators in both the US and the EU have indicated that for acceptance of UAS delivery operations, TPR is the most critical parameter [11, 10]. One can also categorize risk by different failure modes, resulting in different descent profiles [28]. Concerning the latter, Clothier et al. developed a Barrier-Bowtie-Model to understand the different types of failure modes, and their associated probabilities, better [29, 27]. Furthermore, Primatesta et al. have researched the effects different failure modes would have on the descent trajectory [30].

Another high-fidelity ground risk analysis is done by Bertrand et al., who narrowed the focus on their environment such that they only consider a crash on a single stretch of road or railway [31, 32]. Using a 100x100 meter area, compromising the business district of a metropole, Kim et al. researched the effect of crashes due to two drones colliding mid-air [33]. Rattanagraikanakorn et al. also developed a method for assessing the consequences of UAS collisions [34].

JARUS, a consortium of EU countries, developed SORA, which is a qualitative method for assessing the ground risk following from UAS operations [35]. SORA was compared to a high-fidelity ground risk method by LaCour-Harbo, concluding that the risk is of the same order of magnitude [36]. It should be noted that this high-fidelity analysis mostly considered rural areas.

Considering the field of path-finding, Rudnick-Cohen et al. worked on planning takeoff trajectories to reduce the risk and only considered the small area around the takeoff location in their work [37]. Primatesta et al. extended the A* algorithm to RiskA*, minimizing risk, rather than path-length [38].

2.3 Summary of related work

Summarizing, we identify three research gaps. Firstly, many works on UAS Safety focus on only one or a few components of risk, but not on the integrated risk calculation. Secondly, most safety analyses consider small environments, being only a small neighbourhood, a road, or a very approximate city. We found little work aimed at using high-fidelity environments to understand UAS Safety. Thirdly, we found no research that studies the interaction between a high-fidelity environment and an integrated risk model.

3 Overview of Methodology

Our modeling approach follows the procedure of an agent-based safety risk analysis, which is a process that analyses the risk of a proposed concept of operations [13]. A part of this procedure is to develop an agent-based safety risk model that quantifies the risks arising from the operations. Section 3.1 explains the process of an agent-based safety risk analysis. Subsequently, section 3.2 outlines the agent-based modeling paradigm used for the safety risk model.

3.1 Agent-Based Safety Risk Analysis

In an agent-based safety risk analysis, agent-based techniques are used to study risks on a global level, emerging from the local behaviour of individual agents [13]. Blom et al. describe this analysis as an iterative process of eight steps, as presented in Figure 1. These eight steps can be categorized in two phases. The first phase is qualitative and establishes the context, including the operations and the associated hazards. The second phase is quantitative, and describes how the operations are modelled and evaluated.

Within the first phase, step one identifies the objective, including the scope, safety criteria and goals. Step two describes the operation, which encompasses all procedures and relevant operational context. Thirdly, all hazards are identified. One method for determining hazards is Functional Hazard Identification, which means that all possible failure conditions and their effects are listed for each function. Another method is Pure Hazard Brainstorming, in which experts brainstorm on possible hazards in an open, unrestricted setting [13, 39]. In the fourth step, scenarios are constructed based on the hazards. In section 4, we elaborate on these four qualitative steps as used in our research.

Within the second phase, step five considers developing the model. Different models can be created, such as models based on statistical data, or agent-based models. In step five, this model is evaluated and the risks are calculated, for example, by using a Monte Carlo simulation. Subsequently, step six evaluates whether or not the risk is acceptable, and step seven scrutinizes this value to find specific bottlenecks contributing to the risk. However, as mentioned in the introduction, the absence of accurate data for several parameters made it impossible to derive accurate risk values. Therefore, no conclusions were drawn regarding acceptability. However, we still scrutinize the (relative) values in step eight to understand how different elements contribute to these risks, and to compare the risk emerging from different model variants and scenarios. In the remainder of this section, we explain how the safety risk model was developed.

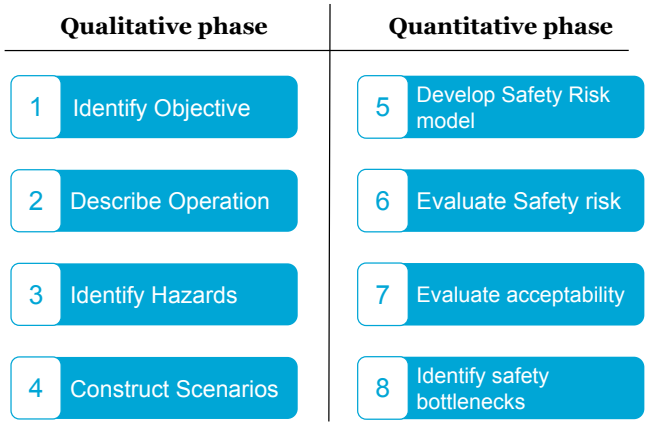

Figure 1: Safety Risk Assessment Steps (adapted from [39]).

3.2 Agent-Based Modeling and Simulation

To develop the safety risk model, we use ABMS as modeling paradigm. An agent-based model consists of three elements [40]. The first element is a specification of the agents and their local properties. Agents can have behavioral properties (stimulus-response) or cognitive properties (based on an internal state, e.g. goal-oriented behavior). Secondly, the environment is specified, which encompasses all non-agent objects. The environment can be static or dynamic, accessible or inaccessible, and deterministic or non-deterministic. The third component is the specification of the interactions. These interactions can be between agents, but they can also be between agents and the environment.

By simulating an Agent-Based Model for a number of steps, one can observe the model's emerging behavior on a global level, such as risks. When a model has stochastic elements, Monte-Carlo simulations can be used to obtain numerical results for these global properties [41].

4 Case Description

This section elaborates on the qualitative phase of the agent-based safety risk assessment. We do so by first identifying the objective, second describing the operations, third selecting the hazards, and finally, presenting the scenario that was considered.

The objective of the operation: The objective is to deliver packages as safely and efficiently as possible. We measure *safety* by the expected number of fatalities per flight-hour and delivered package. We measure *efficiency* by the number of packages delivered per unit of time, or by measuring the extra distance flown compared to the shortest-path distance. These two objectives are to be balanced against each other, as it could be possible that the safest operations are not the most efficient.

Concept of Operations: Regarding operations, we assume BVLOS operations, without the need for prior approval from a UTM to fly a specific route. Regulators have indicated that it is feasible that such operations would be approved in the future [18, 19]. Also, Locascio et al. concluded that BVLOS operations are a requirement for a package delivery business case to be viable [8]. We focus on delivering small packages (<2.5kg), which are to be delivered by drones from hubs to delivery-points (DP). These DPs are points scattered throughout the city, from which people can pick-up their package. Research has shown that this scenario is more realistic than delivering directly to houses, as it is easier from a technological point of view [37]. Therefore, we chose this scenario for our research. Demand is generated using a Poisson-process and is distributed evenly over all hubs and DPs. As it is likely that drones spend a significant part of their short route in the climb-&descent phase, routes are modelled in 3D, although drones always fly their route at a fixed altitude. The FAA has indicated a maximum altitude of 400 feet above ground level, or 120 meters. Therefore, the cruise altitude is limited to 100 meters. It is also assumed that a UAS does not encounter any other traffic, whether it is another UAS or other types of traffic, such as helicopters.

Doole et al. presented an approach for forecasting the package demand in urban metropoles in the future, and forecasts a yearly demand of 35 packages per capita in Paris in 2035 [9]. Using this approach for New York and Delft leads to a yearly demand of respectively 63 and 55 packages per year per capita. We assume that there is an infinite number of drones available at each hub so that new tasks can be assigned to drones immediately.

For the drone's specifications, we selected the most advanced drone developed by DJI, the Matrice 300 RKT. This drone weighs about 9kgs, can carry 2.7kg payload, can fly up to 55 minutes, and has operating horizontal and vertical speeds of respectively 10m/s and 5m/s [42].

Lastly, we define the cities in which the operations take place. As no two cities are alike, one city's behaviour cannot necessarily be generalized to other cities. Regarding this dilemma, recent and relevant publications take one of two approaches. One approach is to analyze dummy cities, such as done by the Metropolis research project [22]. The other approach is to (arbitrarily) choose a city to analyze, such as done by two recent research works that chose to analyze Torino or Aalborg, which are home to the universities of the respective authors [36, 38]. In this work, we chose to analyze Delft, Paris and New York. These cities are of widely ranging histories, as these cities are founded millennia apart, they are on different continents, have different setups and are of different sizes. Therefore, if results generalize among these three cities, this is an indication that the results might generalize to many more cities.

Hazard: For UAS operations in urban areas, many hazards can be identified. These include a systems-failure in which the UAS fails without interacting with any other vehicle or object, but it can also collide with buildings, other vehicles, or even birds. As much is still unknown about the different failure modes for UAS operations, their corresponding probabilities and implications, the scope of this research is limited to an all-engines-out failure causing in a ballistic descent.

Scenarios: Although we have investigated many different model variants and parameters, the only safety relevant scenario considered in this research is a generic failure of all systems and engines.

5 The Agent-Based Safety Risk Model

This section outlines the steps needed to create the safety risk model, which simulates the risk associated with the operations described in the previous section. We made our implementation of the safety risk model available as an open-source tool, available on Github [43]. It is based on Mesa, an ABMS library in Python [44].

The safety risk model consists of three elements. Section 5.1 describes the agents and their interactions. Subsequently, section 5.2 describes the model of the urban environment. Lastly, section 5.3 describes the process

of calculating the risk of these operations in the urban environment.

5.1 Agents and their interactions

Three types of agents are defined in the model: the UAS, the Hub, and the Delivery Points (DPs). This section elaborates on all the states and actions of these three agents. Figure 2 gives an overview of the agent-based model. In section 5.1.1 we elaborate on the DP agent. Section 5.1.2 outlines the Hub agent. Lastly, section 5.1.3 presents the UAS agent.

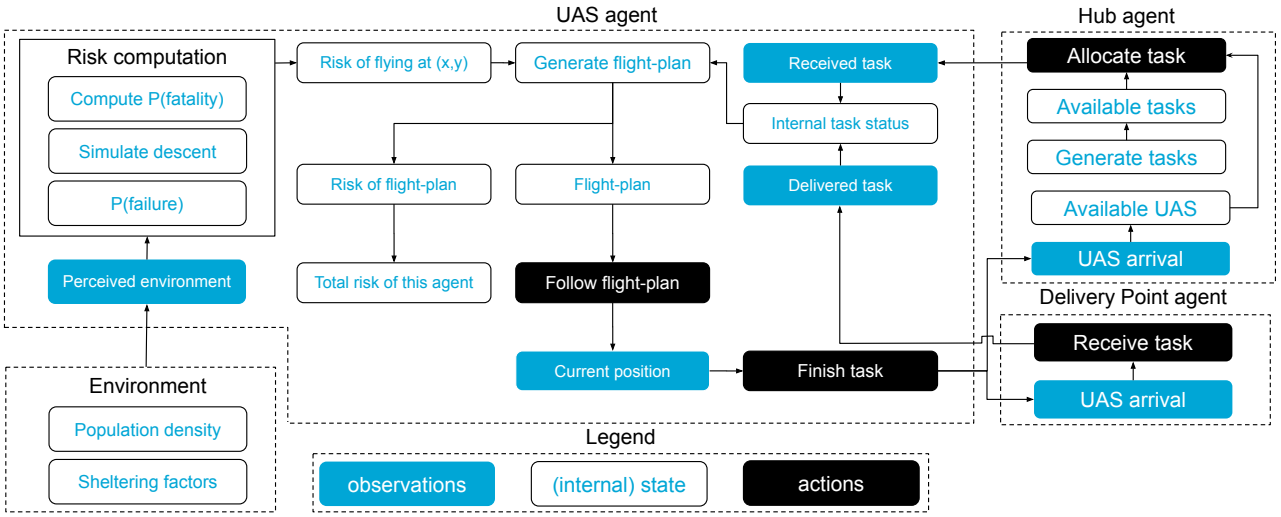

Figure 2: Overview of the Agent-Based Model

5.1.1 Delivery Points (agent)

The objective of DPs is to receive packages. This section specifies the initialization and behaviour of the DPs.

Initialization: As described in the case description, every few blocks of a city has one DP. Therefore, 100 DPs were simulated. In reality, there would be more DPs in densely populated areas than in sparsely populated areas. The number of DPs at a location should thus correlate with the population density at that area. To achieve this, we sampled the population density at each location, and then selected random positions using this population density as weights. The process for placing DPs is described in Algorithm 1.

Receive task: Upon observing the arrival of a UAS, the DP acts to receive the task. This action signals to the UAS that it has delivered a task, and it stores several parameters associated with this flight, such as the distance traveled and the risk associated with this flight.

Algorithm 1 Process for placing Delivery Points

```
DP ← set with all DPs that need to be placed
Pos ← set with all positions in the environment
PosWeightsDict ← empty dictionary
// For each possible position, store the population density around it
for P ∈ Pos do
    PosWeightsDict[P] ← No. of people within 100m of P
end for
// Randomly select position, using population density as weights
for dp ∈ DP do
    dp.pos ← random(options = PosWeightsDict.keys, weights= PosWeightsDict.values)
    Do PosWeightsDict.remove(dp.pos)
end for
```

5.1.2 Hubs (agent)

The objective of hubs is to generate tasks, and allocation these to UASs. This section discusses the initialization of the hubs, the generation of tasks, and the task allocation.

Initialization: In reality, hubs would be placed in suburbs, which are often on the edges of cities. Possibly, hubs would even be beyond the boundaries of our environment. To best reflect this situation, hubs were placed on the edges of our environment. The four hubs are each placed on one edge at a random position, provided that they cannot be placed on water, and that each edge must have exactly one hub.

Task generation: Demand is generated evenly at all hubs, following a Poisson-process. If the overall demand is denoted by λ, the demand at a hub is thus modelled by Poisson($\frac{1}{4}\lambda$). All tasks are identical, apart from their destination. This destination is determined by randomly choosing a DP. Generated tasks are added to the hub's internal state *available_tasks*.

Task allocation: Whenever there is a task in *available_tasks*, the hub attempts to assign a UAS to it. It does so by randomly selecting a UAS from its internal property *available_UAS*, which lists all available UASs.

5.1.3 UAS (agent)

The objective of a UAS is to fulfil assigned tasks while minimizing risk and maximizing efficiency. It does so by generating a flight-plan, which it then follows.

Risk computation: Based on its observations of the environment, the UAS agent can compute the risk of flying at any node (x, y). This risk computation is based on the process outlined in section 5.3, and is used both to find a flight-plan, and to store the risk associated with the operations of this agent.

Internal task status: The UAS observes whether it has received or delivered a task. Following this observation, the internal state *Internal task status* is updated accordingly.

Flight-plan generation: Based on the internal task status, the UAS generates a flight-plan to fulfil this task. To find the flight-plan, the A* algorithm is used, which is developed by Hart et al. [45]. In A*, the cost of moving $F(x, y) = G(x, y) + H(x, y)$ is used to find a solution. The algorithm's objective is to find the route that minimizes $F(x, y)$. Here, $G(x, y)$ is the risk associated with a move from $x \rightarrow y$. $G(x, y)$ is determined using the internal state *risk_x_y*. This risk is determined using the risk computation module, which is discussed in detail in section 5.3. $H(x, y)$ is a heuristic function that estimates the cost of a path. The solution of the A* algorithm is guaranteed to be optimal if $H(x, y)$ is admissible. Mathematically, this implies that $h(x_n) \leq h*(x_n)$ for each node x_n, with $h*(x_n)$ being the lowest cost of reaching the goal from x_n. As, in our environment, there is always at least one gridcell without population density and thus without risk, the heuristic has to be 0 in order to be admissible, which takes a significant toll on computing power. In our work, we set $h(x, y) = \text{add\_value} * distance(x, y)$. By varying *add_value*, the model can prioritize for distance traveled or TPR. In [46], *add_value* is varied between 0 and 1 to investigate the effect of this method on the behavior of the model.

Much effort was spent on optimizing the data-structures in the Python implementation of the A* algorithm. By using a heap-queue for the open-list and a set for the closed-list, the complexity per visited node was decreased from $\mathcal{O}(2 + logN)$ to $\mathcal{O}(2)$, with N being the size of the open list. In practice, this resulted in an approximately 200x faster computation of the paths. Without this performance increase, many of the experiments would have been computationally infeasible.

Following the flight-plan: The flight-plan is carried out in three steps: the initial climb to cruise altitude, the cruise phase, and the descent phase. In the first- and last-phases, the drone climbs or descends at a fixed vertical speed, without any horizontal speed, until it has arrived at its target altitude. In the cruise phase, the drone moves horizontally as specified in the flight-plan. The model is discretized in steps of one second, so each model step, the drone moves a distance equal to its velocity. However, this distance does not necessarily correspond with the size of the grid. Therefore, the position of the drone is defined in continuous space. At each step of the model, the drone moves in the direction of the next waypoint on its flight-plan. An example of this procedure is presented in Figure 3. The algorithm is specified in Algorithm 2. In this procedure, it is theoretically possible that the drone flies a shorter distance than specified in the flight-plan, as it "cuts off corners". However, simulation results show that the difference in flight-plan-distance and actual distance flown, is less than 0.1%.

Finish task: The final action of a UAS agent is to finish a task. This action is triggered upon arrival at the last node of the flight-plan, which corresponds with the destination of its current task. Depending on whether it arrives at a DP or at the hub, it delivers a package, or it reports itself to the hub as being available for a new task.

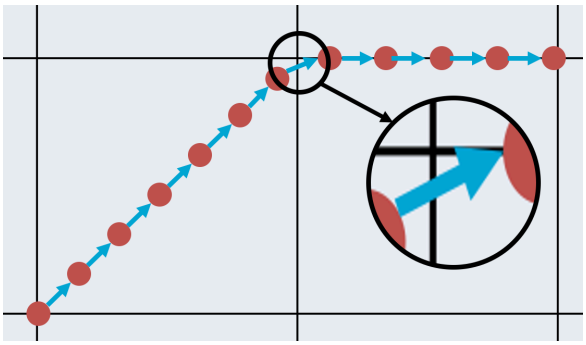

Figure 3: Example of drone movement on the grid. Waypoints are on intersection of (black) gridlines, and drone is represented as a red dot. A zoomed-in view of "corner cutting" is presented.

Algorithm 2 Process for moving in continuous and discrete steps

```
posGrid ← (discrete) position on the grid
posVirtual ← (continuous) position
flightPlan ← list of (discrete) points to be followed
destinationPos ← last item in flightPlan
// Move step in direction of the nextPos in continuous space
// If within one step of nextPos, move posGrid
while posGrid is not destinationPos do
    nextPos ← next element in flightPlan
    heading = angle(posVirtual, nextPos)
    posVirtual = posVirtual + speed * heading
    if distance(posVirtual, nextPos) < speed then
        posGrid = nextPos
    end if
end while
```

5.2 Model of the urban environment

After presenting the agents and their interactions, the remaining component of the agent-based model is the environment. In this work, a static and accessible environment is modeled, implying that agents can access all information at all times.

Constructing accurate environments is vital for a realistic safety analysis of UAS operations. La Cour-Harbo proposed a method in which they form the environment by stacking four maps on top of each-other [36]. These comprised of firstly the population density map, which specified how many people are present at each grid-cell. The second map specifies the sheltering factor, which is a factor describing to what extent people are protected from an impact because of buildings, trees, or vehicles. The third and fourth map respectively specified the no-fly zones and all obstacles present. La Cour-Harbo used realistic population density data from parts of Denmark but used a constant sheltering factor. Primatesta used this approach with accurate data from a part of Turin, Italy [30].

In this research, we decided to focus our efforts on the population density map and the sheltering map, as these maps influence the TPR. We have generated accurate representations of these maps for not one, but three cities. These maps are then discretized in cells of approximately 10*10 meters, as described in Figure 4.

In respectively sections 5.2.1 and 5.2.2, it is explained how these maps were constructed as realistically as possible. Because humans are constantly on the move, the population density map can assume many forms. Section 5.2.3 explains how this is included in the model.

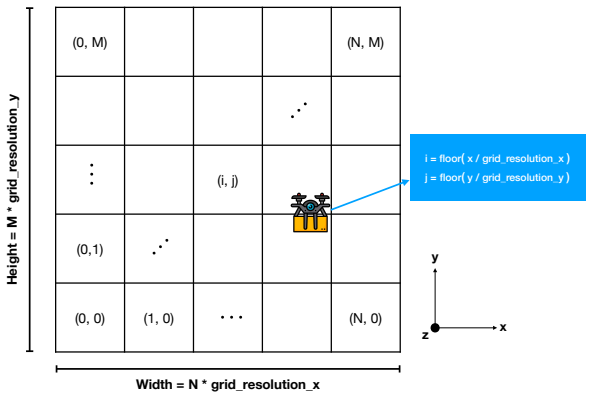

Figure 4: Discrete representation of the environment in a grid.

5.2.1 Population Density

A literature survey of 17 UAS Safety studies shows that models commonly assume a uniform population density [28]. In recent years, some models have implemented population density based on Census data [47, 38, 30]. Although not perfect, as it only reflects where people live and not where they are, using Census data is a state-of-the-art approach which is used in this work. We retrieved the Census data from New York, Paris, and Delft from the relevant authorities [48, 49, 50]. A graphical representation of the population density map was created, which was subsequently cropped to achieve an exact match down to pixel level with the sheltering map. This input map is presented in fig. 5(a). Each unique color was assigned a value. For example, a black pixel indicates a population density of 0. In this way, all pixels are converted into values for population density. The exact process is described in the appendix [46]. As an example, the population density map for New York is given in fig. 6(a).

5.2.2 Sheltering Factor

The sheltering factor describes to what extent a person is shielded from an impact by objects, such as buildings, trees or vehicles. The precise model underlying the sheltering factor is discussed in the section on Risk (section 5.3). In this work, OpenStreetMap data was interpreted using visuals to arrive at sheltering data such as in fig. 6(b) [51]. It was decided to use visuals, rather than APIs, as it is necessary that the positions on the sheltering map exactly match the positions of the population density map. An example of such an input map is in fig. 5(b). Just as for the population density map, each unique color is mapped to a sheltering factor. For example, a green pixel corresponds to a sheltering factor of 2.5.

Another issue is that maps-data does not distinguish between low- and high buildings. Therefore, a third map is generated, based on Google Streetview observations, which was then used to classify buildings as either high or low [52]. Every time a location (x, y) is classified as a building based on the sheltering map, the high- vs low-buildings map is then consulted to correctly classify the sheltering factor. The resulting heatmap of the risk per location is given in fig. 6(c). In the appendix, we elaborate on the specifics of converting the input data into an environment [46].

Sheltering	Area
0	No obstacles
2.5	Sparse trees
5	Vehicles and low buildings
7.5	High buildings
10	Industrial buildings

Table 1: Sheltering factors as used in [30].

5.2.3 Environment Modification

In the real-world, the sheltering factors are approximately constant, while the population density would change constantly. It would change both on a daily level, such as in phases as 'office-hours' and 'rush-hour', and it would differ greatly on a longer time scale, for example during holidays, seasons and extreme events such as pandemics. Even weather may significantly impact the population density map. Furthermore, the input data

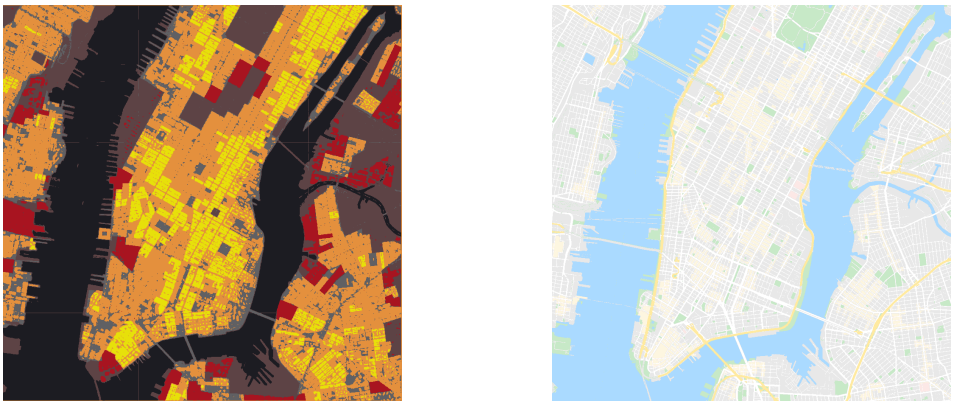

(a) Input map for population density [50]

(b) Input map for sheltering factors [51]

Figure 5: The figures used as input data for New York.

(a) Heatmap of population density (people / m^2)

(b) Heatmap of sheltering factors

(c) Heatmap of risk (expected fatalities per impact)

Figure 6: The environmental representation of New York.

for the population density maps only consider where people live. This section describes steps taken to increase the accuracy of the population density map.

In order to more accurately reflect a population's movement, Melnyk et al. proposed to incorporate statistical data on the whereabouts of people [53]. They calculated the average time spent at a location (such as an office) throughout the day. However, human movement differs greatly per moment, with stark differences emerging between measurements at for example rush-hour, working times, nighttimes or weekend. To reflect these differences, we created a bottom-up estimation of the same data, for these different phases of the day, according to the procedure specified in the appendix [46]. These estimations are presented in Table 2. Using this data, we can simulate several scenarios and compare the resulting model behavior.

Information	**Melnyk et al.**	**Office hours**	**Nighttime**	**Rush hour**	**Office hours with increased population on water**
In residence	68.7%	54.9%	95.8%	62.6%	50.9%
Office / Factory	7.6%	16.9%	3.5%	4.2%	14.9%
Vehicle	5.5%	14.1%	0.4%	3.5%	12.1%
Outdoors	5.4%	4.9%	0.3%	8.8%	4.9%
Transportation	12.8%	9.1%	0%	16.9%	9.1%
Water	0.0%	0.0%	0.0%	3.9%	8.0%

Table 2: Data on the percentage of people per type of location in different phases of the day.

5.2.4 Scaling the Environment

So far, this section explained the process of creating maps based on images of the cities. Often, these images would have over 1 million pixels, causing the environment to have over 1 million nodes. As our model has stochastic elements, it is required to run Monte Carlo simulations, and for several tasks, such as path-finding,

it is computationally infeasible to do so on this large map. Therefore, a process for scaling these maps was set up in such a way that as much data as possible is retained. For scaling a map with a factor N, each pixel in the scaled map was based on an NxN pixels view, denoted V, in the original map (Figure 7). After experimenting with different methods, it was decided that the population density map should be scaled by taking the average of V, while the sheltering map should be scaled by taking the mode of V. The remainder of this section explains the rationale behind this process, and Figure 7 shows an example of how a 4x4 map is scaled by a factor 2, using both average and mode. Lastly, a sensitivity analysis was performed, which is presented in the appendix [46]. This analysis concluded that downscaling the environments with a factor 5, thus scaling a 1,000 x 1,000 pixels to 200 x 200, has only a minimal effect on the results of the simulation [46]. It is thought that this effect is so minimal because the resolution of the original map is higher than the data itself (pixels are defined on nearly a house-by-house level, while population density and sheltering are defined on a block-by-block level).

For the sheltering map, using the mode of the view V to generate the scaled pixel is advantageous as doing so retains the discrete sheltering category of the old map. Taking the average would lead to some in-between category, where a pixel is not, for example, a vehicle nor a building, but something in-between. Knowing the distinct sheltering category is especially important for keeping track what type of location the UAS is flying over, which is relevant for statistics.

The population density map was scaled by taking the average of an area. This was deemed the most realistic solution, because using the average ensures that the total number of people in an area does not change because of scaling.

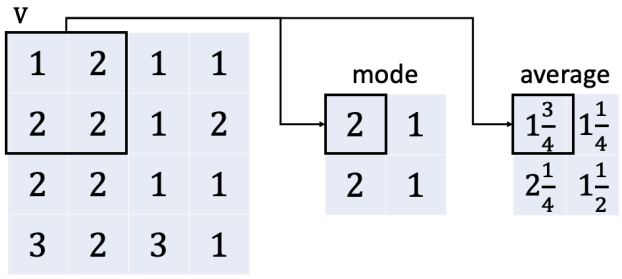

Figure 7: Example of scaling a 4x4 map by factor 2 with both "mode" and "average" method.

5.3 Model for assessing Third-Party-Risk

Recent work on assessing the TPR following from (urban) UAS operations use slightly different, but mathematically equivalent, expressions [36, 30, 54]. In this research, we define the TPR such as in eq. (1). In this equation, a path p represents an individual flight operation, and all paths of all vehicles in the environment are combined in the set P, such that $p \in P$. M is the set of all (x,y) coordinates on the map. Furthermore, eq. (2) represents the total risk of all paths in a simulation. This equation comprises of three elements. The first is the probability of failure. The second element is the probability that the drone will crash at a location (x,y), given a failure. Or in other words, it describes what the descent trajectory looks like. The third element considers the expected number of fatalities, given this impact. The model for determining these three elements is discussed in the following three sections. The final section reflects on the computational implementation of these methods.

$$\mathbb{E}[\text{Risk(p)}] = \underbrace{\mathbb{P}(\text{failure})}_{\text{section 5.3.1}} \times \int_{(x,y)\in M} \underbrace{\mathbb{P}(\text{impact at x,y})}_{\text{section 5.3.2}} \times \underbrace{\mathbb{E}\left[\text{fatality} \mid \text{impact at (x,y)}\right]}_{\text{section 5.3.3}} \tag{1}$$

$$\mathbb{E}[\text{Risk}] = \sum_{p\in P} \mathbb{E}[\text{Risk(p)}] \tag{2}$$

5.3.1 Probability of Failure

In the absence of historical failure data, one should revert to other ways of estimating the probability. The space industry, where failure data is also limited, provides inspiration on how to approach such situations [55, 56]. Approaches include the decomposition of the system into subcomponents. This allowes to obtain the failure rate on system-level by analyzing the failure rate on the level of its subcomponents. This method has also been applied to UAS research [57]. Another method is using Bayesian Belief Networks, which is discussed in several recent publications in this field [58, 59, 60]. A literature-survey of 17 recent papers discussing UAS TPR concludes that all papers assume the failure rate to be constant throughout the flight. However, in GA, almost half of all accidents happen in the very first- or last minutes of a flight [28, 61].

Despite ongoing development in accurately modelling the probability of failure, recent UAS Safety analyses mostly use "expert opinion". Studying recent publications shows mean-time-between-failures MTBF values differing by a factor 1,000 (10^3 - 10^6 hours), and therefore these "expert opinions" can be interpreted as wild guesses [62, 63, 64, 36, 30]. Because the goal of this research is to draw conclusions by analysing relative differences between scenarios, and not to find an estimation of the absolute value of risk, using a different failure rate has no impact on the overall findings. Therefore, the arbitrary value of one crash per 500 hours is chosen. It is noteworthy that no literature was found that studied the effect of varying failure rates per flight phase on the safety of UAS operations, and therefore, it was worth studying this by means of a sensitivity analysis in our research.

5.3.2 Descent Trajectory

When analyzing UAS safety, one can distinguish different modes of failure. In literature, four failure modes are usually identified [29, 27, 30]. These are (1) an unpremeditated descent trajectory (gliding descent), (2) loss of control, (3) controlled flight into terrain, and (4) dropped or jettisoned components. A literature survey by Washington et al. shows that the first two failure modes are used in practically all research, and relevant literature published since then confirms this conclusion. Furthermore, a gliding descent is used for fixed-wing UAS, and a loss-of-control scenario (resulting in a ballistic descent) is used for multi-rotor UAS [65, 28, 30, 54]. The list of descents is not exhaustive. Other types exist as well, such as parachute descents (studied in [66, 67, 30]) and flyaways [68, 30].

As mentioned in the case description, we limit us to a ballistic descent following from an all-engines-out failure. This ballistic descent is based on a second-order drag model as in eq. (3) [69], where m is the mass of the UAS, g is the gravitational acceleration, C_D is the drag coefficient, A_D is the frontal surface area, ρ is the air density, and v is the velocity of the vehicle. The drag coefficient is taken to be $C_D \sim \mathcal{N}(0.7,\, 0.1)$.

$$m\dot{\boldsymbol{v}} = mg - C_D A_D \rho |\boldsymbol{v}| \boldsymbol{v} \tag{3}$$

Even though it is assumed that wind does not influence the cruise-speed nor the probability of failure, it is assumed that the drone can be subject to wind during the descent trajectory. In this scenario, the aerodynamic forces are manipulated to take this wind into account. Wind is only considered in the scenario of Delft. Here, five years of real wind data are used, consisting of 50,000 samples of wind-speed and corresponding direction, as provided by the Dutch Meteorological Institute (KNMI) [70]. However, this data considers the free-field wind, which is not the same as the wind in an urban setting. Research has shown that the free-field wind can be extrapolated to urban wind using eq. (4). Here, v_0 is the wind at the height of H_0 (typically 10 meters), in an unobstructed field. Measurements show that $\alpha = 0.3$ gives a good approximation of the wind in urban areas [71].

$$\frac{v}{v_0} = \left(\frac{H}{H_0}\right)^{\alpha} \tag{4}$$

5.3.3 Probability of Fatality

After having modelled the probability of a failure and the descent trajectory following that failure, this impact should be converted to an expected number of fatalities. This calculation comprises of three components, as can be seen in eq. (5). Here, A_{impact} is the area concerned in this impact (in m^2), $\rho(x, y)$ is the population density at this location, and $\mathbb{P}(\text{fatality} \mid \text{impact})$ encompasses the probability that a person at this location dies, following this impact. As an example, if the impacted area is 1 m^2, the probability of fatality is 50%, and the population density is two people per m^2, then the expected number of fatalities is one. As the method for obtaining the population density is already explained previously, the remainder of this section presents our model for determining the impact area and the probability of fatality.

$$\mathbb{E}\left[\text{fatality} \mid \text{impact at (x,y)}\right] = A_{\text{impact}} * \rho(x, y) * \mathbb{P}(\text{fatality} \mid \text{impact}) \tag{5}$$

In our literature survey [46], we presented an extensive comparison of several methods of assessing this impact area, and concluded that the equation such as in eq. (6) is the most realistic. R_p is the radius of a person, R_{uav} is the radius of the UAV, H_p is the height of a person, and γ is the glide angle. Experimentation shows that during the cruise-phase, impacts have an angle of approximately 78 degrees, while the impact angle is exactly 90 degrees during the climb-&descent-phase. This results in an impacted area of about 0.50 m^2 during the climb-&descent-phase, and 0.80 m^2 during the cruise-phase.

$$A_{\text{exp}}(\gamma) = \pi \left(R_p + R_{\text{uav}}\right)^2 \sin(\gamma) + 2 \left(R_p + R_{\text{uav}}\right) \left(h_p + R_{\text{uav}}\right) \cos(\gamma) \tag{6}$$

All methods for assessing the probability of fatality, given an impact, are based on research by Feinstein et al., who derived a log-normal distribution back in 1968. This relation correlates impact energy to the probability of fatality. They arrived at this distribution using data provided by the US Department of Defense data, which linked injury severity to missile impact data [72]. Based on this work, the Blunt Criterion (BC) was introduced, connecting blunt impacts to the chest to fatality [73]. The BC was then extended to head injuries and applied to UAS operations in 2013 [74, 75], and was compared to experimental data to connect the BC to the Abbreviated Injury Scale (AIS), to categorize injuries from "none" to "virtually unsurvivable" in 6 discrete classes [76, 77].

Subsequently, Dalamagkidis et al. provided a novel relation specifically aimed at the field of UAS research, which also incorporates a sheltering factor. This equation is considered state-of-the-art, being used in recent UAS research [36, 30], and is useful for urban operations because of the presence of a sheltering factor. Common values for the sheltering factor are given in Table 1.

The fatality model is presented in equations 7 and 8. In it, p_s is the sheltering factor, taking values $(0, \infty)$, with higher values implying a better protection. α is the energy required for a fatality probability of 50% when $p_s = 6$, and the model approaches other models when $\alpha = 32kJ$. Furthermore, β is the impact required to cause a fatality when $p_s \to 0$, which is 34J [78, 79].

$$P(\text{ fatality } | \text{ exposure }) = \frac{1-k}{1-2k+\sqrt{\frac{\alpha}{\beta}}\left[\frac{\beta}{E_{\text{imp}}}\right]^{\frac{3}{p_s}}} \tag{7}$$

$$k = \min\left[1, \left(\frac{\beta}{E_{\text{imp}}}\right)^{\frac{3}{p_s}}\right] \tag{8}$$

There are several limitations to this approach. Two of the most important ones are, firstly, that this relation was still calibrated using missile impacts, which are substantially different from UAS impacts, and secondly, that the sheltering factors are chosen based on little to no evidence underlying their validity.

5.3.4 Computational considerations

So far, we have established the theory for computing TPR. This section elaborates on the implementation in the model. For each step a UAS takes on its path, the expected value of the risk is computed by simulating a failure at that location. The ballistic descent trajectory is discretized to find the crash location. The drag coefficient and the wind influence this descent trajectory, and an option was modeled to either choose fixed values for the drag and wind, or to model them as stochastic elements. In the latter scenario, an MC simulation is performed, in which values for the drag and wind are sampled from the distributions specified in section 5.3.2. Stable results were found using 5,000 MC simulations and $dt = 0.1$ for steps in the descent trajectory.

This risk computation is computationally expensive. However, the descent trajectory only depends on the velocity-vector $\boldsymbol{v}$ and initial altitude z and is thus invariant to the position of failure (x, y). Because of this, a descent trajectory that is calculated for position (x_1, y_1, z_1) and $\boldsymbol{v_1}$, is identical to the descent trajectory calculated at (x_2, y_2, z_1) and the same speed $\boldsymbol{v_1}$. Due to the shape of the grid, constant cruise altitude and constant velocity of the drones, many risk profiles are thus identical. We employed an algorithm that stored the descent trajectories in a database, and by doing so, the time spent computing the risk was decreased by approximately a factor of 100 per simulation.

6 Verification and Validation

This section elaborates on the verification- and validation-approaches. These approaches were used to thoroughly scrutinize the model to increase confidence in the conclusions of this research. Klügl et al. proposed an iterative validation framework, in which verification and validation are conducted during model development [80]. The code was verified by performing unit- and system-tests, and by resolving all compile errors [81].

In the absence of real data, one should resort to different methods for validation. We validated the model both through sensitivity analysis, and through a critical inspection of the results found. In the next sections, discussing the experiments conducted, the results of each experiment are scrutinized. For any result that could not easily be explained through reasoning, we performed additional experiments to find explanations.

In the future, it is likely that more real world data will become available that can be used to validate elements of this model. For example, data from flight-tests could be used to validate the probability of failure and the dynamics of the drone. Also, more detailed information on the movement of people would allow the validation of the population density map. Another validation direction considers the model of the probability of fatality, which could be validated by conducting crash-tests into buildings and human crash-dummies [82].

7 Experimental approach

This section presents an overview of the experimental approach. Firstly, section 7.1 presents an overview of all experiments conducted. Secondly, section 7.2 discusses the performance indicators used throughout the experiments. Finally, section 7.3 presents the statistical methods used to analyse the results.

7.1 Overview of Experiments

This section provides an overview of the seven experiments discussed in this research. These experiments are centered around the main contributions of our work. The first contribution considers the behavior of a state-of-the-art TPR computation in a detailed model of the urban environment. This is investigated in experiment one (E1), by studying how local features of an environment influence model behavior. The second contribution studies the impact of individual elements of the risk computation, on the overall risk. This is analysed in experiments two to five (E2 - E5). Furthermore, two experiments were conducted that investigate variations in the operational concept. These are experiments six and seven (E6 and E7), and their findings are summarized briefly in this paper and discussed in detail in the appendix [46].

Table 3 presents a tabular overview of the aforementioned seven experiments. In fig. 8, the place of each of the experiments in the different flight-phases is shown.

	Changed variable	Discussed in
E1	Several environmental parameters	section 8.1
E2	Failure rate in the climb&descent-phase	section 8.2.1
E3	Velocity at impact	section 8.2.2
E4	Several parameters that are part of the fatality model	section 8.2.3
E5	The presence of wind and drag during the descent trajectory	section 8.2.4
E6	The cruise altitude	section 8.3
E7	Balancing the weight of effectiveness and safety in the path-finder	section 8.3

Table 3: Overview of all experiments.

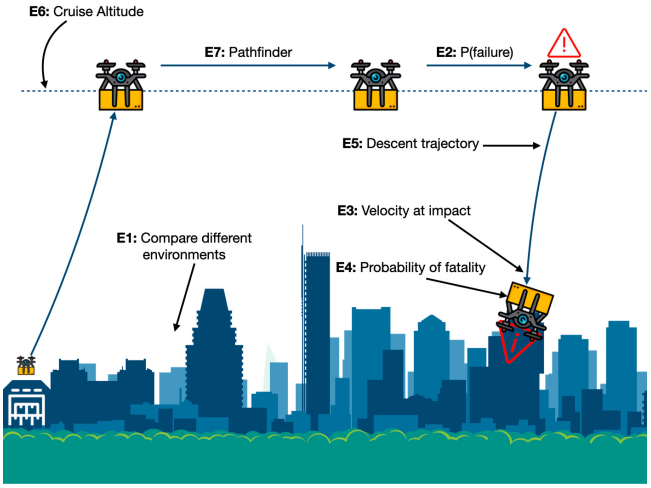

Figure 8: Schematic overview of all experiments.

7.2 Performance indicators

We identify three performance indicators that are studied in the experiments. The first is the number of fatalities per 1 million flight-hours (FAT-1M), which provides insight into the overall safety. The second is the ratio of the flown distance, divided by the shortest-path distance, which is a measure of the efficiency of the system (DIST). The third is the number of fatalities per completed task (FAT-TASK), a measure which combines both the objectives on safety and efficiency.

7.3 Statistical approach

For each simulation, the coefficient of variation was calculated to evaluate the necessary number of simulations. Subsequently, we concluded that the results of all experiments follow a normal distribution, by using both the

Shapiro-Wilk test and by inspecting the QQ-plots [83]. The Shapiro-Wilk test was chosen as it is more powerful than similar tests, such as the Kolmogorov-Smirnov test [84]. After having established normality, a t-test was used to compare the results of two experiments head-on. Also, performing multiple tests on the same dataset increases the probability of making a Type 1 error. This was mitigated by using a Bonferroni correction [85].

To study how model behaviour changes upon changing parameters in the model, a sensitivity analysis (SA) was carried out. The most straightforward way is to vary elements one at a time (OAT) while keeping other elements constant. One can then measure the change in output to a 1% change in input, which is also called a local sensitivity analysis (LSA). However, Saltelli et al. argue that this method has certain shortcomings, that become more pronounced when multiple parameters are inter-related and are non-linear [86]. In such a situation, a global sensitivity analysis (GSA) is a more appropriate method [87, 88]. However, the downside of a GSA is that it is computationally expensive. Comparing several input parameters simultaneously leads to what is called *"the curse of dimensionality"*, where the sample space grows exponentially for each additional parameter. To counteract this problem, several sampling techniques exist that efficiently sample values from the domains of the input parameters. These techniques ensure that the result of the GSA converge to its true mean as quickly as possible. One of these techniques is Latin Hypercube Sampling (LHS) [89], which we used in the GSA in this research. Furthermore, in situations where linearity cannot be assumed, Saltelli et al. propose to use rank correlation, in which samples are replaced by their rank [90]. Spearman's rank correlation coefficient is then used to measure the importance of each parameter [88]. An elaboration on the tools mentioned in this section can be found in the appendix [46].

8 Experimental set-up and results

This section presents the detailed set-up of the experiments, the results following from these experiments, and a critical assessment of these results. Section 8.1 presents experiment one, which considers the environment. Section 8.2 presents experiments two through five, considering the risk computation. Finally, experiments six and seven are mentioned briefly in section 8.3.

8.1 Influence of environmental parameters on model behavior (E1)

To what extent do different elements in the environment influence the results of our experiments? Knowing the answer to this question could help advance the safety of UAS operations, for example, through the smart placement of hubs, delivery-points, or (no-)fly-zones. This section elaborates on experiments carried out to answer precisely this question. First, we will discuss how environments can be characterised using different parameters. Second, we present an approach for changing these parameters. We observed how changes in these environmental parameters correlate with changes in the performance indicators. Finally, we present the results following from these experiments. In the following, the risk of an individual grid-cell is defined by the expected number of fatalities, given an average UAS impact at this grid-cell.

8.1.1 Characterising the environment

To analyse the influence of individual environmental parameters on the global model behaviour, five aggregate parameters were developed that characterise the environment. These five parameters consider general properties of the environment, properties focusing on the low-risk areas, and properties relating to the location of hubs and DPs. Table 4 lists these parameters, and we elaborate on each of these parameters in the remainder of this section. The first parameter is an aggregate of the entire environment:

1. *avg_risk_env:* This parameter represents the average risk of the environment, which is calculated as the average risk of all grid-cells. We hypothesise that the lower the average risk in the environment, the lower the Fat-1M will be.

Parameters two and three consider low-risk areas. These are the cells on the map where the risk is the lowest. While observing the model, UAS agents showed emergent behaviour of gravitating towards these low-risk areas. This observation was confirmed quantitatively by the finding that agents fly between 30% and 50% of their mission over the gridcells in the lowest risk-quartile. To classify which gridcell is low-risk and which gridcell is not, a threshold value is used, called T. A grid-cell is considered to be "low-risk", if it falls in the T%-percentile, i.e., the grid-cell belongs to the lowest T% of grid-cells. In this experiment, T is set to 20%, and a sensitivity study was performed with values of 10% and 25%. In Figure 9, histograms of the risk of all grid-cells are presented for two population density scenarios in New York. In these figures, the percentiles are marked. Parameters two and three are defined as follows:

2. *avg_risk_lowrisk:* This parameter computes the average risk of the low-risk areas, and it is hypothesised that the average risk of the low-risk areas correlates with the FAT-1M.

3. *ann_lowrisk:* It is hypothesised that the more clustered these low-risk-areas are over the map, the higher the risk is, as it makes these low-risk areas less accessible for drones. The clustering (or dispersion, oppositely) of low-risk areas is calculated using the Average-Nearest-Neighbor-index (ANN), and is called *ann_lowrisk*. This metric was developed by Ord et al. and is frequently used in the geographical- and topological studies [91, 92]. The ANN is calculated as in Equation 9. Here, $\bar{D}_O$ is the average distance to the nearest neighbour for all involved gridcells, and is calculated per Equation 10, with d_i being the distance to the nearest neighbour from node i. $\bar{D}_E$ is the distance to the nearest neighbour, assuming all nodes are distributed uniformly, and is calculated using Equation 11. A lower value implies more clustering. An example of two situations of the ANN is given in figures 10 and 11. Here, it is visible that indeed, the routes gravitate towards the low-risk areas.

$$ANN = \frac{\bar{D}_O}{\bar{D}_E} \tag{9}$$

$$\bar{D}_O = \frac{\sum_{i=1}^{n} d_i}{n} \tag{10}$$

$$\bar{D}_E = \frac{0.5}{\sqrt{n/A}} \tag{11}$$

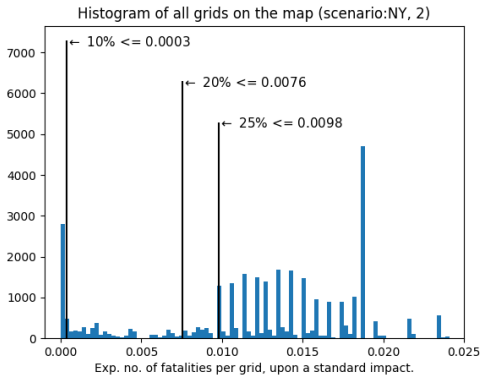

(a) No population on water

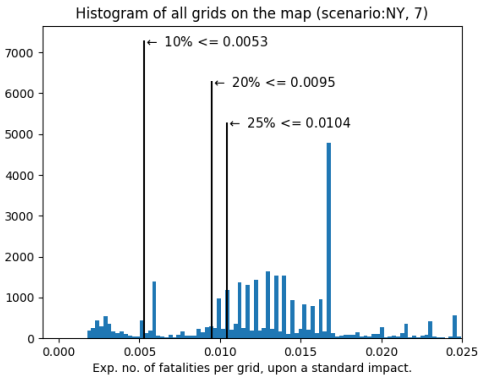

(b) No population in parks

Figure 9: Histogram of the expected number of fatalities upon an average impact for all grid-cells in the city of New York, for two population density scenarios.

Parameter	Description	Category
avg_risk_env	Average risk in the environment	General
ann_lowrisk	Average Nearest Neighbour index of low-risk areas (measure of dispersion)	Low-risk areas
pct_lowrisk	Percentage of gridcells that is in the category "low-risk"	Low-risk areas
hub_risk	Average risk in area around hubs (<50 meters)	Hub / DP
DP_risk	Average risk in area around DPs (<50 meters)	Hub / DP

Table 4: Overview of parameters used to characterise the environment.

Parameters four and five consider the hubs and DPs. We have selected these parameters, as these locations are part of every flight. Knowing the influence of these parameters on the global risks gives insight into the importance of hub- and DP-placement on the risks. Parameters four and five are defined as follows:

4. *hub_risk:* This parameter represents the average risk of the grid-cells around all hubs.
5. *DP_risk:* This parameter represents the average risk of the grid-cells around all DPs.

8.1.2 Changing the environment through variations in population density

To study the hypotheses presented in the previous section, it is needed to vary the environment. The environment comprises two elements: the topological structures that are the input for the sheltering map and the population density. The environment can only be changed by changing one of these two elements.

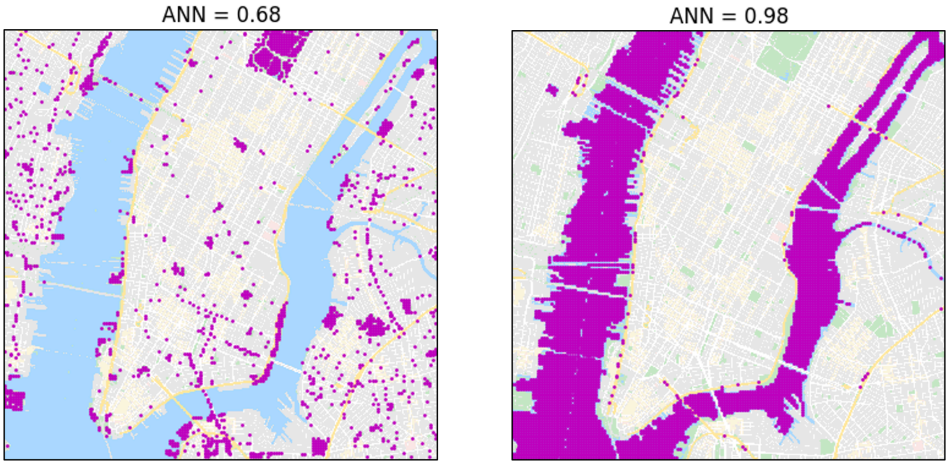

(a) Scenario with no population in parks (b) Scenario with no population on water

Figure 10: ANN of two scenarios in New York City with low-risk areas scattered in purple.

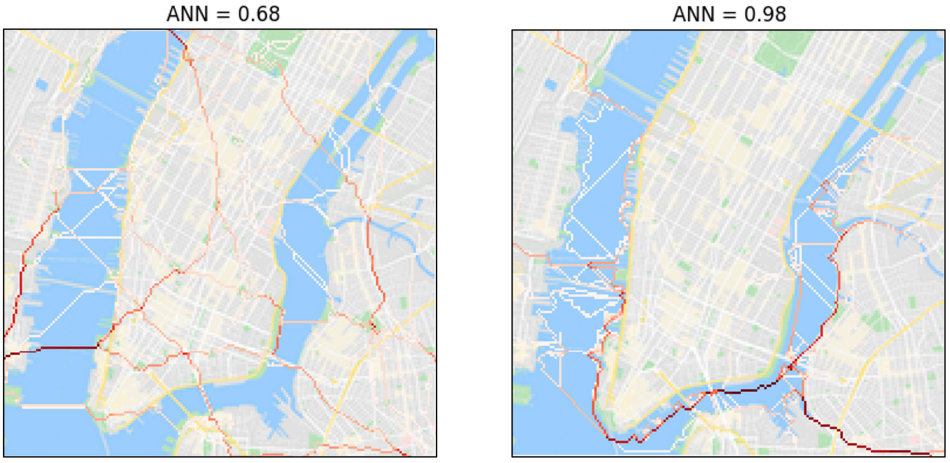

(a) Scenario with no population in parks (b) Scenario with no population on water

Figure 11: Heatmap of flight movements in two scenarios in New York (dark red = more flight movements).

The environment should be changed as realistically as possible. Changing topological structures in a city would not be realistic. Take, for example, New York. If we would remove high buildings or the main river from the environment, the city would be dramatically different, and it would not make any sense to use it in our analysis.

However, changing the population density map is quite realistic, as this happens all the time in reality. Therefore, the environment is changed by changing the population density. In the case description, Table 2 gives four scenarios of how population density could be distributed in a city, based on a high-fidelity bottom-up estimation. As transitions between scenarios such as "rush hour" and "daytime" happen gradually, it is also sensible to take the average between these scenarios, giving us additional scenarios for this experiment. We thus have four scenarios, being "daytime", "rush-hour", "nighttime" and "additional water presence". We also have six additional scenarios, being the average between the four aforementioned scenarios, such as "daytime"+"rush-hour". Only the combination of "nighttime" and "daytime" is not averaged, as there is no continuous overlap between the two, in reality. That leaves us with nine scenarios, comprising four direct scenarios, and five averaged scenarios. To improve robustness of the experiment, noise is added by adding $X \sim \mathcal{N}(0, 0.015^2)$ to each of the percentages. Subsequently, the percentages in each scenario are normalised to ensure that they still sum to 100%.

For the nine scenarios, the environmental parameters described in the previous section can be calculated. Instead of combining the data of Delft, New York and Paris, this experiment was run separately in each of the three cities. Afterwards, the results were compared, to observe whether or not findings from one city, generalise for other cities.

8.1.3 Results for changing environmental parameters

Nine scenarios have been evaluated for three cities, being 27 scenarios in total. We have simulated each scenario 60 times, leading to a total of 1620 simulations. One simulation takes about 20 minutes to run on an Intel Xeon Gold 6200 CPU (3.70 GHz). The results are presented in Table 5. In this table, Spearman's rank correlation coefficient is provided for each of the five parameters, in each of the three cities. This coefficient measures the correlation of a change in a parameter to the fatalities per 1m flight-hours (Fat-1M). A correlation coefficient between 0.3 - 0.7 indicates a moderate relationship, while a value between 0.7 and 0.1 indicates a strong relationship [87]. Furthermore, the table indicates whether or not the observed correlation coefficient is statistically significant after applying a Bonferroni correction, where a result is deemed significant for p-values <0.01.

For two of the parameters, a statistically significant correlation coefficient is observed in all three cities. Despite the correlation coefficient indicating a stronger relationship in one city than another, the correlation is in the same direction (negative or positive) in each city nonetheless. Furthermore, this experiment was repeated for different threshold-values T. This value is used to classify low-risk areas and is defined as 20% by default. A sensitivity study, which is presented in the appendix [46], shows that the conclusions regarding the correlation remain unchanged after varying the threshold value to 10% and 25%.

In the remainder of this section, we will discuss the results for each of the parameters individually. The first of these is *avg_risk_env*, with a positive correlation coefficient for Delft and New York. This implies, unsurprisingly, that when the average risk of the gridcells becomes higher, the system risk does so, too. However, it is noteworthy that no statistically significant correlation was found in the city of Paris. A possible explanation can be found by comparing the maximum- and minimum-value of *avg_risk_env* for all three scenarios. In New York, there is a 31% difference between the maximum- and minimum value, in Delft, the difference is 40%, but in Paris, the difference is only 6%. It is possible that no correlation is observed in Paris because *avg_risk_env* is similar in all simulations.

We find a negative correlation between *ann_lowrisk* and the fatalities per flight-hour in all cities. This negative correlation implies that whenever the low-risk areas become more dispersed over the environment, the system risk decreases, as was hypothesised. This could be explained by the fact that agents have easier access to the low-risk areas if these are more dispersed.

In all environments, *avg_risk_lowrisk* shows the strongest correlation of all parameters. This result confirms the hypothesis that an increase in the risk of low-risk areas causes an increase in the observed system risk. It is thought that this correlation is stronger than *avg_risk_env* because drones spend a disproportionate amount of the flight-time over the low-risk areas.

The remaining two parameters consider the risk around DPs and Hubs. In Delft, a higher *avg_risk_DP* leads to a lower system-risk. No satisfactory explanation is found for this observation. Also, a sensitivity study, which is presented in the appendix [46], showed little or no correlation of *avg_risk_DP* for different values of threshold value T. In Delft and New York, a higher *avg_risk_hubs* leads to a higher system risk. We believe that this is the case because the hubs are encountered on every flight, and therefore, a higher risk at the hub directly increases the risk of each flight.

Correlation coefficient	y = fatalities per 1m flight-hours (FAT-1M)					
	Delft		New York		Paris	
x parameter	*conclusion*	*r-value*	*conclusion*	*r-value*	*conclusion*	*r-value*
avg_risk_env	↑↑	0.42	↑↑	0.51	–	0.06
ann_lowrisk	↓↓	-0.64	↓↓	-0.35	↓↓	-0.32
avg_risk_lowrisk	↑↑	0.44	↑↑↑	0.71	↑↑↑	0.87
avg_risk_DP	↓↓	-0.35	–	-0.01	–	-0.10
avg_risk_hubs	↑↑	0.34	↑↑	0.45	–	0.16

Table 5: Spearman's rank correlation coefficient of environmental parameters with respect to number of fatalities per task. Arrow indicates direction and strength of correlation, and when an arrow is indicated, correlation is statistically significant after applying Bonferroni's correction. One arrow = weak correlation, two arrows = moderate correlation, three arrows = strong correlation. – implies no (statistically significant) correlation.

8.2 Risk Computation

This section presents experiments two through four, which each study a component of the risk computation. The probability of failure (experiment 2) is studied in section 8.2.1, the impact velocity (experiment 3) is discussed in section 8.2.2, and the fatality model (experiment 4) is presented in section 8.2.3.

8.2.1 Varying the failure rate (E2)

In General Aviation (GA), close to 50% of all failures happen either shortly after takeoff or shortly before landing [61]. As explained in the section on the model, little is yet known on failure rates for UAS, making it impossible to conclusively attest if a similar effect would be realistic during the flights of UAS. To discover how much such a change would influence the model, two scenarios are compared. In scenario 1, the failure rate is constant throughout the flight. In scenario 2, the failures are split evenly between the c&d-phase (climb&descent-phase) and the cruise phase. This is achieved by manipulating the failure rate in the c&d-phase. From experiments, it is observed that the vehicles on average spend 7.2 times as much time in the cruise phase compared to the c&d-phase. Therefore, to obtain an equal number of failures in both phases, the failure rate in the c&d-phase should be 7.2 times higher than in the cruise phase.

The hypothesis is that the overall Fat-1M is higher, when the failure rate is higher in the c&d phase. This is indeed found to be the case, with a 35.1% increase in Fat-1M (95% CI [31.3%, 38.9%]). A 35% change in Fat-1M seems modest, as the failure rate in the c&d-phase is increased by 610% (a factor of 7.1). This is caused by the following two reasons:

1. Only 12.2% of the flight-time is spent in the c&d-phase. This means that the higher failure rate is only encountered in a short part of the flight. This effect is clearly visible in fig. 12(a), where the c&d phase only accounts for a small part of the total risk.
2. As can be seen in fig. 12(b), a crash in the c&d-phase is only 46.9% as risky as an impact during the cruise-phase. This is caused by three reasons:
   a) The areas around hubs and DPs are safer than the average area on the map. This is mainly caused by the hubs, which are placed in rural areas.
   b) Failures during the c&d-phase happen at a lower altitude, leading to a lower impact energy.
   c) The impact area following a failure in the c&d-phase is only half the size compared to an impact in the cruise phase, as the impact angle is 90 degrees, instead of the approximately 78 degrees impact angle during the cruise-phase.

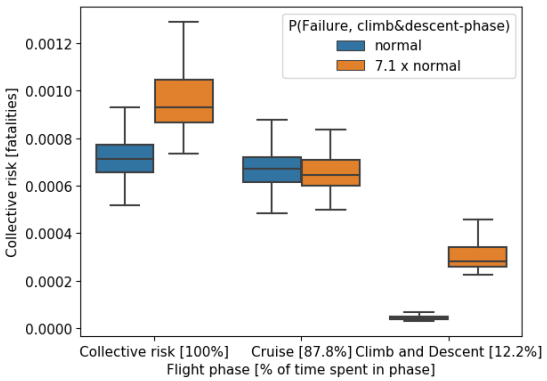

(a) Collective risk (total no. of fatalities)

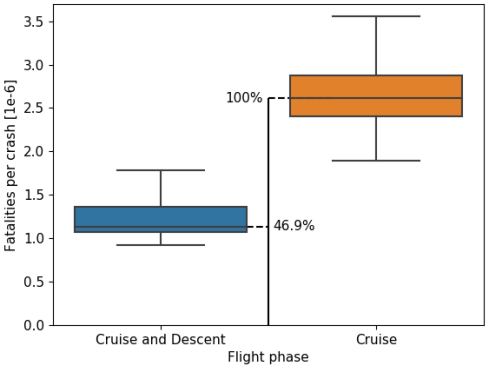

(b) Number of fatalities following an average crash

Figure 12: Results of changing the failure rate in the c&d phase.

Furthermore, a local sensitivity analysis has been performed, assessing the impact on Fat-1M caused by a 1% change in failure rate in either the cruise- or c&d-phase. Inspecting these values, which are presented in Table 6, shows that the sensitivity to a change in failure rate during the cruise phase is 13 times higher.

It can thus be concluded that even though data from GA suggests higher failure rates in the initial- and final phase of a flight, this failure rate is not as important as the failure rate in the cruise-phase, as the majority of the flight is spent in the cruise-phase, during which crashes impact a larger area, and a more densely populated area. Lastly, these changes had no influence on the efficiency of the operations.

	Cruise-phase	Climb-&descent-phase
Time spent in phase (% of total flighttime)	87.8%	12.2%
LS to a 1% change in P(failure) to Fat-1M	0.93%	0.07%

Table 6: Local sensitivity analysis

8.2.2 Varying the velocity at impact (E3)

The kinetic energy upon impact is one of the factors determining the probability of a fatality. Assuming a constant mass, the only variable determining the kinetic energy is the velocity upon impact. In the model presented in this work, this velocity is modeled based on the drone dynamics and the descent profile, an approach that is shared by recent work in UAS research [30]. However, other work researching UAS safety proposes to use the terminal velocity [28, 78, 79, 93], and in the past, EU and UK regulators have indicated to use the terminal velocity + a 40% safety margin for a safety analysis of GA operations [94, 95]. The question is raised what the difference in observed risk would be in each of the three scenarios. In this section, we answer that question by comparing these three scenarios.

As terminal velocity is the velocity where there is no more increase in speed, i.e. $dv/dt = 0$, it can be derived to be as in Equation 12. Using the speed obtained with this relation to do simulations yields the results presented in Figure 13. These results are also presented in Table 7, where the impact speed and confidence intervals are provided, as well. As seen, using terminal velocity leads to an increase in the observed number of fatalities with about 15%, while using terminal velocity + 40% leads to an increase of about 26%. Again, p-values are lower than 1e-10, raising no doubt that a significant difference is found.

$$V_{imp} = \sqrt{\frac{2mg}{\rho_\alpha A C_d}} \tag{12}$$

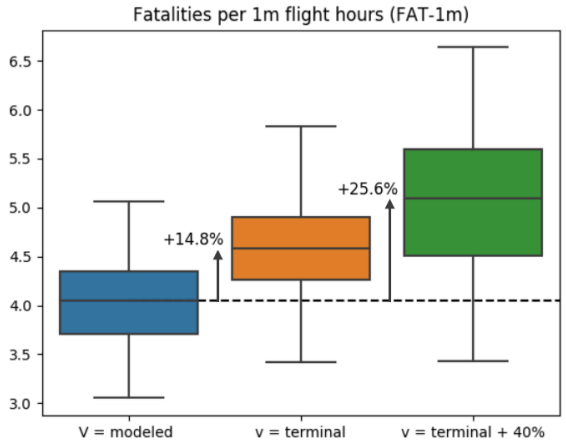

Figure 13: Boxplot for different impact speed scenarios.

	modeled	terminal	terminal + 40%
Average impact speed [m/s]	44.3	71.8	100.5
Average impact energy [kJ]	8.8	23.1	45.4
Fat-1M	4.0	5.1	10.1
% w.r.t modeled [95% CI-interval]	-	14.8% [12.2%, 17.3%]	25.6% [22.7%, 28.6%]
% w.r.t. terminal [95% CI-interval]	-14.8% [-12.2%, -17.3%]	-	9.5% [6.8%, 12.1%]

Table 7: Results of different impact speed scenarios.

It is also worth noting that using terminal velocity + 40% results in 5.5x higher impact energies than in the modelled scenario. Why, then, does a 5.5x higher impact energy only yield an approximately 26% increase in fatalities? This can be explained by means of Figure 14. The fatality model is based on a logistic function, which shows asymptotic behaviour. In this asymptotic region, increasing the impact energy has little influence on the probability of fatality. For example, when comparing impact energies of 8.8kJ and 44.1kJ (as drawn from the scenarios with v=modeled and v=terminal + 40%), the probability of fatality only increases by 3%

at a sheltering factor of 2.5, and 65% when the sheltering factor is 7.5. By means of experiments, this fatality model is scrutinised in more detail in the following section.

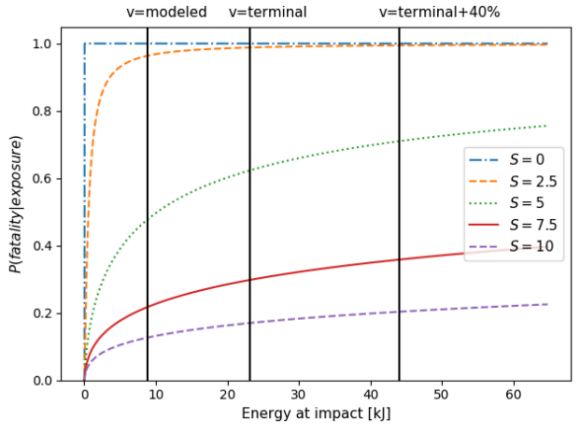

Figure 14: Probability of fatality for different sheltering factors (S), given impact energy.

8.2.3 The Fatality Model (E4)

As explained in the model description, the model used to map an impact to a probability of fatality is based mostly on missile-impact data from the 1960s. In this section, the sensitivity of the components of this fatality model to the fatalities per 1m flight-hours is investigated. The four parameters encompassing this model are the sheltering factor, the kinetic energy upon impact, and α and β. A more detailed description of these parameters, as well as of the equations, can be found in section 5.3.3. This model is evaluated both through a local- and a global sensitivity analysis, of which the results are in Table 8.

The LSA was conducted for a drone weighing 9kg and a drone weighing 1kg. A drone weighing 9kg is according to the case studied in this research, and a drone weighing 1kg is included, as the behaviour of the fatality equation seems to be substantially different at low impact energies (as can be seen in Figure 14).

The results of the LSA are presented in Figure 15 for a drone of 9kg, and in Figure 16 for a drone of 1kg. In the first scenario, a change in the sheltering factor yields the highest sensitivity. The α and impact energy have sensitivities of similar magnitude but different direction, while changing the β has virtually no impact. When considering a drone of 1kg, the major difference is in the sensitivity to impact energy, which is much higher. This result is readily understood by inspecting Figure 14, where it is obvious that the probability of fatality has a higher gradient for lower impact energies.

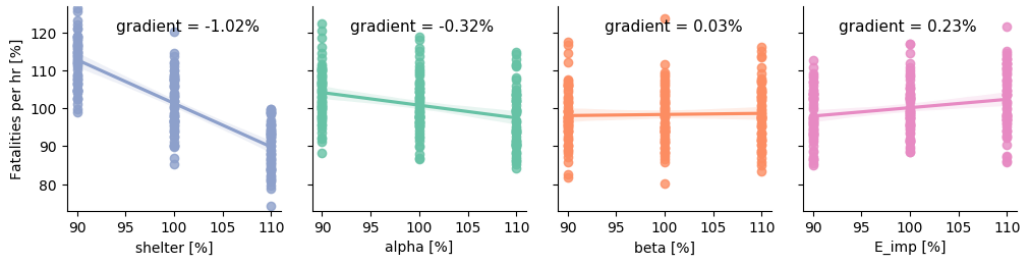

Figure 15: Local sensitivity analysis of fatality equation (UAS mass = 9kg).

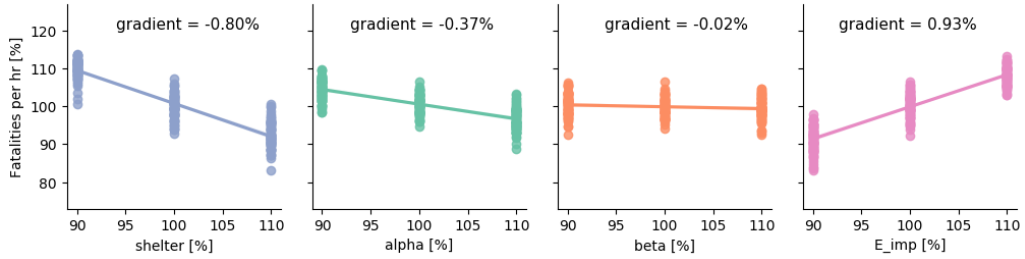

Figure 16: Local sensitivity analysis of fatality equation (UAS mass = 1kg).

To understand the behaviour of the formula over the entire spectrum and account for any interrelations, a GSA is conducted, as well. In this analysis, an important consideration is the range in which the parameters are varied. It is chosen to vary α, β and E_{imp} between 50% and 150% of their original value. Even though the variables used in our model are not well validated, it is unreasonable to think that the error in the current values is more than 50%. The E_{imp}, however, is varied between 10% and 400% of its original value. On the one hand, drones could be much lighter, for example, for surveillance applications, driving the impact energy down. On the other hand, in section 8.2.2, we presented literature where the impact energy is based on the terminal velocity + 40%. In such a scenario, E_{imp} is close to 4x higher than the value obtained through simulations. Therefore, the entire space between [0.1, 4.0]x the value original E_{imp}, was considered. As explained in section 7.3, Latin Hypercube Sampling was used to sample variables in the sample space. For this experiment, 300 data points have been sampled.

Analysing the results of the GSA, we observe significant correlations for the E_{imp} and the Sheltering factor, which is in accordance to the findings from the LSA, with the sheltering factor being the most sensitive according to all analyses. This makes the use of this equation to calculate the probability of fatality especially problematic, as sheltering factors are selected rather arbitrarily. Even though all literature uses a sheltering factor of 7.5 for high buildings, there is no consensus on the definition of a high building, and no experiments have been carried out that validate whether a value of 7.5 is warranted. It should be noted that this analysis assumes that the current shape of the equations is correct, something that is also to be studied.

	Local SA (effect of a 1% change on Fat-1M)		Global SA		
	% gradient (mass = 9kg)	gradient (mass = 1kg)	R	R [95% CI]	p-value
α	-0.18%	-0.70% *	-0.13	[-0.24, -0.02]	0.025
β	0.00%	0.2%	0.02	[-0.10, 0.13]	0.77
E_{imp}	0.10%	0.9% *	0.63 *	[0.56, 0.70]	<1e-10
Sheltering	-0.65% *	-1.1% *	-0.74 *	[-0.79, -0.69]	<1e-10

Table 8: LSA and GSA of four elements of the fatality computation (* = statistically significant correlation)

8.2.4 Wind and stochastic drag (E5)

After failure, the UAS is assumed to be in a ballistic descent until impact. During this descent trajectory, it is only subject to gravitational- and aerodynamic forces. As explained in the model description, these aerodynamic forces are stochastic in nature, because the wind data is sampled from historical data, and the drag coefficient follows a normal distribution. In this section, we present the results from an experiment in which we study the influence of the stochasticity of these aerodynamic forces. Two scenarios are considered. In the first scenario, the wind- and drag-coefficient are determined according to the model description. In the second scenario, the wind-speed is assumed to be zero, and the drag-coefficient is assumed to be constant.

Examining 108 simulations per scenario, it is found that the first scenario (with wind & drag) has 0.32% higher Fat-1M than scenario two with a 95% CI of [-2.9%, 3.5%]. The null-hypothesis cannot be rejected, but it can also be concluded with high certainty that even if there would be a difference between the scenarios, that this difference does not exceed a few percentage points. The same conclusion can be drawn for the flown distance, where the 95% CI is [-0.26%, 0.21%].

The question is raised why the presence of wind does not change the risk. After all, it is reasonable to think that wind causes the crash locations to be more dispersed, making it harder to plan around risky areas. However, comparison of the descent-profiles of scenario 1 and 2 shows that this effect is not large enough to have any significant impact. If the crash locations in scenario 1 are compared to the crash location in scenario 2 for the same failure location, it is seen that in approximately 80% of the times, the crash location is within 12.5 meters of the no-wind crash location, and it is virtually always within 25 meters. The introduction of wind and stochastic drag thus leads to a crash on the same gridcell, as where the UAS would have crashed without wind and stochastic drag. And a crash on the same gridcell yields the same number of fatalities, provided that the kinetic energy is the same. In Figure 17, a heatmap and a 3D-plot is shown of the descent trajectories and -locations with and without wind and drag. In this plot, purple indicates the descent trajectory in the absence of wind and drag. With the red square being a gridcell with 25 meters edges, here, too, it is apparent that crashes are often on the same gridcell in both scenarios.

One might propose to reduce the scaling factor of the environment, creating gridcells with edges of 5 meters instead of 25 meters. Even though this experiment was not conducted due to computational constraints (this would result in a grid of about 1 million cells), it is hypothesised that the result would not be very different. This is because the raw data itself, both regarding sheltering and population density, is not defined with a resolution smaller than 25 meters.

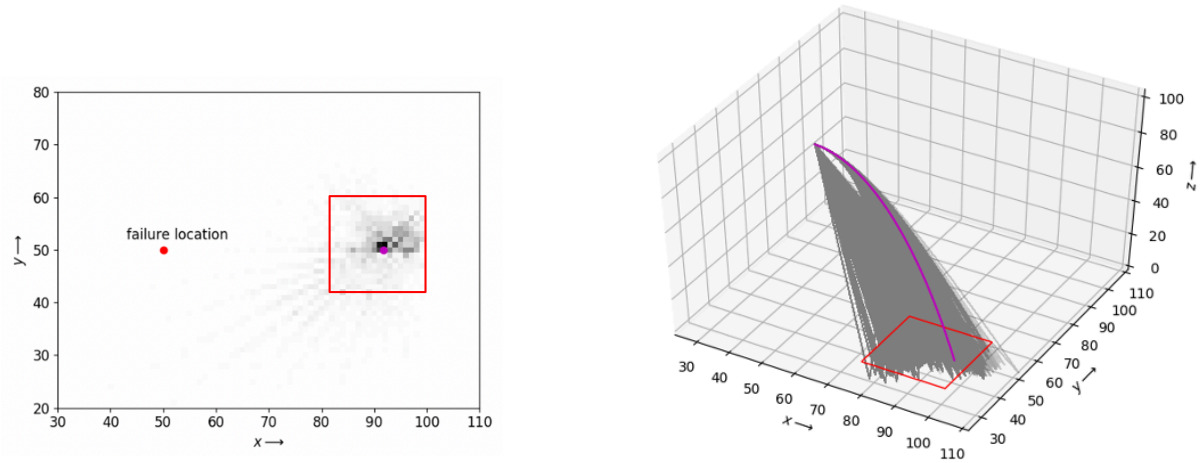

(a) Heatmap of crash locations

(b) 3D-plot of descent trajectories

Figure 17: Crash locations following a failure at (x=50, y=50, z=100), (vx=10, vy=0, vz=0). Purple = descent trajectory without wind. Red = 25x25meters grid-cell.

8.3 Other experiments (E6 and E7)

Two experiments that have been conducted are presented in detail in the appendix [46]. This section briefly states the main results.

The first experiment compared the effect of a cruise altitude of 100 meters with a cruise altitude of 500 meters on the risk and efficiency. It was concluded that a cruise altitude of 500 meters leads to 14% more fatalities per flight-hour and as much as 47% more fatalities per delivery. This is explained by the fact that the deliveries also take longer as more time spent in the climb- and descent-phase.

The second experiment analyses the possibility to optimise for efficiency rather than risk. This experiment is conducted through varying the *add_value*, which stimulates the path-finders to take the shortest path, rather than the path with the least risk. For the city of New York, an *add_value* of 1e-9 results into 1.8 Fat-1M and takes about 45 minutes per task, while an *add_value* of 1 leads to about 13 Fat-1M, but takes only 20 minutes per task. The model is eventually calibrated for an *add_value* of 0.001, yielding almost the same risk as a value of 1e-9, but taking only 30 minutes per task. It can be concluded that, indeed, it is possible to lower the risk by a significant amount by optimising for risk rather than path-length. Decision-makers could use this data to discuss what amount of safety is deemed acceptable, at what cost of efficiency.

9 Discussion

The previous section presented the results of our experiments. These experiments examined the influence of the environment on UAS safety and analysed several methods for calculating TPR. This section discusses the limitations and broader context of these results. The discussion is structured in four topics, being the case description, the environment, the failure modes and the risk computation.

9.1 Limitations of Case Description

This section discusses three limitations regarding the case description and explains how they have influenced the results.

1. Structuring of the Airspace: It was assumed that every part of the urban airspace was accessible for UASs. However, this is likely to be different in reality. One can expect the urban airspace to be scattered with no-fly zones. These no-fly zones would, for example, separate UAS traffic from government buildings, airports and accident scenes. Also, some (e.g. [22] and [23]) propose to deliberately structure the urban airspace to reduce the complexity of urban UAS operations. Such structures reduce the solution space for flight-paths. In a smaller solution space, optimal routes may be no longer available. However, airspace structures might also have advantages. In particular, Hoekstra et al. showed that this could reduce the number of CD&R actions, thereby reducing complexity and thus increasing safety [96]. Furthermore, airspace structures could steer UAS away from dangerous areas. This benefits safety, especially for UASs which do not have either the objective or

the capabilities to minimise the risk of their operations, themselves.

2. Interactions between UASs and other traffic: This research assumed that UASs are invisible to each other, and thus ignored conflicts. In reality, such conflicts would require drones to conduct collision-avoidance manoeuvres, which, again, may make it impossible to fly the optimal flight plan. These manoeuvres may thus increase the risk of the operations.

3. Model of vehicle dynamics: The vehicles are modelled very simplistically in this research. For example, instantaneous acceleration is assumed. Without this assumption, manoeuvres would be slower, thus increasing the risk. Even though the most realistic solution would be to implement a full 6-degrees-of-freedom model of the drone, intermediate possibilities exist, such as smoothening the paths using Dubins curves [97].

9.2 Environment

Despite extensive efforts to model an environment that is as detailed as possible, it is still far from realistic, and several limitations are identified.

1. Resolution of impact data: Population density data is based on Census data. Even though this differs slightly per city, this is defined per neighbourhood. Also, it is not defined exactly, but rather in categories such as 25.000 - 50.000 people / square kilometre. The sheltering data is defined with a higher resolution, as OpenStreetMap data makes it possible to distinguish between different buildings. However, it still uses discrete categories, such as "homes" or "offices", and does not discriminate between a one-story home or a ten-story home. However, the risk of an impact in either of those would be very different.

2. Data based on where people live: The population density data does not include the movement of people throughout the day. Even though bottom-up estimations of the movements of people through cities are incorporated, this is not perfect. Also, we modelled cities as close environments, where people cannot enter or leave. Again, this limitation does not invalidate conclusions drawn from comparisons between different scenarios. However, it does make it impossible to derive exact values of the risk in a city.

3. Cities are modelled as static environments: This means that an optimal flight plan computed at $t = 0$, remains optimal throughout the execution of this flight plan. In reality, people moving through the city would likely cause this solution to be suboptimal already before the UAS started its engines. Therefore, a dynamic environment would make it harder to minimise the risk.

The identified limitations make it impossible to realistically obtain an absolute figure of the risk. However, when comparing, for example, different methods of computing the impact speed, the current level of environmental detail does not influence the conclusions much. However, for experiments in which one is interested in comparing the risk of impact at two locations close to each other, a higher level of detail is needed to draw robust conclusions. Consider, for example, the experiment considering the impact of wind on the descent trajectory. We concluded that wind caused the impact location to move 10 meters, on average, however, due to the low resolution of the data, this did not significantly influence the overall risk.

9.3 Failure Modes

This work considered only one failure mode, namely an all-engines-out failure, resulting in a ballistic descent. This is a limiting assumption, and this section discusses the implications following from this assumption.

1. Other modes of failure: We identify three other modes of failure that this research did not include. Firstly, although wind is incorporated in the descent-trajectory, wind, and adverse weather conditions in general, could also be a reason for failure. Secondly, the presence of other vehicles could also increase the probability of failure, as drones might crash into each other. Such a scenario is studied by Kim et al. [33], who not only concludes that the probability of failure would increase, but also concludes that energy dissipates in the crash, thereby reducing the energy upon impact on the ground. Thirdly, the UAS may fail in many more ways than only an all-engines-out scenario. For example, it may lose a few engines, it may catch fire, or it may lose its payload. Several methods to analyse the implications of these failure modes include boolean trees [98], Bayesian Belief Networks [59], or Barrier Bowtie Models [29].

2. Secondary impact following a crash: Our model only quantifies fatalities resulting from the initial crash of the UAS. However, secondary effects may occur. For example, a building may catch fire, or debris may fly

around, causing a second impact. These effects can increase the risk following a crash.

3. Risk mitigation: A UAS can still mitigate risk after it has failed. For example, Wyllie et al. proposed to use a parachute-recovery system [66]. Also, some failure modes may cause the drone to lose only partial control, making it possible to crash on a safe location. Another mitigation-factor lies in human behaviour. It is feasible that humans will identify a crashing UAS and take actions to avoid an impact.

9.4 Risk Computation

Lastly, we discuss two limitations following our risk computation.

1. Estimated probability of failure: The previous section already indicated limitations of the probability of failure because many failure modes are ignored. It is also limited because even for the scenario considered in this research, no exact values are available. This makes it impossible to derive an accurate absolute value for the risk. In the absence of real data, a sensitivity analysis can be used to analyse the model, as was done in this research. However, future work should include better models of the probability of failure. Inspiration can be drawn from the space industry, which is experienced in estimating probabilities failure in the absence of data on the reliability of systems [55, 56].

2. Probability of fatality calculation: We conducted both a local- and global-sensitivity analysis on the model of the probability of fatality. Even though both analyses were performed thoroughly and converged to the same results, different forms of equations should also be studied by means of a structural analysis.

10 Conclusions and Future Research

In this research, we conducted an agent-based safety analysis of the third party risk of UAS operations in the urban environment. More specifically, we modelled a drone package delivery concept of operations in Delft, New York and Paris. These cities were modeled by combining data from four sources to arrive at an environment that includes the movement of people throughout the day, and their sheltering because of vehicles, buildings and trees. The objective was to identify the most essential elements of the risk computation, and to study the impact of individual environmental parameters on safety. This research yields four important conclusions.

The first conclusion is that the dispersion of low-risk areas in environments has a mitigating effect on the system-wide risk. We found that if these low-risk areas are more dispersed, UASs can incorporate them more easily into their flight plans, and in such a way decrease the total, system-wide risk. Secondly, the results show that understanding the failure rate in the cruise phase is much more important than understanding it in the climb-&descent-phase, as the UAS spends more time in the cruise phase, and the risk per unit of time is also higher in the cruise-phase. Thirdly, modelling the impact speed using the terminal velocity or terminal velocity + 40%, approaches which are both commonly used in research, leads to 15 - 26% higher values for risk. Fourthly, both a local- and global sensitivity analysis suggests that the sheltering factor is the most sensitive component of the fatality computation. Even though all components of this computation are unvalidated, this finding implies that it is most important to focus on this parameter. It is also easier to study the sheltering factor than to study human impact models, as it is easier to test impacts on buildings than on humans. Lastly, wind has no meaningful impact on the resulting safety, as the 95% CI on the difference is [-2.9%, 3.5%]. We provide two explanations for this small difference. First, the wind has little influence on the drone's path as the descent-time is only about five seconds. Second, we have modelled grid-cells with dimensions of 25x25meters, and the effect of stochastic wind and drag is often too small to cause a crash in a different grid-cell than where the drone would have crashed without the influence wind.

This research owes its importance to the regulators, who demand a thorough assessment of the safety of urban UAS operations before they would certify these operations. Therefore, future research should be aimed at answering the regulators' question. Based on our findings, we have established 14 directions of future research, which are discussed in the appendix [46]. Here, we present the three most important directions. First, a better understanding should be obtained of the different failure modes, their consequences, and most importantly, their associated probabilities of failure. This research should be focused on failures in the cruise phase, because these have the most impact on risk. Secondly, more accurate representations of urban environments should be developed. Most important is achieving a better resolution and obtaining data on the whereabouts of people during the day. Thirdly, the model relating the probability of fatality to impact should be validated. We found that it is most important to focus on the sheltering factor, which is the most sensitive parameter.

References

[1] Uber, "Uber Copter infopage," 2020. [Online]. Available: https://www.uber.com/nl/nl/ride/uber-copter/

[2] L. Hausmann, J. Murnane, M. Neelesh, and F. Neuhaus, "The future of parcel delivery: Drones and disruption The Next Normal," *McKinsey & Company*, 2019. [Online]. Available: https://www.mckinsey.com/~/media/McKinsey/FeaturedInsights/TheNextNormal/The-Next-Normal-Thefuture-of-parcel-delivery-vF

[3] M. Oonk, "Using Cargo Drones in Last Mile Delivery," *Deloitte*, 2018. [Online]. Available: https://www2.deloitte.com/nl/nl/pages/consumer-industrial-products/articles/using-cargo-drones-in-last-mile-delivery.html

[4] J. Elias, "Alphabets Wing launches first commercial drone delivery," 2019. [Online]. Available: https://www.cnbc.com/2019/10/18/alphabets-wing-launches-first-commercial-drone-delivery.html

[5] J. Porter, "Alphabets Wing drones get FAA approval to make deliveries in the US," 2020. [Online]. Available: https://www.theverge.com/2019/4/23/18512658/google-alphabet-wing-drone-delivery-service-faa-approval-commercial-deliveries

[6] A. Palmer, "Amazon wins FAA approval for Prime Air drone delivery fleet," 2020. [Online]. Available: https://www.cnbc.com/2020/08/31/amazon-prime-now-drone-delivery-fleet-gets-faa-approval.html

[7] Federal Aviation Administration (FAA), "Package Delivery by Drone (Part 135)," 2020. [Online]. Available: https://www.faa.gov/uas/advanced_operations/package_delivery_drone/

[8] D. Locascio, M. Levy, K. Ravikumar, B. German, S. I. Briceno, and D. N. Mavris, "Evaluation of Concepts of Operations for sUAS Package Delivery," *16th AIAA Aviation Technology, Integration, and Operations Conference, 2016*, no. June, pp. 1–16, 2016.

[9] M. Doole, J. Ellerbroek, and J. Hoekstra, "Drone delivery: Urban airspace traffic density estimation," *SESAR Innovation Days*, no. December, 2018.

[10] EASA (European Aviation Safety Agency), "European Aviation Safety Agency Introduction of a regulatory framework for the operation of unmanned aircraft," pp. 1–50, 2015.

[11] J. L. Rios, I. S. Smith, D. R. Smith, S. Jurcak, and R. Strauss, "UTM UAS Service Supplier Development," *NASA Technical Memorandum*, no. December, pp. 1–21, 2018.

[12] S. Kaplan, "The words of risk analysis," *Risk Analysis*, vol. 17, no. 4, pp. 407–417, 1997.

[13] H. A. Blom, S. H. Stroeve, and T. Bosse, "Modelling of potential hazards in agent-based safety risk analysis," *Proceedings of the 10th USA/Europe Air Traffic Management Research and Development Seminar, ATM 2013*, 2013.

[14] J. Epstein, *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press, 2007.

[15] E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. SUPPL. 3, pp. 7280–7287, 2002.

[16] R. L. Rothfeld, M. Balac, K. O. Ploetner, and C. Antoniou, "Agent-based simulation of urban air mobility," *2018 Modeling and Simulation Technologies Conference*, pp. 1–10, 2018.

[17] S. H. Stroeve, H. A. Blom, and G. J. Bakker, "Contrasting safety assessments of a runway incursion scenario: Event sequence analysis versus multi-agent dynamic risk modelling," *Reliability Engineering and System Safety*, vol. 109, pp. 133–149, 2013. [Online]. Available: http://dx.doi.org/10.1016/j.ress.2012.07.002

[18] T. Prevot, J. Rios, P. Kopardekar, J. E. Robinson III, M. Johnson, and J. Jung, "UAS Traffic Management (UTM) Concept of Operations to Safely Enable Low Altitude Flight Operations," no. June, pp. 1–16, 2016.

[19] P. Kopardekar and S. Bradford, "UTM Research Transition Team Plan," NASA Ames Research Center, Tech. Rep., 2017.

[20] NASA, "NASA UTM in 2019: Highlights," NASA, Tech. Rep., 2019. [Online]. Available: https://utm.arc.nasa.gov/utm2019.shtml

[21] Federal Aviation Administration (FAA), "Unmanned Aircraft System Traffic Management Concept of Operations," FAA, Tech. Rep., 2020.

[22] E. Sunil, J. Hoekstra, J. Ellerbroek, F. Bussink, D. Nieuwenhuisen, A. Vidosavljevic, and S. Kern, "Metropolis: Relating airspace structure and capacity for extreme traffic densities," *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar, ATM 2015*, no. June, 2015.

[23] M. F. Bin Mohammed Salleh, D. Y. Tan, C. H. Koh, and K. H. Low, "Preliminary concept of operations (ConOps) for traffic management of unmanned aircraft systems (TM-UAS) in urban environment," *AIAA Information Systems-AIAA Infotech at Aerospace, 2017*, no. January, pp. 1–13, 2017.

[24] M. F. Bin Mohammed Salleh, W. Chi, Z. Wang, S. Huang, D. Y. Tan, T. Huang, and K. H. Low, "Preliminary concept of adaptive urban airspace management for unmanned aircraft operations," *AIAA Information Systems-AIAA Infotech at Aerospace, 2018*, no. 209989, pp. 1–12, 2018.

[25] N. Pongsakornsathien, S. Bijjahalli, A. Gardi, A. Symons, Y. Xi, R. Sabatini, and T. Kistan, "A performance-based airspace model for unmanned aircraft systems traffic management," *Aerospace*, vol. 7, no. 11, pp. 1–25, 2020.

[26] P. Wu and R. A. Clothier, "The development of ground impact models for the analysis of the risks associated with unmanned aircraft operations over inhabited areas," *11th International Probabilistic Safety Assessment and Management Conference and the Annual European Safety and Reliability Conference 2012, PSAM11 ESREL 2012*, vol. 7, pp. 5222–5234, 2012.

[27] R. A. Clothier, B. P. Williams, and K. J. Hayhurst, "Modelling the risks remotely piloted aircraft pose to people on the ground," *Safety Science*, vol. 101, pp. 33–47, 2018.

[28] A. Washington, R. A. Clothier, and J. Silva, "A review of unmanned aircraft system ground risk models," *Progress in Aerospace Sciences*, vol. 95, no. August, pp. 24–44, 2017. [Online]. Available: https://doi.org/10.1016/j.paerosci.2017.10.001

[29] R. Clothier, B. Williams, and A. Washington, "Development of a Template Safety Case for Unmanned Aircraft Operations over Populous Areas," *SAE Technical Papers*, vol. 2015-Septe, no. September, 2015.

[30] S. Primatesta, A. Rizzo, and A. la Cour-Harbo, "Ground Risk Map for Unmanned Aircraft in Urban Environments," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 97, no. 3-4, pp. 489–509, 2020.

[31] S. Bertrand, N. Raballand, F. Viguier, and F. Muller, "Ground risk assessment for long-range inspection missions of railways by UAVs," *2017 International Conference on Unmanned Aircraft Systems, ICUAS 2017*, pp. 1343–1351, 2017.

[32] S. Bertrand, N. Raballand, and F. Viguier, "Evaluating Ground Risk for Road Networks Induced by UAV Operations," *2018 International Conference on Unmanned Aircraft Systems, ICUAS 2018*, pp. 168–176, 2018.

[33] S. H. Kim, "Third-party risk analysis of small unmanned aircraft systems operations," *Journal of Aerospace Information Systems*, vol. 17, no. 1, pp. 24–35, 2020.

[34] B. Rattanagraikanakorn, A. Sharpanskykh, M. Schuurman, D. Gransden, H. A. Blom, and C. De Wagter, "Characterizing UAS collision consequences in future UTM," *2018 Aviation Technology, Integration, and Operations Conference*, 2018.

[35] Joint Authorities for Rulemaking of Unmanned Systems (JARUS), "JARUS guidelines on Specific Operations Risk Assessment (SORA), EDITION 2.0," 2019. [Online]. Available: http://jarus-uas.org

[36] A. la Cour-Harbo, "Quantifying Risk of Ground Impact Fatalities for Small Unmanned Aircraft," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 93, no. 1-2, pp. 367–384, 2019.

[37] E. Rudnick-Cohen, S. Azarm, and J. W. Herrmann, "Planning unmanned aerial system (UAS) takeoff trajectories to minimize third-party risk," *2019 International Conference on Unmanned Aircraft Systems, ICUAS 2019*, pp. 1306–1315, 2019.

[38] S. Primatesta, G. Guglieri, and A. Rizzo, "A Risk-Aware Path Planning Strategy for UAVs in Urban Environments," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 95, no. 2, pp. 629–643, 2019. [Online]. Available: http://dx.doi.org/10.1007/s10846-018-0924-3

[39] H. A. P. Blom, "AE4448 Agent Based Safety Risk Analysis (Lecture 4)," 2013.

[40] G. Weiss, *Multiagent Systems*, 2nd ed. MIT Press, 2013. [Online]. Available: https://mitpress-mit-edu.tudelft.idm.oclc.org/books/multiagent-systems-second-edition

[41] N. Metropolis and S. Ulam, "The Monte Carlo Method," *Journal of the American Statistical Association*, vol. 44, no. 247, pp. 335–341, 9 1949. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/01621459.1949.10483310

[42] DJI, "Matrice 300 RTK." [Online]. Available: https://www.dji.com/nl/matrice-300

[43] B. Zwanenburg, "UASriskSim Github repository." [Online]. Available: https://github.com/bastiaanzwanenburg/UASriskSim

[44] D. Masad and J. Kazil, "Mesa: An Agent-Based Modeling Framework," *Proceedings of the 14th Python in Science Conference*, no. April, pp. 51–58, 2015.

[45] P. E. Hart, N. J. Nilsson, and B. Raphael, "Formal Basis for the Heuristic Determination eijj ,," *Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[46] B. Zwanenburg, "An Agent-Based Safety Analysis of UAS operations in a high-fidelity urban environment," Delft University of Technology, Tech. Rep., 2021.

[47] G. Guglieri and G. Ristorto, "Safety Assessment for Light Remotely Piloted Aircraft Systems," *2016 INAIR - International Conference on Air Transport*, vol. 1, no. December, pp. 1–7, 2016. [Online]. Available: http://www.mavtech.eu/site/assets/files/1261/safety_assessment_for_light_remotely_piloted_aircraft_systems_inair_template_2016_final.pdf

[48] Centraal Bureau voor Statistiek, "Kadastrale Kaart Delft," 2018. [Online]. Available: https://kadastralekaart.com/gemeenten/delft-GM0503

[49] Atelier Parisien d'Urbanisme (APUR), "Population Density Paris," 2019. [Online]. Available: https://www.apur.org/en/geo-data/population-density

[50] US Census Bureau, "New York City population density," 2019. [Online]. Available: https://www.arcgis.com/home/webmap/viewer.html?url=https%3A%2F%2Fportal.sgm.gob.mx%2Farcgis%2Frest%2Fservices%2FDatosAbiertos%2FDatosAbiertos%2FMapServer&source=sd

[51] "OpenStreetMap." [Online]. Available: https://www.openstreetmap.org/

[52] Google, "Google Streetview." [Online]. Available: https://www.google.com/streetview/

[53] R. Melnyk, D. Schrage, V. Volovoi, and H. Jimenez, "A third-party casualty risk model for unmanned aircraft system operations," *Reliability Engineering and System Safety*, vol. 124, pp. 105–116, 2014. [Online]. Available: http://dx.doi.org/10.1016/j.ress.2013.11.016

[54] U. C. Kaya, A. Dogan, and M. Huber, "A probabilistic risk assessment framework for the path planning of safe task-aware uas operations," *AIAA Scitech 2019 Forum*, no. January, pp. 1–17, 2019.

[55] S. Guarro, B. Bream, L. K. Rudolph, and R. J. Mulvihill, "The Cassini mission risk assessment framework and application techniques," *Reliability Engineering and System Safety*, vol. 49, no. 3, pp. 293–302, 1995.

[56] S. D. Guikema and M. E. Paté-Cornell, "Bayesian Analysis of Launch Vehicle Success Rates," *Journal of Spacecraft and Rockets*, vol. 41, no. 1, pp. 93–102, 2004.

[57] D. A. Burke, C. E. Hall, and S. P. Cook, "System-level airworthiness tool," *Journal of Aircraft*, vol. 48, no. 3, pp. 777–785, 2011.

[58] A. Washington, R. A. Clothier, and B. P. Williams, "A Bayesian approach to system safety assessment and compliance assessment for Unmanned Aircraft Systems," *Journal of Air Transport Management*, vol. 62, pp. 18–33, 2017. [Online]. Available: http://dx.doi.org/10.1016/j.jairtraman.2017.02.003

[59] A. Washington, R. Clothier, N. Neogi, J. Silva, K. Hayhurst, and B. Williams, "Adoption of a Bayesian Belief Network for the System Safety Assessment of Remotely Piloted Aircraft Systems," *Safety Science*, vol. 118, no. March, pp. 654–673, 2019. [Online]. Available: https://doi.org/10.1016/j.ssci.2019.04.040

[60] M. V. Bendarkar, A. Behere, S. I. Briceno, and D. N. Mavris, "A Bayesian Safety Assessment Methodology for Novel Aircraft Architectures and Technologies Using Continuous FHA," no. June, 2019.

[61] F. Jackman, "Nearly Half of Commercial Jet Accidents Occur During Final Approach, Landing," 2014. [Online]. Available: https://flightsafety.org/asw-article/nearly-half-of-commercial-jet-accidents-occur-during-final-approach-landing/

[62] R. Clothier, R. Walker, N. Fulton, and D. Campbell, "A casualty risk analysis for unmanned aerial system (UAS) operations over inhabited areas," *Second Australasian Unmanned Air Vehicle Conference*, pp. 1–15, 2007.

[63] A. T. Ford and K. J. Mcentee, "Assessment of the risk to ground population due to an unmanned aircraft in-flight failure," *10th AIAA Aviation Technology, Integration and Operations Conference 2010, ATIO 2010*, vol. 1, no. September, pp. 1–13, 2010.

[64] J. D. Stevenson, S. OYoung, and L. Rolland, "Estimated levels of safety for small unmanned aerial vehicles and risk mitigation strategies," *Journal of Unmanned Vehicle Systems*, vol. 3, no. 4, pp. 205–221, 2015.

[65] H. Usach, J. A. Vila, and . Gallego, "Trajectory-Based, Probabilistic Risk Model for UAS Operations," *Risk Assessment in Air Traffic Management*, 2020.

[66] T. Wyllie, "Parachute recovery for UAV systems," *Aircraft Engineering and Aerospace Technology*, vol. 73, no. 6, pp. 542–551, 2001.

[67] A. Panta, S. Watkins, and R. Clothier, "Dynamics of a small unmanned aircraft parachute system," *Journal of Aerospace Technology and Management*, vol. 10, pp. 1–14, 2018.

[68] L. C. Barr, R. L. Newman, E. Ancel, C. M. Belcastro, J. V. Foster, J. K. Evans, and D. H. Klyde, "Preliminary risk assessment for small unmanned aircraft systems," *17th AIAA Aviation Technology, Integration, and Operations Conference, 2017*, no. June, 2017.

[69] A. la Cour-Harbo, "Ground impact probability distribution for small unmanned aircraft in ballistic descent," *2020 International Conference on Unmanned Aircraft Systems, ICUAS 2020*, no. February, pp. 1442–1451, 2020.

[70] Ministerie van Infrastructuur en Waterstaat, "Dataplatform KNMI." [Online]. Available: https://dataplatform.knmi.nl/catalog/index.html

[71] F. Bauelos-Ruedas, C. Angeles-Camacho, and Sebastin, "Methodologies Used in the Extrapolation of Wind Speed Data at Different Heights and Its Impact in the Wind Energy Resource Assessment in a Region," in *Wind Farm - Technical Regulations, Potential Estimation and Siting Assessment*. InTech, 6 2011, no. tourism, p. 13. [Online]. Available: https://www.intechopen.com/books/advanced-biometric-technologies/liveness-detection-in-biometricshttp://www.intechopen.com/books/wind-farm-technical-regulations-potential-estimation-and-siting-assessment/methodologies-used-in-the-extrapolation-of-wind-sp

[72] D. Feinstein, W. Heugel, and M. Kardatzke, *Personnel Casualty Study*. Defense Technical Information Center, 1968.

[73] V. R. Clare, J. H. Lewis, A. P. Mickiewicz, and L. M. Sturdivan, "Blunt Trauma Data Correlation," no. May, p. 54, 1975. [Online]. Available: https://apps.dtic.mil/docs/citations/ADA012761

[74] D. Raymond, C. Van Ee, G. Crawford, and C. Bir, "Tolerance of the skull to blunt ballistic temporoparietal impact," *Journal of Biomechanics*, vol. 42, no. 15, pp. 2479–2485, 2009. [Online]. Available: http://dx.doi.org/10.1016/j.jbiomech.2009.07.018

[75] A. Radi, "Human Injury Model for Small Unmanned Aircraft Impacts," pp. 1–31, 2013. [Online]. Available: https://www.casa.gov.au/sites/g/files/net351/f/_assets/main/airworth/papers/human-injury-model-small-unmanned-aircraft-impacts.pdf

[76] L. GREENSPAN, B. A. McLELLAN, and H. GREIG, "Abbreviated Injury Scale and Injury Severity Score," *The Journal of Trauma: Injury, Infection, and Critical Care*, vol. 25, no. 1, pp. 60–64, 1 1985. [Online]. Available: http://journals.lww.com/00005373-198501000-00010

[77] C. Bir and D. C. Viano, "Design and injury assessment criteria for blunt ballistic impacts," *Journal of Trauma - Injury, Infection and Critical Care*, vol. 57, no. 6, pp. 1218–1224, 2004.

[78] K. Dalamagkidis, K. P. Valavanis, and L. A. Piegl, "On unmanned aircraft systems issues, challenges and operational restrictions preventing integration into the National Airspace System," *Progress in Aerospace Sciences*, vol. 44, no. 7-8, pp. 503–519, 2008.

[79] ——, *On Integrating Unmanned Aircraft Systems into the National Airspace System*. Dordrecht: Springer Netherlands, 2012. [Online]. Available: http://link.springer.com/10.1007/978-94-007-2479-2

[80] F. Klügl, "A validation methodology for agent-based simulations," *Proceedings of the ACM Symposium on Applied Computing*, no. January 2008, pp. 39–43, 2008.

[81] H. Zhu, P. A. Hall, and J. H. May, "Software unit test coverage and adequacy," *ACM Computing Surveys*, vol. 29, no. 4, pp. 366–427, 1997.

[82] N. Kurczewski, "Smarter Crash Test Dummies," 1 2011. [Online]. Available: https://www.roadandtrack.com/car-culture/a16755/smarter-crash-test-dummies/

[83] S. S. Shapiro and M. B. Wilk, "An Analysis of Variance Test for Normality (Complete Samples)," *Biometrika*, vol. 52, no. 3/4, p. 591, 1965.

[84] M. Saculinggan and E. A. Balase, "Empirical power comparison of goodness of fit tests for normality in the presence of outliers," *Journal of Physics: Conference Series*, vol. 435, no. 1, 2013.

[85] J. M. Bland and D. G. Altman, "Statistics notes: Multiple significance tests: the Bonferroni method," *BMJ*, vol. 310, no. 6973, pp. 170–170, 1 1995. [Online]. Available: https://www.bmj.com/lookup/doi/10.1136/bmj.310.6973.170

[86] A. Saltelli and P. Annoni, "How to avoid a perfunctory sensitivity analysis," *Environmental Modelling and Software*, vol. 25, no. 12, pp. 1508–1517, 2010. [Online]. Available: http://dx.doi.org/10.1016/j.envsoft.2010.04.012

[87] B. Ratner, "The correlation coefficient: Its values range between 1/1, or do they," *Journal of Targeting, Measurement and Analysis for Marketing*, vol. 17, no. 2, pp. 139–142, 2009.

[88] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola, *Global Sensitivity Analysis. The Primer*, 2008.

[89] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 42, no. 1, pp. 55–61, 1979.

[90] A. Saltelli, S. Tarantola, and F. Campolongo, "Sensitivity analysis as an ingredient of modeling," pp. 377–395, 2000.

[91] D. W. Allen, *GIS Tutorial 2: Spatial Analysis Workbook*. New York: ESRI Press, 2011.

[92] J. K. Ord and A. Getis, "Local Spatial Autocorrelation Statistics: Distributional Issues and an Application," *Geographical Analysis*, vol. 27, no. 4, pp. 286–306, 9 1995. [Online]. Available: http://doi.wiley.com/10.1111/j.1538-4632.1995.tb00912.x

[93] G. Guglieri, A. Lombardi, and G. Ristorto, "Operation Oriented Path Planning Strategies for Rpas," *Amer. J. Sci. Technol.*, vol. 2, no. 6, pp. 321–328, 2015. [Online]. Available: http://www.aascit.org/journal/ajst

[94] EASA, "Policy for unmanned aerial vehicle (UAV) certification," *Management*, no. 16, pp. 1–42, 2005.

[95] D. R. Haddon and C. J. Whittaker, "AIRCRAFT AIRWORTHINESS CERTIFICATION STANDARDS FOR CIVIL UAVs. D.R.Haddon, C.J.Whittaker - Civil Aviation Authority, UK," *Aviation*, vol. 107, no. August, pp. 79–86, 2002. [Online]. Available: http://www.caa.co.uk/docs/1416/srg_acp_00016-01-120203.pdf

[96] J. Hoekstra, S. Kern, O. Schneider, F. Knabe, and B. Lamiscarre, "Metropolis Simulation Results and Analysis Report," Tech. Rep., 2015.

[97] L. De Filippis, G. Guglieri, and F. Quagliotti, "A minimum risk approach for path planning of UAVs," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 61, no. 1-4, pp. 203–219, 2011.

[98] B. Levasseur, S. Bertrand, N. Raballand, F. Viguier, and G. Goussu, "Accurate Ground Impact Footprints and Probabilistic Maps for Risk Analysis of UAV Missions," *IEEE Aerospace Conference Proceedings*, vol. 2019-March, pp. 1–10, 2019.

II

Literature Study
previously graded under AE4020

1

Introduction

Unmanned Aircraft safely roaming urban skies

Vehicles flying through the biggest cities of the world used to be a concept belonging to science-fiction movies, but is quickly becoming reality. Technology enabling these operations is developing quickly, with companies like Amazon, Google and, Uber having conducted fly-tests with vehicles that would eventually be able to transport packages or even passengers through cities [51, 97, 132]. These operations are performed by a *UAS*, which is defined by EASA as [33]:

> An Unmanned Aircraft System (UAS) comprises individual system elements consisting of an "unmanned aircraft", the "control station" and any other system elements necessary to enable flight, i.e. "command and control link" and "launch and recovery elements". There may be multiple control stations, command & control links, and launch and recovery elements within a UAS.

Currently, UAS operations are only allowed within Visual Line of Sight of the operator (VLOS) and are restricted to special use-cases [36]. However, regulators are working hard on changing this. The FAA and NASA have jointly proposed a Concept of Operations that would allow beyond-VLOS operations in dense traffic situations over densely populated areas, and the European Aviation Safety Agency is working on similar proposals. Both organizations have indicated that mitigating the risk following from urban UAS operations is an important consideration while developing new regulations [29, 110].

In urban UAS operations, the people most at risk are third-parties, which are the people who have neither control over the operations or derive any benefit from these operations. Furthermore, they have no choice other than being at risk if unmanned aircraft fly over them through a city. Therefore, it is of utmost importance to analyze and mitigate the risk for these people as much as possible.

Over the past years, a lot of work has been done both inside and outside the academic field on mitigating this risk through safer vehicles. This work includes the engineering of vehicles that are less prone to failures, and that have robust systems that reduce the impact following a failure. Futuristic concepts like parachutes that deploy upon failure and air-bags are designed for this purpose [144]. However, the risk can also be mitigated from an operational standpoint. By flying smarter routes, the UAS could avoid densely populated areas and thus reduce the third-party risk.

To gain a better understanding of the effectiveness of operational measures that reduce third-party risk, and how those measures could be successfully implemented into ConOps for operations by UAS in urban areas, a literature survey is conducted.

Report Structure

Chapters 2 - 5 together are the core of the literature study, which investigates relevant aspects of (the risk following from) urban UAS Operations. In this survey, chapter 2 analyzes the practicalities of UAS Operations, such as the regulations, possible concepts of operations, and how UAS Operations can be modeled. Chapter 3 analyzes three state-of-the-art research projects that have attempted to conduct traffic-management by

employing (airspace) structures. Following this research, three research gaps have been identified. Chapter 4 presents an exhaustive collection of state-of-the-art literature regarding (third-party) risk following (urban) UAS operations. Lastly, chapter 5 provides a brief view of popular path-finding algorithms and discusses how these algorithms can be extended to minimize the risk, rather than the path-length.

Following the core of the literature study, chapter 6 summarizes the most important findings following from it. These findings are then extended in chapter 7 to a research plan for an MSc thesis that follows this literature study. This plan could lead to a contribution of the field of UAS operations and third-party-risk assessments, and hopefully brings the field a small step closer to a wide-scale introduction of Unmanned Aircraft safely roaming the urban skies.

#	Title	Purpose
2	UAS Operations	Present and evaluate regulations and ConOps relevant for future UAS applications
3	Structure of Airspace	Discuss research the effect of (airspace) structures on UTM
4	Risk	Discuss the role of Risk in UAS Operations, and present state-of-the-art third-party risk models
5	Multi-Agent Path Finding	Present common path-finding algorithms, including their extension to risk-based planning
6	Conclusion	Summarize the findings of chapters 2-5
7	Research Plan	To present a research plan for a MSc thesis, following from research gaps and knowledge established in earlier chapters

Table 1.1: Tabular overview of the report, including the purpose of each chapter.

2

UAS Operations

Unmanned Aircraft Systems are increasingly used for commercial operations. For example, in May 2020, Port Authorities completed the first successful package delivery by drone to a ship in the Port of Rotterdam [101]. In the Eastern-African country of Rwanda, a company called *Zipline* has delivered over 33,000 blood samples to medical facilities by drone since 2013 [111, 145]. However, what will it take to introduce these commercial operations in densely populated urban areas? That is the focus of this chapter.

Firstly, section 2.1 gives an insight in current- and future-regulations on this topic. Subsequently, section 2.2 discusses the different types of missions for which UAS can be used. In section 2.3 Traffic Management architectures are discussed that specify the interaction between the different actors in the urban airspace. Lastly, section 2.5 discusses several modeling techniques and argues why the Agent-Based Modeling paradigm is well-suited for modeling these operations.

2.1. Regulations

Few industries are as strictly regulated as the aviation industry. Therefore, an essential part of UAS operations in Urban Areas is adherence to regulations. In the U.S., the Federal Aviation Authority (FAA) released rules for small UAS under part 107 of the Federal Aviation Regulations (FARs). Some of the important characteristics of FAR 107 include [36]:

- Maximum aircraft weight is 55lbs ($\approx$ 25kg).
- Maximum altitude is 500ft AGL (above ground level).
- Operation in Class G airspace is allowed without ATC permission.
- Operation in Class B - E airspace allowed with ATC approval.
- Visual Line of Sight (VLOS) operation is required - no beyond VLOS operations.
- Operations are restricted to daylight.

These rules are very restrictive regarding urban UAS operations. It is hard to see how urban UAS operations can gain serious momentum without the possibility of BVLOS operations. An economic study comparing the feasibility of VLOS vs BVLOS package delivery by UAS in cities indeed showed that this is infeasible [85].

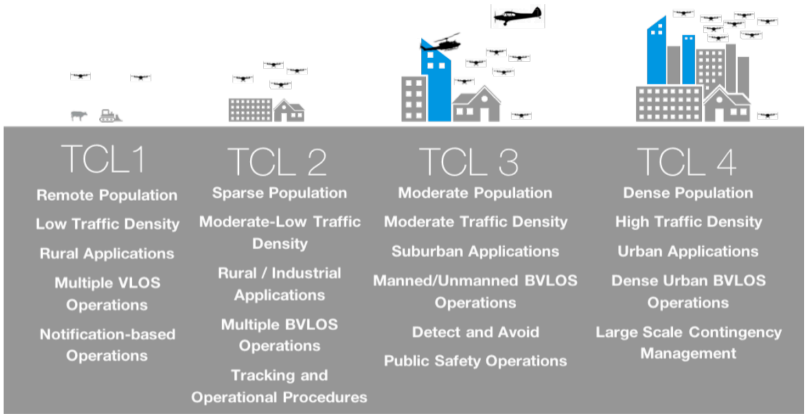

Figure 2.1: Planned roll-out of NASA Technical Capability Levels. Figure adapted from [70].

However, it is not expected that these rules will be the rules that govern the roll-out of Urban UAS operations in the coming decades. NASA is working on a concept of UAS Traffic Management (UTM) which consists of four stages, such as described in Figure 2.1. The ideas proposed by NASA are highly relevant for Urban UAS operations, as BVLOS operations over highly populated areas are included. The roll-out is defined in four *Technical Capability Levels*, which were introduced in 2015, and in the years since, NASA has proven the feasibility of each of the four stages with 21 industry companies. The feasibility of TCL 4 was proven in August 2019, with flight-tests in Nevada and Texas [93].

In the EU, similar steps are being taken. In 2015, EASA initiated a wide set of research activities for UAS operations, including the search for ways to safely accept operations of UAS in urban areas [29]. Since then, its research activities have been channeled into JARUS *(Joint Authorities for Rulemaking on Unmanned Systems)*, which is a group of experts from National Aviation Authorities from 61 countries, representing large parts of the world, including the US, China, India, and the EU [63]. A major achievement by JARUS is the development of SORA, which is a method for analyzing and mitigating the risk of an (urban) UAS ConOps. In subsection 4.4.4, this method is described in greater detail [71].

Reflecting on all regulations, the conclusion that can be drawn is that regulators all over the world are working on ways in which they can allow BVLOS operations with a high traffic density over dense populations. However, the rules are not yet set in stone. Therefore, in developing a ConOps for Urban UAS operations, one should not be overly focused on complying with either the current rules or even proposals of future rules, as they will likely be changed. One should instead focus on ways to create ConOps that are as safe as possible.

2.2. Applications

In cities, Unmanned Aircraft Systems can be employed for a wide range of applications. In this section, five of those applications will be discussed, being Package Delivery, Food Delivery, Passenger Transport, Surveillance and Medical Transportation.

Package Delivery

Globally, the E-commerce market is valued at around $3.5 trillion dollars, and growing at a rate of 18% per year [96]. Research shows that while people grow increasingly demanding of delivery times of the articles they purchased, the sheer physics of the delivery network are increasingly constraining, with congested cities delaying deliveries. Following these trends, many studies have indicated that using drones for "last-mile" delivery is a promising concept of operations [51, 69, 97, 117].

Currently, companies such as Amazon, UPS, and Alphabet (Google's parent company) are experimenting with drones that can deliver packages. Those initiatives focus either on last-mile delivery in densely populated urban areas (which can be commercially attractive because this scenario accounts for the most deliveries) or on deliveries in very remote areas (which can be commercially attractive because such deliveries are the most expensive).

Regarding delivery in urban areas, two possible scenarios are currently envisioned. The first is delivery to the front door of each home, while the second considers delivery to specific delivery-points throughout

the city. The second case would operate much like current pick-up points already being used, such as the system depicted in Figure 2.2. Both cases have their (dis)advantages. While delivering to each home is more customer-friendly, it would mean that drones should be capable of landing on unknown territory each landing, requiring advanced computer-vision capabilities. On the other hand, pick-up points may be less convenient to consumers, but are easier from a technological point of view, as one can create a closed system where all possible landing spots are known and can be pre-programmed [85].

Figure 2.2: Amazon pick-up point in the United Kingdom [130].

Demand forecast

To more accurately analyze the concept of operations, it is necessary to have an estimation of the amount of deliveries drones might do in the future. Such an estimation was made by researchers at TU Delft in 2018 and their findings will be presented in the remainder of this section [27, 117].

As this ConOps considers deliveries in urban areas, the degree of urbanization can be used to estimate packages delivered in cities. For this, the assumption is made that whether or not someone lives in an urban area does not influence the number of packages they order. For the US, Japan, Germany, UK, and France, this number is respectively 83.8, 94.6, 78.6, 85, and 81.4 percent. Also, Amazon indicated that 86% of all packages it delivers, are light enough for drone-delivery. Using these figures leads to the demand forecasted in Table 2.1 [3, 100, 102].

	US	Japan	Germany	UK	France
No. of parcels	12.0	9.5	3.0	2.2	1.5
No. of parcels in urban areas	10.1	9.0	2.4	1.9	1.2
No. of parcels in urban areas, <2.2kg	8.6	7.8	2.0	1.6	1.0

Table 2.1: Packages delivered in 2018 (billions) [27].

Food Delivery

The worldwide food delivery market is valued at $120 billion and is expected to grow swiftly over the coming decade. Furthermore, Mckinsey, a consultancy, conducted an extensive study into consumer behavior, which indicated that speed and punctuality are the most important variables in customer satisfaction [54, 123]. This section investigates whether UAS will able to serve (a part of) this market in the future.

With delivery-speed and punctuality being so important, drones delivering food are expected to add significant value in this market. In the media, many reports can be found of companies experimenting with drones that can be used to deliver food [52, 75].

Furthermore, research shows that delivering food using drones is potentially cheaper than using an E-bike. Doole et al. compared the expected cost of delivery by drone with delivery by E-Bike and concluded that delivery by drone costs between 0.23 and 1.35, while delivery by E-Bike costs between 2.02 and 2.04.

This implies that delivery by drone is between 1.5 - 9x cheaper than delivery by E-Bike. Although that is still quite a broad range, drones surely have the potential of being commercially viable in the delivery scene [27].

Even though several possible applications are being proposed for food delivery, it is most anticipated that there will be companies that act as the middle-man between a customer and a restaurant, arranging everything from taking the order and payment, to the delivery. They will thus employ a fleet of drones that are sent to take orders from many different restaurants across a city, and visit their hub in between operations for things like charging.

Passenger transport

In the US alone, traffic congestion forces drivers to waste over eleven billion liters of fuel and keeps them trapped in their cars for seven billion extra hours each year [129]. The notion that air transportation could decrease travel time was probably introduced not long after the first helicopter flight in 1939 and maybe did Leonardo Da Vinci already think of such a concept of operations while he made the first sketch of a flying machine in 1485. As of 2020, helicopter transportation around cities is on the rise, with companies like Voom and Uber offering the possibility for ordering helicopter rides around dense urban areas such as the San Francisco Bay Area or Manhattan, New York [132, 136].

However, mass aerial transportation is a much more interesting concept of operations than the helicopter services that are currently being offered, and that area is also heavily being researched by companies and researchers alike. For example, Uber is working on *Uber Elevate*, which (perhaps over-ambitiously) aims to commence commercial mass aerial transportation of passengers from 2023 onwards. For this project, it is cooperating with over 70 companies, including big names in the Aeronautical industry such as NASA and Boeing. Not only has the company generated futuristic renders of vertiports and electric vehicles with VTOL capabilities (such as depicted in figures2.3 and 2.4 [133].

In 2017, researchers at NASA's Langley Research Center together with researchers from Duke University analyzed possible concepts of operations for aerial passenger transportation. Their main conclusion was that the used ConOps is most influenced by the degree of autonomy of the vehicle used. On one hand of the spectrum, they envision vehicles which are fully controlled by a pilot and are designed to transport passengers across cities. From a perspective of functionality, such vehicles will be similar to current helicopters. On the other hand of the spectrum, a fleet of fully autonomous vehicles is envisioned, which fly without a pilot and can be ordered using an app. It is most likely that aerial transportation will start at the former scenario, and through the years, will evolve in the direction of the latter scenario [95].

As with many futuristic concepts, it is difficult to predict the future traffic volume that will be handled by such a concept of operations. Nevertheless, back-of-the-envelope estimations agree that it will be a significant part of traffic volume in cities. For example, an estimation done during the Metropolis research project, which researched radically new airspace design concepts and to which four research institutes contributed, among which TU Delft, estimated that between 15%-20% of traffic will be through the air in 2050 [55].

Surveillance

Visually inspecting objects on the ground from the air is not uncommon. In remote areas, drones are used to inspect crops or pipelines, with the European Commission naming drones as *one of the hottest trends in agriculture* [107]. In urban areas, aerial surveillance is mostly limited to helicopters. However, there definitely is potential for the use of UAS in aerial surveillance in cities.

Fontanella et al. surveyed for which surveillance tasks UAS will be employed in cities, and published a list of tasks like law enforcement, infrastructure inspection, traffic monitoring, aerial photography and mapping [39].

As these applications often have a high added value to society and are controlled by governments, it is expected that they will be introduced relatively easily. Indeed, as of 2020, drones are already used in, for example, the UK to break up gatherings [53]. Still, it is expected that surveillance drones will often fly customized missions using a low number of drones.

Medical Transportation

Since 2016, a US company called Zipline has been using drones to deliver blood across Rwanda, an east African country. This is more than just a test-case, as 65% of all blood-deliveries outside of the capital are done using Zipline's drones. As of March, 2020, the company made over 33,000 commercial deliveries [111, 145].

Next to drones currently being used to transport medical goods in rural areas, future scenarios also include medical drones in urban areas. An example of such a concept of operations is the *Ambulance Drone*,

which was developed at TU Delft and involves a drone that is equipped with an AED. For every minute extra it takes for an AED to arrive at a patient with a heart attack, survival rates decrease by 10%. Therefore, such a drone could significantly boost survival rates [91, 135].

It is expected that only a few of these drones would operate simultaneously at any given moment in a city. However, their operation would get the highest possible priority and other vehicles should move out of the way. Around incident zones, temporary no-fly zones will be created.

Reflection on Applications

In this section, five applications of drones for urban scenarios have been discussed. When analyzing the efficiency of different UTM structures on the safety of UAS operations, it is helpful to analyze operations that have a high traffic density and have simple missions. High traffic density because only then the effect of a UTM structure emerges, and little complexity because the easier the scenario, the more research focus can be directed to the UTM structures.

It is concluded that the application of Package Delivery is the best scenario. This conclusion was reached using Table 2.2. In this table, the complexity and traffic density of the five applications is discussed to find the best application for a research study on Urban UTM Structures. Complex applications are denoted with a (-), and high traffic densities with a (+). Package delivery is a relatively simple application, with an expected high traffic density. Food Delivery is a bit more complex because drones often need to make an extra stop compared to package delivery, namely at the restaurant. Passenger Transport requires (much) bigger vehicles than transporting small packages or food, and transporting people is a more dangerous business than transporting packages. Surveillance can lead to many custom missions and has a lower number of operations than the former three applications. The same holds for the traffic density of medical transportation.

	Complexity	Traffic Density
Package Delivery	+	+
Food Delivery	ś	+
Passenger Transport	-	+
Surveillance	-	-
Medical Transport	ś	-

Table 2.2: Comparison table of five urban UAS applications

2.3. Traffic Management in the Urban Airspace

In an Urban Airspace ConOps, it is important to consider which actors are present, and how they interact with each other. Several ConOps, published in academic papers, elaborated on this question. This section will reflect on the findings of this research, and concludes which (interaction between) actors in the Urban Airspace is not only the most sensible but also one that is generally accepted by researchers in the UTM field.

Proposed ConOps by NASA

As mentioned in section 2.1 on Regulations, NASA is working on *TCL 4*, which would allow BVLOS operations of UAS in dense traffic over densely populated areas. As part of the work towards *TCL 4*, NASA researchers published one conference paper, and two technical documents on the ConOps that would enable it [78, 103, 110].

The NASA research proposes a UTM Architecture that is divided into ANSP functions (Air Navigation Service Provider), an Operator Function, and a third category for other stakeholders. The first category is generally to be occupied by government organizations, while the second concerns industry operators [103]. Responsibilities of the ANSP include:

- Define and update airspace constraints and structures (including no-fly zones, geofenced areas, corridors, altitude requirements, etc).
- Manage an information-exchange that is widely accessible throughout the system with information on other UAS Operators and airspace constraints (Often referred to as FIMS, which stands for Flight Information Management System).
- Regulate access to airspace, for example, to manage capacity.

NASA divides the industry operators into UAS Service Suppliers (USS), UAS Operators, and Supplemental Data Service Provides. The first two have a vague boundary, but USSs are thought of as "managers" of UAS Operators, allocating tasks to them, and relaying information between the ANSP and other USSs to the UAS Operator, and vice-versa. UAS Operators are more concerned with safely executing the assigned task. Together, the USS and UAS Operator are responsible for [103]:

- Safe flying by avoiding other aircraft, terrain, obstacles, people, and animals.
- Respecting any constraints set by the ANSP.
- Detect, sense and avoid manned aircraft and/or other UAS
- Broadcast intent and operational plan according to ANSP guidelines.

Lastly, supplemental data providers are concerned with providing information that is relevant to a mission but not provided by the ANSP. This could include data on for example weather data, surveillance data, or performance-enhancing data [103].

Comparison with other ConOps

In the previous section, a UTM Architecture proposed by NASA is presented. By virtue of the reputation of NASA and the extensiveness of the research program of which this UTM Architecture is a part, the proposed structure is a good baseline. However, a comparison with other proposed UTM Architectures is still relevant. In this section, two other proposed architectures are considered. The first is by Jiang et al., which considered plans for drone deliveries by Google and Amazon and based its UTM Architecture on those plans [68]. The second, by Lin et al, proposes a UTM plan for Taiwan [84].

When comparing the different UTM Architectures, it is seen that there are many parallels. The most significant difference is that while NASA proposes that UAS Operators are free to choose their own flight plans, as long as they satisfy constraints set by ANSP and broadcast their plans, the architectures discussed in this section require the UTM to approve all flight plans. Even so, the NASA approach is given the advantage. The reason is twofold: firstly, the difference is only very subtle, because when a UTM decides whether or not to approve a flight plan, it does so by comparing the plan with its own constraints. When those constraints are always adhered to by UAS Operators in the first place, the flight plan will always be approved anyway. Secondly, in dense urban areas, it is expected that UAS Operators will generate flightplans very frequently, and will deviate from them often, as well. These deviations are unpredictable as they are caused by probabilistic events such as urban wind and maneuvers to avoid other traffic, but also less predictable events such as birds.

It is thus concluded that a UTM Architecture, such as proposed by NASA in [78, 103, 110], and depicted in Figure 2.6, is a good approach for managing the UTM in dense urban areas, especially for a drone delivery ConOps. The most important stakeholders are the ANSP, the UAS Operator together with the USS, and any supplemental data providers. This section has presented the responsibilities of each of these.

2.4. The Urban Environment

UAS Operations in Urban Areas are heavily influenced by the environment of a specific area. How does a system interact with this environment, and how is the environment modeled? Also, which environment is chosen? Operations will likely be vastly different when comparing New York and Delft.

When analyzing literature on UAS Operations in urban areas, it becomes obvious that authors often choose to model the city of their university. Research by M. Salleh et al. at NTU in Singapore model the city of Singapore [9, 10], work by La-Cour Harbo et al. at Aalborg University models (the surroundings of) Aalborg [79], and Primatesta et al. also models the city where the university of the authors resides, namely Turin [106]. As these papers are all considered as recent and state-of-the-art work on UAS Operations in Urban Areas, it can be concluded that there is not yet a sound method for choosing an appropriate city to represent in a model.

Different types of cities

American tourists visiting The Netherlands often notice that many cities are very similar, with a church in the center of the city surrounded by one or several *grachten*. Dutch tourists visiting the US notice the same phenomenon, with many US cities displaying a Manhattan-like structure with tall buildings in the city center.

This is however not just something that occurs to tourists, but a phenomenon that is mathematically proven by Louf et al., who discovered that there are only four different types of cities around the world and published about this in the *Journal of the Royal Society Interface* [87].

The work by Louf et al., characterized cities around the world by separating city blocks in small blocks (Area between 10^3 and 10^4 m^2) and medium-sized blocks (Area between 10^4 and 10^5 m^2). It then looked at the shape of the blocks in each set. The four different types of cities are characterized in figure Figure 2.7. It concluded that there are the following four different types of cities:

- Group 1: Blocks of medium size with almost only square shapes. The only city in this group is Buenos Aires, Argentina.
- Group 2: Blocks are predominantly of small size with a wide distribution of shapes. Athens, Greece, is mentioned as an example.
- Group 3: Similar to group 2 in the diversity in shapes of blocks, but there is a wider distribution in the size of the blocks. New Orleans, US, is the given example.
- Group 4: For example Mogadishu, Somalia, with very small square-shaped blocks.

It is thought that many European (And Dutch) cities fall in category 3. Although a categorization has drawbacks that there are always examples of cities that are in-between categories, or cities that do not fit in any of the existing categories (As may be the reason for the strange fact that Buenos Aires has its own category), the general conclusion that cities around the world are of very similar structure, is relevant for UAS research, as findings for one city can generally be applied to other cities.

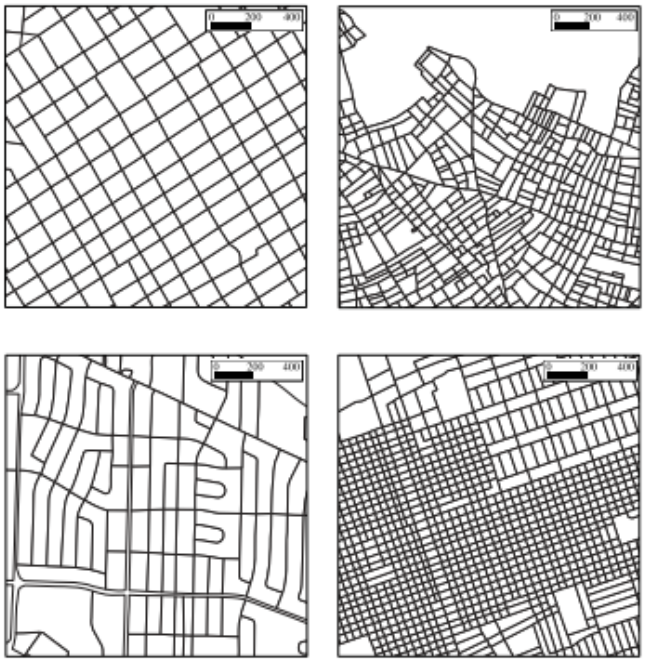

Figure 2.7: Four different types of city structures. (Top left (1), top right (2), bottom left (3), bottom right (4). Figure from [87].

Conclusions

It is concluded that there is little consensus on the most appropriate city to model an environment for Urban UAS Operations. As a starting point for work in this area, it is relevant that cities around the world generally fall in one of four categories. Next to the choice of a city, one should be concerned with how the environment inside a city is modeled. Two important characteristics of that environment are the structure of the Urban Airspace, which is discussed in chapter 3, and the modeling of the Risk to other actors in the environment, primarily humans, which is discussed in chapter 4.

2.5. Modeling Urban UAS Operations

The purpose of this section is to argue why the system described so far in this chapter can be best represented as a Multi-Agent System (MAS), which can be modeled using the Agent-Based Modeling paradigm (ABM). Firstly, a definition will be given of a Multi-Agent system, and subsequently, an argument will be presented that explains why the characteristics of an Agent-Based Model are suitable for modeling the proposed model.

In lectures on ABM in the course AE4422 at the Faculty of Aerospace Engineering, Dr. Sharpans'kykh defines a MAS as follows [120]:

> A multiagent system is a set of agents **interacting** in the **environment** to **solve problems**, **achieve goals**, or **execute tasks** that are difficult or impossible for a single agent

Because of the interaction between the agents in a MAS to achieve a common goal, these systems are defined as being *distributed* and *decentralized*. Furthermore, in the same lecture series, an Agent is defined as:

> An agent is an **autonomous**, computational entity that **perceives** its environment through its sensors and **acts** upon its environment through its effectors

Books by Weiss and Gilbert that elaborate on MAS and ABM provide similar definitions [41, 142].

The question is whether a system modeling UAS operations in the urban environment matches these definitions. Three arguments can directly be identified that support this view. Firstly, UAS and hubs can be viewed as agents because of their autonomous behavior. Examples of this is the autonomous decision of the UAS on how to fly through the environment, based on its perception of it, and based on interaction with other entities, such as the UTM, the hub and perhaps through interaction with other UAS. Secondly, these two examples also indicate that several separate entities collaborate through interaction to achieve a common goal, which fits to the definition of a MAS.

The third argument is that these operations can be modeled in a decentralized, distributed approach, which is customary to a MAS. The view that a decentralized and distributed approach is suitable to this model is supported not only by the examples regarding coordination between drones and several other entities but also by a literature review of similar academic projects, as presented in Table 2.3, which shows that a decentralized approach is applicable to this type of problem.

Another advantage of Agent-Based Models is their inherent modularity, allowing for easy addition of features to agents, of new agents, or the introduction of new elements in the environment. As will be described in chapters 3 and 4, elements such as Airspace Structures or a Safety analysis are logical additions to a model of Urban UAS Operations.

Not only does the inherent modularity allow for easy addition of these two elements, but these elements are also suitable to be represented by an ABM. Changing the environment through the introduction of airspace structures leads to changes to the system of which the effect on agent-behavior is not easily predicted. Through analyzing the emergent behavior of this system in an ABM, the effects of these structures can easily be deducted. The same holds for a Safety analysis. Dr. Sharpanskykh argued that safety properties often stem from complex, nonlinear dependencies [121]. Again, the ABM paradigm is suitable for capturing the emergent properties following from this system.

It can thus be concluded that the components of Urban UAS Operations can be modeled as entities that behave autonomously by interacting with the environment, called *agents*. Furthermore, these agents can cooperate to achieve common goals that are more difficult for an individual agent. Also, such a system can be set up in a decentralized and distributed fashion. All these characteristics indicate that an Agent-Based Model is a suitable paradigm for this concept of operations. Also the emergent properties following the introduction of other components, such as Safety or Airspace Structures, can be captured efficiently and adequately by an ABM.

Author	Research Purpose	Number of UAS	Modeling paradigm
Hoekstra, 2015 [55–59]	Model high-traffic scenarios in the urban airspace	Many	Decentralized
Salleh, 2017 [9]	Effect of # UAVs on # of conflicts	Many	Decentralized
Salleh, 2018 [10]	Effect of airspace structure on airspace capacity	Many	Decentralized
Primatesta, 2020 [106]	Ground risk estimation	One	Centralized

Table 2.3: Modeling paradigms used in relevant Urban UAS Operations research

2.6. Synthesis on Urban UAS Operations

In this chapter, five areas of UAS Operations have been investigated. These areas include the regulations governing Urban UAS Operations, the different applications of UAS in Urban areas and the corresponding mission profiles, the different actors in Urban UAS Operations, how the urban environment is modeled, and which modeling paradigm is most suitable for Urban UAS Operations. This section synthesizes these findings, and together, the conclusions for these five areas make up a Concept of Operations that can be used for research into (the safety of) future urban UAS Operations.

Regulations

Even though current regulations in both the EU and the US prevent BVLOS operations over dense Urban areas, regulators are currently working on proposals for future regulations that would allow such operations. Therefore, one should not be limited by the current regulations, but instead, create models that are likely to be in line with future regulations. As the most important aim of future regulations is Safety, not only to the ones directly involved with the operations but also (and especially) to the ones indirectly involved (such as people walking cities at risk of being impacted by a crashing UAS), it is paramount that future UAS concept of operations makes safety a central aspect of their approach.

Applications

Although UASs can be useful in many different applications in Urban areas, the application of drone delivery is a good starting point for research on UAS Operations in Urban Areas, as the application is relatively not complex and there is a high expected traffic density.

Actors in the Urban Airspace

It is found that a UTM architecture such as in Figure 2.6 is a good approach for the Urban airspace. Two important differences to regular ATM are (1) that UAS Operators can fly without the need of waiting for the approval of a flight-plan, as long as they adhere to constraints set by the UTM, and (2) that information on the airspace and on the intent of individual UAS is widely available, and not only to the ATC.

The Urban Environment

Current research on Urban UAS Operations is often simulated in the city of the university of the primary researcher. It is difficult to judge how well the findings would generalize to other cities. However, researchers on the topology of cities found that (almost) all cities adhere to one of four basic patterns. Assuming that results generalize between cities in the same topology-category, one can find results for cities all over the world by simulating just a few cities.

Modeling UAS Operations

The different entities participating in Urban UAS operations can be modeled as autonomous agents, which aim to achieve a common goal through interaction. These properties make it suitable for the model to be described as a Multi-Agent System. Furthermore, elements such as airspace structures and a safety-analysis can be easily added to this system in a later stage, and Agent-Based Modeling is a suitable method for analyzing the emergent properties of these elements.

Figure 2.3: A vertiport.

Figure 2.4: An eVTOL vehicle.

Figure 2.5: Renders by Uber Elevate of their mass aerial transportation concept of operations [133].

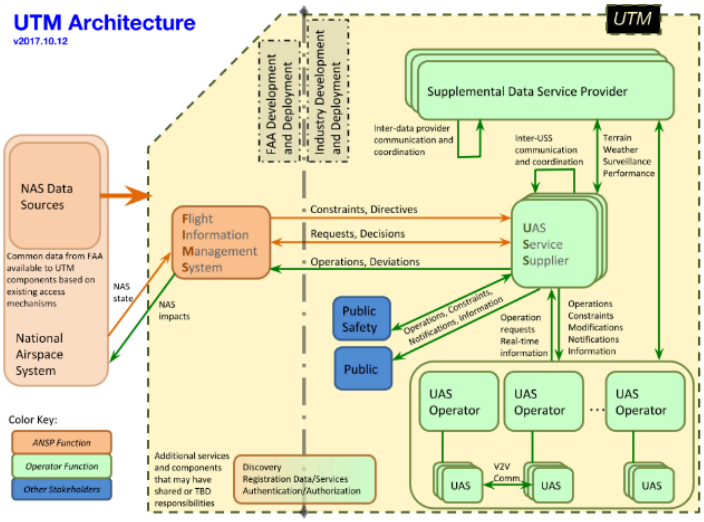

Figure 2.6: Proposed UTM Architecture, figure from [110].

3

Structure of airspace

As described in the previous chapter, the only certainty regarding regulations for UAS in Urban areas, is that these regulations will change in the coming years. Many different parties, including Aerospace giants like Airbus and governmental organizations such as the FAA and NASA, have suggested that the same applies to airspace structures [37, 115, 117].

The goal of this chapter is to investigate how the (urban) airspace might be structured in the future. It starts with a brief overview of current airspace structures in section 3.1, and continues with an evaluation of proposals for- research on future airspace structures in section 3.2. Subsequently, a method of conducting traffic management through airspace structured is discussed in section 3.3, and the chapter is wrapped up in section 3.4, where three research gaps are identified that could be starting points for future research.

3.1. Current airspace structure

This section gives a brief overview of the current airspace structures. These were established in 1990 when ICAO adopted the airspace classification system that is used today [61]. In this system, airspaces are classified in classes denominated by the letters A - G, however, in practice, F is not used. Even though there are subtle differences in how the classes are applied in different countries, they are usually categorized as follows:

- **Class A:** High-altitude airspace (Between 18,000ft and 60,000ft above ground-level in the US)
- **Class B:** Airspace around major airports
- **Class C:** Airspace around medium-sized airports
- **Class D:** Airspace around small airports (that still have a control tower)
- **Class E:** Controlled airspace outside of airports
- **Class G:** All airspace outside of classes A-E, uncontrolled

Next to Class G airspace, Air Traffic Control does not cover the airspace that is below 400ft AGL. In both of these areas, flight operations are cooperatively managed between airports, and interaction with ATC is limited [35, 37].

3.2. Future airspace structures

This section discusses possible future structures of the urban airspace in section 3.2, and presents research that evaluates the (dis)advantages of these airspace structures for urban UAS operations in section 3.2.

Possible structures of urban airspace

This section will discuss four ways of structuring the airspace in urban environments. These four ways are a *free airspace*, an airspace with *zones*, an airspace with *tubes*, and a *layered airspace*.

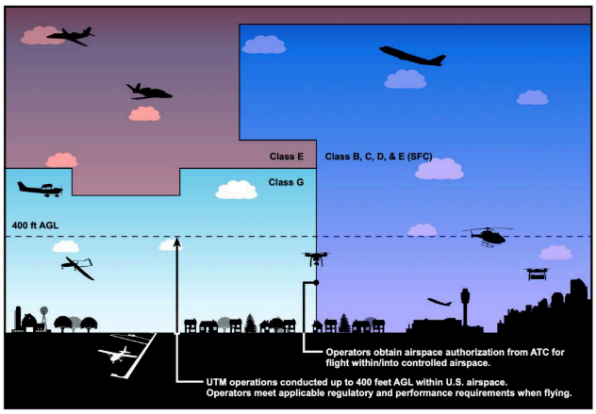

Figure 3.1: Current airspace structure [37].

Free skies

The first concept of airspace structures has no structures at all. This is the free skies concept. In literature, it is also described as the *Full Mix* concept [56] or the *AirMatrix Network* [10]. It is needless to say that such an airspace provides a lot of possibilities for routing, however, it also provides the most possibilities for conflicts. Those conflicts can exist not only between UAVs but also between UAVs and other air traffic, such as planes and helicopters. Also, obviously, this airspace is static so there is no need for a manager to update it.

Layers

In the concept of layers, each altitude range corresponds to a heading range. This forms a pattern that can repeat itself. The goal of this structure is to allow for maximum freedom of routing while lowering the difference in speed vectors, which benefits safety [56].

Zones

The zones-concept is inspired by current Air Traffic Principles, in which different zones are used for different types of vehicles and different types of missions. This concept allows for many variations, including:

- Specific zones to separate specific types of aircraft. For example cargo- and passengers-aircraft could be separated in the airspace. This could make sense because different aircraft types have different characteristics in terms of maneuverability and speed.
- Specific zones for one type of drone. For example around hospitals, a zone can be created that can only be accessed by one type of drone. Or only drones with a high safety rating are allowed to enter zones that exist over densely populated areas.
- No-go zones. These zones cannot be accessed by any drone.

These zones could be created and adjusted dynamically, based on external factors. One could envision a no-go zone around a park which is only active if there is a certain amount of people in the park, causing a higher risk, or a zone around an accident that can only be accessed by drones of the emergency-responders [56].

Tubes

Instead of specifying where aircraft cannot fly, such as in the Zones concept, the Tubes concept specifies the area (or tube) where the aircraft is allowed to fly. This creates a dense route structure, similar to a network of (rail)roads. This type of network allows for a very good integration with the regular infrastructure. Some variations of this type of network were proposed by Salleh et al. and Hoekstra et al. [10, 56].

- Over-roads network, in which there is a tube above every road in a city.
- Over-buildings network, in which the rooftops of all buildings in cities are connected with tubes.
- Aircraft-specific tubes, in which each type of aircraft has its own set of tubes.

The tubes-concept is highly dynamic and comes with a lot of variations. Tubes could be dynamically generated to for example account for route-demand of operators, or to avoid third-party-risk.

Evaluation of different airspace structures

The applicability of the airspace structures introduced in the previous section on Urban UAS operations was researched extensively as part of the *Metropolis* research project. This project was conducted between 2014 - 2016, and aimed to investigate *"radically new airspace design concepts for scenarios, which are extreme when compared to today in terms of traffic density, complexity, and constraints"*. The project was coordinated by Prof. Hoekstra at TU Delft, and NLR, DLR and ENAC contributed to it [55–59]. The findings of this research project are presented in the remainder of this section.

Research Scenario

In Metropolis, four different airspace structures were analyzed for four different urban scenarios. The four airspace structures were *full mix*, *layers*, *tubes* and *zones*. The layout of these structures is as described in section 3.2. Also, four scenarios were created for urban density. These scenarios are shown in table Table 3.1. When comparing the scenarios with Amsterdam and London, it is clear that the Metropolis scenarios have both a higher population and a significantly lower population density. The population density in Metropolis is more similar to the Randstad, but then still, the population in the busiest scenario is over three times larger. The Metropolis scenario thus considered an extremely large and populated area, leading to a high number of flight movements. Furthermore, it considers both passenger-traffic and delivery of all sorts of goods, such as food and cargo.

	Population [million]	Population density [population / km^2]
Scenario 1	14	974
Scenario 2	18	1253
Scenario 3	22	1531
Scenario 4	26	1809
Amsterdam	0.8	5000
London	8.8	5900
Randstad	8.2	1018

Table 3.1: Population and population density in four Metropolis scenarios, and in Amsterdam and London [21, 55, 98, 131].

*

Experiment setup

The Metropolis-experiment compared the performance of the air structures using a wide range of metrics, which were categorized in three categories:

Organizational metrics

Characterises the *complexity* of the traffic situation that emerges from the used airspace structure. A more complex situation implies that it is more difficult to control in a safe manner. The researchers measured complexity in two ways, namely using a *proximity indicator* and a *convergence indicator*. The former indicates whether there are areas where there is a high level of aggregation of aircraft, while the latter compares the speed vectors between aircraft to analyze whether or not they are converging.

Operational metrics

The end-user is most concerned with the operational performance of a flight. This operational performance can be characterised by safety, stability, efficiency, and capacity. Operational metrics are calculated in the following four categories:

- **Safety**
  This metric mainly considers whether an airspace structure facilitates in maintaining safe separation between aircraft. This is assessed by measuring three elements: the number of conflicts, the number of intrusions, and the severity of intrusions. Here *intrusion* is defined as a loss of the minimum separation requirements, the severity measures by how much the minimum separation requirements are violated and conflicts are defined as "predicted intrusions", meaning that the track of an aircraft will intrude the minimum separation requirements of another aircraft within 60 seconds.

- **Efficiency**
  Six metrics were used to analyze the efficiency of a concept. Many are similar, and come down to comparing the route flown with the distance as the crow flies between origin and destination.

- **Stability**
  The stability is a measure that quantifies how many conflicts arise, because previous conflicts are solved. Such conflicts are a "domino-effect", and the parameter measuring this effect is very fittingly called the *Domino Effect Parameter* (or DEP). It measures how many additional conflicts arise, because earlier conflicts are solved, and can be measured by analyzing a simulation with and without conflict resolution. All conflicts that arise with conflict detection on, but do not arise with conflict detection off, are thus caused by the domino effect.

- **Capacity**
  The Metropolis authors concluded that it is difficult to analyze the capacity of an airspace structure, and observed that there is some "inflection point" where increased traffic density results in big changes in safety and efficiency metrics. Therefore, the capacity is analyzed qualitatively, by analyzing the results from previous categories.

Environmental metrics
The impact of flights on the environment is already a hot topic in Air Traffic Management today. The issues that we currently discuss for ATM, such as noise, pollution/emissions, and third party risk will be just as important to consider for drones in urban environments. As the authors of Metropolis expect emissions to be less of an issue in 2050 (the date for which the scenarios are created), their main focus is in Noise and Third-Party Risk. They are calculated as follows:

- **Noise**
  The LDEN noise footprint is calculated. This is a metric that calculates the cumulative effect of all individual flights.

- **Third Party Risk**
  TPR is the involuntary exposure of people on the ground to an aircraft accident. In the Metropolis research, the following calculation is proposed: the chance of an aircraft crashing, the location of the crash area, in case of an accident, and the lethality chance based upon the aircraft state and the location parameters. However, this calculation was not performed to obtain the results, instead, the authors chose to (over-)simplify the analysis by assuming that TPR is linear with distance flown.

*

Results

In this section, the results of the Metropolis experiments are presented in each of the three categories that were introduced in the previous section.

Organizational metrics
The proximity- and convergence metrics are shown in Figure 3.2. It can be seen that there is a clear ranking between the concepts, with the tube concept resulting in the least complex system and the zones concept being the most complex. The authors indicate that the zones concept performs especially poor because they set up the zone-concept with concentric rings and radials. These radials have convergence points, where all traffic comes together, causing a lot of complexity. It is expected that the layers-concept performs a little better than the free-mix, as aircraft are separated with regards to the heading they are flying.

The Metropolis researchers also investigated what would happen if aircraft deviate from their 4D-trajectory in time, i.e., how robust the metrics are. Inspecting these results, which are displayed in Figure 3.3, shows that the metrics are the least robust for the tube-concept. This is because the tube-routes are planned beforehand, and are flying closely to each other (in the same tube), such that small changes have a bigger impact on the proximity-metric.

However, there is a reason to be critical of these results. Looking only at Figure 3.2, it seems easy to decide which concept is the least complex. However, in the tubes-concept, the authors used pre-flight conflict-detection and resolution (CD&R) and even reject flight plans if they do not fit into the airspace, while the layer- and full-mix use direct routing and only do CD&R when encountering a conflict. If the same pre-flight CD&R was applied to the other concepts, the results would have been easier to compare, and probably very different. Furthermore, with pre-planned CD&R, one would expect the proximity metric to be zero, as it measures losses of separation. The authors indicate that this is caused by discrepancies in the aircraft model used in the planning-model and simulation-model. One could thus say that the measured proximity is not a result of the inherent characteristics of the tube-model, but rather a measurement of a discrepancy in flight dynamics models used. A lot of caution is thus required before reaching conclusions on which flight structure is the least complex, based on these results.

Operational metrics

As defined in the previous section, operational metrics are calculated in terms of safety, stability, efficiency, and capacity. For brevity, this literature survey only presents the conclusions of the Metropolis research, and does not provide the full results.

- **Safety**
  The researchers presented results for conflicts between two PAVs (Passenger Aerial Vehicle), and between PAV and UAVs. In the former case, the tube-concept performs poorly, however, it performs the best in PAV - UAV conflicts. The researchers conclude that this difference is explained by the tube-topology, which is centered around UAV hubs. The main conclusion is thus that safety is very much dependant on tube-topology. Full Mix and Layers perform similarly, while Zones always perform poorly.

- **Efficiency**
  As explained in the previous section, efficiency relates to whether the drone takes the most efficient route from A to B. The (obvious) result is that with both the Full Mix and the Layers, this is the case. In the Zones-case, routes are between 10% - 25% longer and in the Tubes-case, the routes are twice as long. However, this is as well very much dependant on the topology of the tubes.

- **Stability**
  As the authors did not implement in-flight CD&R for the Tubes-concept, the stability could not be measured. For the other three concepts, there was only instability visible for the zones-concept in the highest traffic density scenario. Other than that, the Domino Effect Parameter was between -0.6 and -0.9, implying a stable situation because resolved conflicts rarely lead to new conflicts.

- **Capacity**
  Capacity was analyzed by qualitatively judging the gradient of safety metrics w.r.t. traffic density. For the Zones, but mainly for the Tubes-concept, the conclusion is (again) that it very much depends on the topology. It is expected, however, that the theoretical capacity of a Tube-concept is less than the theoretical capacity of a Full-Mix concept since more airspace is available.

Environmental metrics

Environmental metrics consider how much the urban UAS operations, such as described in the Metropolis scenario, influence the life in the city. Again, only the conclusions are discussed here for brevity.

- **Noise**
  Two main conclusions arise. Firstly, rather unsurprisingly, that the noise very much adheres to the structure of different airspaces. This implies that with the FM and the layers structure, the noise is spread evenly across the map. In the tubes concept, the noise is also spread evenly over all tubes-locations. With the zones-concept, the noise intensity is very high in some points but low in others. The Metropolis research does not comment on the overall noise intensity in each concept.

- **Third-party Risk**
  The risk-model used assumes a linear relationship between time in cruise and third-party risk. Therefore, as the tube-concept flies about twice as long, it has twice as much risk. This simplification ignores

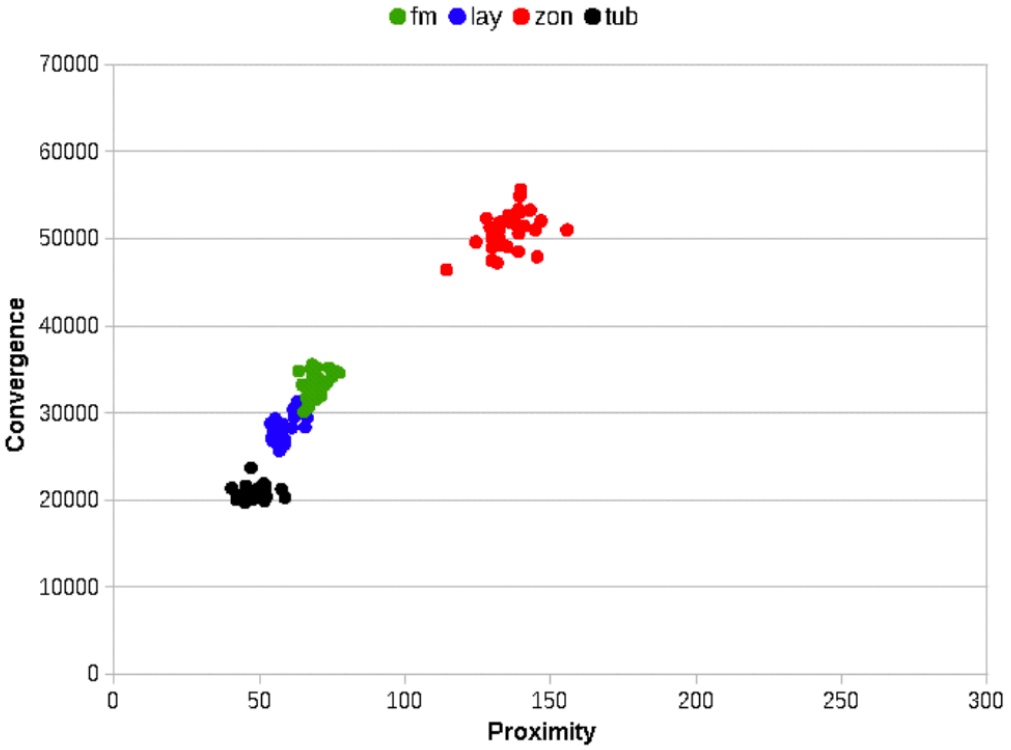

Figure 3.2: Normal conditions.

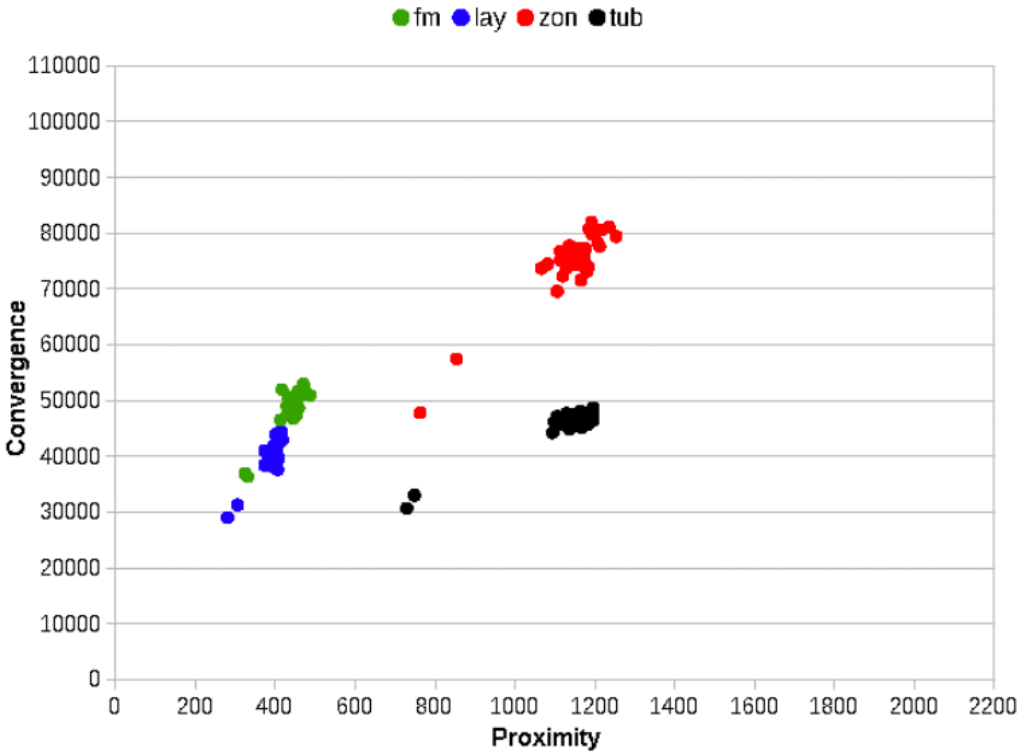

Figure 3.3: Abnormal metrics (robustness).

Figure 3.4: Complexity- and proximity-metrics in four airspace structure scenarios from Metropolis research [58].

whether drones fly over crowded or less-crowded areas during their cruise phase, and therefore, these results are not useful for drawing conclusions on the risk.

3.3. Employing airspace structures for UTM

The research findings presented in the previous section thus conclude that airspace structures can be especially beneficial to reduce the complexity of the traffic. In other words, these structures can be employed as a means of Urban Traffic Management (UTM). Two research projects sought to bring this idea into practice, and these projects are discussed in this section. The first is a project by Salleh et al. from NTU in Singapore, which aims at researching the influence of airspace structures on the capacity and throughput of the environment, and is discussed in section 3.3 [9, 10]. The second project is presented in Table 3.3 and is by Chang et al. from Oregon State University, and employs controller agents that adjust the travel costs at individual parts of the network with the goal of optimizing global throughput [15].

Influence of airspace structures on capacity and throughput

M. Salleh et al. have published two papers that propose a preliminary concept of adaptive urban airspace management for UAS. They focus on operations of small and lightweight ($\leq$ 25 kg) vehicles doing aerial deliveries (food, mail, light parcels) throughout the city of Singapore. The research builds on the authors' belief that in order to manage large traffic densities in cities, it is essential to structure the airspace. Their research then focuses on criteria that indicate when the airspace should transition from free flight to a structured one, and to what structure the airspace should then transition. This section presents their work on the maximum capacity of an airspace, and subsequently presents an experiment done in which the performance of a free airspace is compared with two structured variants [9, 10].

Maximum capacity of an airspace

An experiment was set up using the city of Singapore in order to determine the relationship between the number of UAVs and conflicts. In this, UAVs are of the same type, flying at a constant speed and spawning simultaneously at entry-points which are scattered around the city. A conflict is defined as a case in which the separation between two UAVs is less than 100 feet. The found results are presented in figure 3.5(a), which shows that the number of conflicts rises quadratically with respect to the number of conflicts.

The authors performed a second experiment in which they kept the number of drones per entry-point constant, and instead scaled the number of entry-points. The results of this experiment is shown in Figure 3.5 and also shows an $O(n^2)$ relationship.

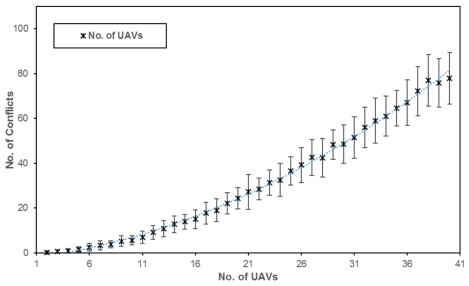

(a) Effect of the number of UAVs on the number of conflicts.

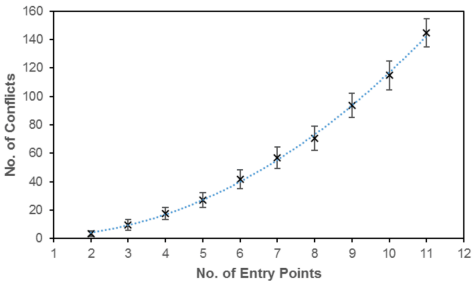

(b) 3D-plot of descent trajectories

Figure 3.5: Effect of the number of entry-points and the number of conflicts.

Capacity and throughput of three urban airspace concepts

As showed in the previous section, Salleh et al. concluded that the number of conflicts scales $O(n^2)$ with regards to traffic density, and made the case that the airspace should become progressively more structured as traffic density rises. In a second experiment, they analyzed the maximum capacity and throughput of three different structures of the airspace. The authors did not choose the networks over buildings and -roads arbitrarily. Flying over these urban infrastructures will reduce the exposure to the urban population, or as the authors put it, *"provide a safe haven for UA operations"*. These three structures are:

- **Route Network in AirMatrix**
  The AirMatrix is a representation of the entire airspace in nodes and links. All nodes and links are permitted, apart from those that conflict with objects (such as buildings and trees). The network thus covers the entire airspace, and is therefore *homogeneous and exhaustive*.

- **Route Network over buildings**
  All nodes are created over buildings, and links are designed in such a way that they follow buildings as much as possible. Roads can only be crossed in order to link the nodes above two buildings to each other. This network obviously does not cover the entire airspace, and is therefore heterogeneous and inexhaustive.

- **Route Network over roads**
  This network is very similar to the previous one, however, the nodes are now placed over roads and not over buildings. The nodes are all placed at a certain height and are linked in such a way that the links follow the natural flow of the roads.

On these networks, a last-mile delivery scenario is simulated to analyse throughput and capacity quantitatively. In this, capacity is defined as *"the maximum number of UA that can be hosted by the network in a given airspace at any point of time"*, and throughput is defined as *"throughput is defined as the number of UA that land to any supply point over a specific time frame"* [10]. The number of UAVs will be increased until all possible routes are saturated. Furthermore, the authors simulated not only a scenario in which all UAVs fly at the same speed, but they also simulated a scenario in which the UAVs fly at either 13, 14, 15, 16 or 17 m/s.

The capacity is calculated using Equation 3.1, in which r is the number of possible routes, $n_i(t)$ is the number of UAVs on route i at time t. Throughput is calculated using equation Equation 3.2, in which p_t is

the throughput over a time duration of ΔT, and $k_i(t)$ is the total number of UAVs that have completed route i.

$$p_c = \lim_{t\to\infty} \sum_{i=1}^{r} n_i(t) \tag{3.1}$$

$$p_t = \Delta T \lim_{t\to\infty} \frac{\sum_{i=1}^{r} k_i(t)}{t} \tag{3.2}$$

A summary with the most relevant results is given in Table 3.2. The (fully connected) Airmatrix shows the best result, while a network over roads shows approximately a 20% drop in both capacity and throughput. Compared to the fully connected network, the network over buildings shows a 60%-70% drop in capacity and 40%-50% in throughput. Even though the specific figureswill differ from city to city, it is clear that the more constrained networks show a significant drop in performance. Furthermore, random cruise speeds also perform poorly when compared to constant cruise speeds. The most likely cause for that is that a lot of extra space is needed between drones to allow for sufficient scheduling.

Route Network	Cruise speed	Capacity	Throughput (5min window)
Airmatrix	Constant	255	751
	Random	88	286
Over Buildings	Constant	77	345
	Random	34	168
Over Roads	Constant	206	604
	Random	69	220

Table 3.2: Summary of the performance metrics of the three networks

Learning dynamic traffic management strategies

The goal of the work of Chung et al. is to that a multi-agent-based dynamic traffic management can successfully improve throughput in a network. They do so by employing controller agents whose actions are to adjust the perceived costs of traveling across different parts of the network. The throughput of the network is then improved if these agents can come up with a policy that modifies the costs in such a way that motivates the robots traveling through the network to take routes that maximize throughput and travel time.

Problem formulation

A network is proposed which is the directed graph $G = (V, E)$, consisting of edges $e \in E$ connecting the vertices, with each edge having a maximum capacity cap_e. The cost of traveling across a node consists of the time to traverse it, $cost_{travel}$, and an additional cost $cost_{add}$. Each robot uses a regular path-finding algorithm to find the route with the lowest cost.

The overall goal of the system is to minimize $G(t)$, being the normalized total travel time of all robots in the system. The travel time consists of both time spent traversing along edges, and time spent waiting at vertices before being able to enter the next edge. This goal is given in Equation 3.3, where $\mathcal{J}(t)$ is the set of all robots having entered the system until time t, and $\mathcal{J}'(t) \subseteq \mathcal{J}(t)$ being the robots having reached their destination.

$$G(t) = \frac{\sum_{j\in\mathcal{J}(t)} \text{total\_time}_j(t)}{\sum_{j'\in\mathcal{J}'(t)} j'} \tag{3.3}$$

In this system, N agents are generated, where N is equal to the number of edges in the network. Each agent thus corresponds to one edge. The state of each agent is the current number of robots traveling along the corresponding edge, summed with the total number of agents waiting to start traveling along this edge. The policy π_i of agent i is then to modify $cost^i_{add}$, being the added cost to traveling along edge i. The overall goal of the multi-agent system is to learn the policies Π for which the total travel time $G(t)$ is minimized.

The policies are defined as a neural network with one input (the sum of robots traveling on an edge and waiting to enter it), one output (the added cost to traveling on that path), and 20 hidden nodes in between.

In order to learn the policies, Cooperative Coevolutionary Algorithms (CCEA) have been applied to learn the policy of each of the agents. Such algorithms are needed because each agent must optimize its own policy in a non-stationary environment. It may be the case that an agent changes its policy in a way that benefits the result, however, because of the modifications of other agents' policies, it receives a negative reward signal.

Experiments and results

To experiment with the effectiveness of the approach in the previous section, Chung et al. developed a network as in Figure 3.6. Robots are spawned at vertices 0 and 1, which are connected by diagonal "highways". When these highways are congested, the system can improve the total travel time by incentivizing robots to travel along the "small roads".

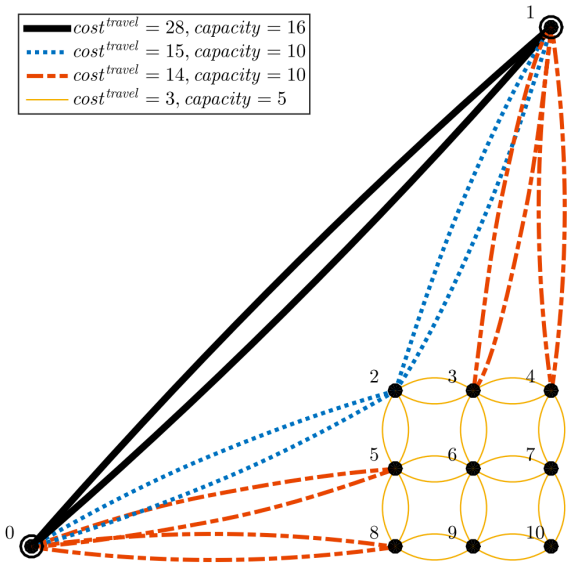

Figure 3.6: Network considered in dynamic traffic management [15].

Although the authors have conducted numerous experiments with varying traffic densities, and even included a "fixed toll" as a control parameter for the added cost by the agents, it is beyond the scope of this literature survey to elaborate on the full results. Nevertheless, the most important conclusion is that using a multi-agent team, trained over 1000 learning epochs, achieves significantly lower travel times than when only the *known costs* are used, which are the actual costs of traveling an edge. Furthermore, the result of the multi-agent teams were compared to a scenario in which the *expected costs* are used, which includes the expected waiting time of the robot caused by delays further along the path. This scenario shows very similar results, and this implies that local agents with only local information and a simple network for their policy, can learn to adjust the costs to account for delays cascading in throughout the system.

It can thus be concluded that a multi-agent setup consisting of agents with only local information and a simple network structure as its policy can learn to incentivize robots to take more optimal paths through the system. Also, the fact that it shows similar results when compared to a cost-structure that incorporates knowledge about delays further along a path, shows that the agents can learn the effects of delays cascading through the network. All in all, such a distributed way of managing traffic in a network is a very relevant way of dynamically managing traffic [15].

3.4. Reflection on Airspace Structures

In this chapter, four ways of structuring the (urban) airspace are presented, as well as research that studied the effectiveness of those structures. It is concluded that structuring the airspace with elements such as no-fly zones and tubes is considered a promising concept, as it would allow reducing the complexity of the airspace. There are drawbacks to those structures, however.

The research by Hoekstra et al. showed that when the topology of the structures is not designed carefully, it could lead to adverse results. Furthermore, it showed that an airspace structured with tubes results in the

least complex traffic situation [58]. Also, in extremely high traffic densities, the less structure, the higher the capacity [10]. Lastly, Chung et al. showed that agents with only local information can learn policies that improve the performance of a system on a global level, even if their policies are only applied locally [15].

*

Research gaps

It is established beyond doubt that the structuring of airspace is an important topic in Urban UAS Operations. But what is next in terms of research? Three research gaps were identified:

1. Having established that the topology of airspace structures is so important, research should be conducted on how different topologies influence system performance, and into the development of algorithms and/or heuristics that can optimize such topologies, based on an environment.

2. Hoekstra et al. acknowledged the influence of airspace structures on TPR, but have not conducted any TPR analysis other than assuming that it is linear with distance flown. An analysis of the influence of airspace structures on TPR using a more accurate representation of the risk is a logical next step.

3. The finding that local agents can successfully improve overall system performance can be applied to the topology of airspace structures. One could conduct research with the objective of creating a multi-agent team of agents that influence local airspace structures, improving system performance on a global level by doing so.

4

Risk

If one took no chances, one would not fly at all. Safety lies in the judgment of the chances one takes. That judgment, in turn, must rest upon ones outlook on life.
– Charles Lindbergh

The purpose of this chapter is to discuss the role of Risk in (urban) UAS Operations, to describe different methods of modeling Risk, and discuss which of these methods is the most relevant for a drone delivery ConOps in an urban environment.

As described in the previous chapters, UAS in urban areas is a rapidly growing field. However, some barriers are to be conquered before drones will massively roam the streets of cities. One of these barriers is the case of risk, and specifically third-party-risk (TPR). The fact that risk is such a barrier to a wide-scale introduction of UAS operations, is supported by the fact that lawmakers working on UAS regulations heavily focus on it. For example, EASA indicated in 2015 that a risk-based approach should be central to the development of a regulatory framework [29], and EU lawmakers indeed accepted such a system into law in 2019 [34]. Their American counterpart, the FAA, has a similar focus on risk. For example, it published a ConOps for UTM in 2020, which refers to risk as one of the most important factors in the design of such a ConOps [37].

But what is exactly meant with the term *risk*? In this work, a definition will be used that was introduced by Kaplan et al. in 1997. It defines Risk Analysis as the combined answer to three questions: (i) What can go wrong? (ii) How likely is it? And (iii) what are the consequences? [73]. Although this definition is used often, it has no monopoly on the field. Another definition of risk is the combination of likelihood & impact, which is still strikingly similar to Kaplan's definition.

Within this definition of risk, risk following from UAS Operations can be categorized into three categories [20]:

- **First parties:** people and property directly involved in the operation of the UAS.
- **Secondary parties:** People and property not involved in the operation of the UAS, but directly driving benefit from its operation (i.e. a passenger transported by a drone).
- **Third parties:** people and property neither involved nor benefiting from the operations.

Considering the acceptance of drones in urban areas, it makes sense to focus primarily on third-party risk (TPR). Not only is the number of people exposed to TPR following drone delivery operations higher than the number of first- and second-parties, they also have no choice but participating in the operations, as they have little influence on whether or not they would be overflown by a delivery drone while strolling through a city.

To discuss TPR, it is useful to have a mathematical definition of it. Recent work on TPR for (urban) UAS operations use slightly different ways for writing expressions that are mathematically equivalent [74, 79, 106]. It is not important to elaborate on these subtle differences, instead, it is most relevant to conclude that modeling TPR such as in Equation 4.1 is equivalent to recent works on UAS TPR for an individual path. In this

equation, a path p represents an individual flight operation, and all paths are combined in the set P, such that $p \in P$. M is the set of all (x,y) coordinates on the map. Furthermore, Equation 4.2 represents the total risk of all paths in an environment.

$$\mathbb{E}[\text{Risk(p)}] = \underbrace{\mathbb{P}(\text{failure})}_{\text{section 4.1}} \times \int_{(x,y)\in M} \underbrace{\mathbb{P}(\text{impact at x,y})}_{\text{section 4.2, 4.3}} \times \underbrace{\mathbb{E}\left[\text{fatality} \mid \text{impact at (x,y)}\right]}_{\text{section 4.4}} \tag{4.1}$$

$$\mathbb{E}[\text{Risk}] = \sum_{p\in P} \mathbb{E}[\text{Risk(p)}] \tag{4.2}$$

The probability of impacting a person can be subdivided into three components: (1) the PDF which relates a failure location to possible impact locations using descent models, (2) the expected population density at that location, denoted by $\rho(x, y)$ and (3) the size of the impact area A_{exp}. Combined, this leads to Equation 4.3.

$$\mathbb{P}(\text{impact at x,y}) = PDF \times \rho(x, y) \times A_{exp} \tag{4.3}$$

Having the definition of TPR established both in plain text and mathematically, the following question is how much TPR should be allowed. JARUS indicated that one death per 10^6 flight hours is reasonable [71]. For comparison, the CAA (UK Civil Aviation Authority) concluded that the number of fatalities in general aviation is 12.7 per million flight hours, taking into account all accidents around the world between 2002 and 2011 [12]. Note that while the former figure by JARUS is focused on third-party risk, while aviation fatalities are almost only first- and second-parties. The figure used by JARUS is thus 10 times more conservative than what is currently accepted in general aviation.

In this introduction, it is concluded that TPR is perhaps the most important barrier to the full-scale introduction of UAS in urban environments. Therefore, this chapter continues with a deep analysis of the literature in this area, using the elements of equations 4.1 and 4.3 as structure. Thus, **section 4.1** discusses ways of modeling the probability of failure, **section 4.2** gives an overview of the different methods to model the descent trajectory after a failure, in **section 4.3**, methods are discussed to model the impact area, and in **section 4.4**, the expected number of fatalities is given.

4.1. Modeling the probability of failure of a drone

The field of UAS Risk Analysis owes its relevance to the fact that drones fail every once in a while. To do a good risk analysis, one should know the probability of failure in any given situation. The infancy of the UAS field causes innovations to follow each other in quick succession and there is not yet much failure data available on UA. Therefore, it is difficult to obtain an accurate probability of failure. This section elaborates on how this probability can be accounted for in a ground risk assessment, given the aforementioned difficulties.

This section is structured as follows. In subsection 4.1.1, different failure modes and their influence on the modeling of failure are discussed. In subsection 4.1.2, methods of dissecting the system into components are described, which are derived from the Space industry. Bayesian Belief Networks have gained popularity in recent years to model the probability of failure and are presented in subsection 4.1.3. Other models just rely on the opinion of experts, which is discussed in subsection 4.1.4. And some risk models do not consider the probability of failure at all, which is discussed in subsection 4.1.5. Lastly, reflections on the different modeling techniques and a recommendation on which method to use in which scenario, is given in subsection 4.1.6.

4.1.1. Considering different failure modes

A starting point could be the definition of possible failure modes. In 2015, Clothier et al. established three failure modes for UAS operations in populated areas [18], and based on guidelines set out by regulators, Clothier et al. added a fourth failure mode to this list in 2018 [20]:

1. **Unpremeditated descent scenario:** A failure which results in the inability of the UA to maintain a safe altitude above the surface or distance from objects and structures
2. **Loss of Control:** a failure (or combination of failures), which results in loss of control of the RPA and may lead to impact at high velocity

3. **Controlled flight into terrain:** when an airworthy RPA is flown, under the control of a qualified remote pilot or certified autopilot system, unintentionally into terrain (water, structures, or obstacles)

4. **Dropped or jettisoned components:** failures that result in a component of the RPA (including its payload or stores) being dropped or jettisoned from the RPA.

Even though these different failure modes are recognized, a survey of 17 failure models used for UAS ground risk showed that only two adopted more than one failure mode [137]. It is nevertheless important to distinguish the type of failure mode being discussed, as an unpremeditated descent is much more likely to result in a safe landing than a loss of control.

Furthermore, the review of 17 failure models showed that all models assume the probability of failure to be constant throughout the flight, regardless of factors like flight phase, environmental conditions. From general aviation, it is known that some flight phases are inherently more dangerous, with for example almost half of all accidents happening shortly after takeoff or before landing [62]. One could think of incorporating failure data per flight phase from General Aviation into a UAS failure model, however, no literature was found that employed this method. One can thus conclude that even though this assumption has a significant impact on the model's performance, it is a very common one in the field.

4.1.2. Decomposition of systems

UAS operations is not the only field where risk assessments have to be made using only a little reliability data and a lot of uncertainty. In the past decades, the Space industry worked around that problem by doing structural and/or functional decompositions of systems, in order to analyze their overall reliability by combining the reliability of the parts. Methods used are for example a Fault Tree Analysis, in which the effect on the overall system caused by a single failing subsystem is deducted using a boolean tree [83]. The advantage of such a method is that a more reliable measure for system reliability can be derived if the reliability of all subsystems is available.

4.1.3. Bayesian Belief Networks

In recent years, much has been published on the application of Bayesian Belief Networks (BBNs) to the UAS safety assessment problem [6, 138, 139]. BBNs are graphical structures that make use of probabilistic reasoning to ascertain information about and are beneficial when expert judgment is ambiguous, incomplete, or uncertain. While other methods based on structural- and/or functional decompositions, such as Fault Tree Analysis, are deductive methods, BBNs are abductive, which means that they attempt to find the most likely cause of a failure. The research on BBNs deems this method advantageous to others for numerous reasons, including the ability to combine objective and subjective factors and the ability to make predictions including incomplete data [76].

4.1.4. Expert opinion

In the absence of data, experts come up with numbers themselves for the failure rate of UA. Although it is not said that the numbers are much more accurate than wild guesses, some works of literature opted for this method. Often, a Mean Time Between Failure (MTBF) is given. Values used in literature are given in Table 4.1. It is obvious that this number varies largely between different "experts", with one expert citing a 300 times higher probability of "non-catastrophic failure" than compared to the probability cited by another expert for "hazardous failure".

Source	**MTBF (hours)**	**Circumstances**
Clothier et al., 2007 [17]	10^5	-
Ford et al., 2010 [40]	10^4 10^5	hazardous failure catastrophic failure
Stevenson et al., 2015 [124]	10^5 10^6	sub-urban urban-areas
La-Cour Harbo, 2019 [79]	36 - 530	non-catastrophic failures

Table 4.1: MTBF used in UAS Risk Assessments

4.1.5. Modeling with non-realistic probabilities of failure

While previous sections described different methods used for modeling the probability of failure, some research on UAS risk assessments uses non-realistic values. This is not an error by the researchers in question, but a well-defendable choice. If the objective of a study is to compare the risk in different scenarios or ConOps, the probability of failure does not influence the comparison of the scenarios. In such a study, the probability of failure could theoretically be neglected at all. Alternatively, as some researchers indeed do, data from General Aviation can be used [79]. Other researchers only use relative failure rates between different scenarios (for example when comparing a ballistic descent scenario with a parachute-descent scenario) [106].

4.1.6. Reflections on the probability of failure

In this section, several methods have been proposed for modeling the probability of failure in drones. Although different methods exist, none is very accurate since the field is both rapidly improving and because little reliability data is available. When choosing a probability of failure model, one should ask whether the goal of the simulation is to find absolute risk levels or to compare different ConOps or modeling techniques. In the former case, the absolute risk level will only be as accurate as the inputs of the models, so one should start with a state-of-the-art method, which is thought to be BBNs. However, if one is interested in comparing models, it does not matter if the used failure rate is an order of magnitude higher or lower. In such scenarios, it is perfectly acceptable to take a short-cut and assume a value, for example from General Aviation, or use an expert-value, such as the ones described in Table 4.1.

4.2. Modeling the descent trajectory after failure

When failure occurs during UAS operation, it is the descent that separates it from a (potentially fatally) collision. To accurately model the risk of fatality, it is necessary to have an adequate model of this descent trajectory. When analyzing literature surveys such as [137] and recent work in which risk of Urban UAS operations is analyzed, such as [79, 106], it becomes clear that five descent models are considered predominantly. Which descent model is appropriate for a given model depends on the ConOps and the vehicle chosen. It is also not uncommon to consider different descent models, where different failure events (such as described in section 4.1) leads to a different descent trajectory.

The models considered in this section are the Ballistic Descent (subsection 4.2.1), Uncontrolled Glide (subsection 4.2.2), Fly-away (subsection 4.2.4), parachute descent (subsection 4.2.3) and a (fully) dynamic descent (subsection 4.2.5). This section is concluded by a comparison of the aforementioned models in subsection 4.2.6.

4.2.1. Ballistic descent

In a ballistic descent, the only forces acting on the UA are gravity and drag force. A Ballistic Descent model is not new to aviation. In 1996 and 2004, it was already implemented in the search of the crash locations of flights Trans World Airlines flight 800 and China Airlines CI611 [45, 94]. In more recent years, numerous studies related to UAS risk analysis in urban environments have implemented the model, as well. Examples of this are implementations by la Cour-Harbo and Primatesta [79, 80, 106].

Mathematically, the ballistic descent can be written as in Equation 4.4. In this equation, c is a coefficient resembling the drag acting upon the vehicle, which often consists of a drag coefficient multiplied with the frontal surface area of the vehicle. Furthermore, this equation can be extended with probabilistic uncertainties on initial horizontal- and vertical velocity, drag coefficient, frontal surface area, initial position, and wind, to arrive with a Probability Density Function of the possible crash locations of the UA.

$$m\dot{\boldsymbol{v}} = mg - c|\boldsymbol{v}|\boldsymbol{v} \tag{4.4}$$

Note that Equation 4.4 is a second-order model with no apparent analytical solutions. The trajectory of the ballistic descent can be modeled iteratively, however, as the trajectory will have to be sampled often to arrive at a PDF, and often PDFs are required for a large number of possible crash locations, this is not the most attractive solution in terms of computational efficiency. La Cour-Harbo has derived an analytical solution to Equation 4.4 to increase computational efficiency, but as this derivation spans 30 equations over 10 pages, it is not included in this literature survey [22].

4.2.2. Uncontrolled Glide

In literature, an uncontrolled glide descent for fixed-wing aircraft is considered the event where all thrust is lost, as well as control over the flight control surfaces. For a (multi-)rotor-aircraft, loss of thrust on the main rotor is considered which results in an auto-piloted auto-rotation descent. Even though the last scenario is not really "uncontrolled", it is considered to be part of the "Uncontrolled Glide" scenario in relevant literature such as [106], and therefore this standard in literature is adhered to.

For a fixed-wing, the trajectory is mainly influenced by the glide ratio, the wind and the configuration in which the control surfaces are "stuck". A simplifying assumption is that the control surfaces are in a neutral position. In such a scenario, the aircraft glides with glide angle γ, and the horizontal distance is simply $x(y) = \gamma y$ with a time to ground of $t(y) = x(y)/v_g$. Wind variations can be added to obtain the PDF.

4.2.3. Parachute Descent

A parachute that deploys whenever a critical failure occurs to the UAS, was already proposed as a recovery mechanism in 2001 by [144]. In recent years, its dynamics have been more accurately modeled for UA operations by [79, 99, 106].

The models for a Parachute descent assume that thrust is turned off, that it takes a short delay of t_d seconds to deploy the parachute, and that the horizontal velocity is reduced to 0 seconds w.r.t the wind. While the first two assumptions are perfectly reasonable in real-life operations, the latter is not. However, the distance traveled during deployment and deceleration is deemed to be small compared to the potential landing area of a parachute descent trajectory. A second-order drag model leads to a drop time t_{drop} from altitude y as per Equation 4.5. In this equation m is mass, g is the gravitational constant, A_p is the parachute area and $C_{d,p}$ is the parachute drag coefficient. As in other descent models, several input values can take probabilistic values, which will result in a PDF of the possible descent locations.

$$t_{\text{drop}} = \frac{y}{v_{\text{drop}}} = y\sqrt{\frac{A_{\text{p}} C_{\text{d,p}}}{2mg}} \tag{4.5}$$

4.2.4. Flyaway

A flyaway scenario results from a complete loss of control of the UA by the operator, while the autopilot continues to operate a stable flight. A flyaway model introduced by [79] consists of two effects.

The first effect is the probability of ground impact, which is considered to be decreasing linearly with distance from the event point. For example, one may argue that the flyaway will continue until the vehicle runs out of energy, and in such a scenario, the crash location will lie at a circle from the event location with the radius being the vehicle's range at the time of the event. On the other hand, it could be that the movements of the vehicle are erratic, which makes it more likely for the ground impact to be closer to the event location. Because that both scenarios are feasible, the simple assumption is taken of a linearly decreasing distance from the event point.

Mathematically, this decreasing probability is modeled as Equation 4.6, where θ is the wind average direction, R_{max} is the range at the moment of the event, v_c is the cruise airspeed, and $\mathbf{p} = (p_N, p_E)$ is a north-east position relative to the event point.

$$f(\mathbf{p}) = \max[0, \underbrace{R_{\text{max}} - \|\mathbf{p}\|}_{\text{linear decrease}} + \underbrace{\cos\left(\arctan\frac{p_{\text{N}}}{p_{\text{E}}} + \theta\right)\frac{\|\mathbf{p}\|}{v_{\text{c}}}}_{\text{modification according to wind}}] \tag{4.6}$$

The second effect is the probability that the vehicle flies a climb- or descent profile from the event location. Because of this effect, the crash location could be closer to the event location. Therefore, this is modeled as a normal distribution around the event location with a mean of 0 and an arbitrarily chosen s.d. of σ_{va}. It is modelled as Equation 4.7.

$$g(\mathbf{p}) = \frac{1}{2\pi\sigma_{\text{va}}^2}\exp\left(-\frac{\|\mathbf{p}\|^2}{2\sigma_{\text{va}}^2}\right) \tag{4.7}$$

The flyaway model then proposes a linear combination of 4.6 and 4.7 with relation factor β, giving each effect an equal contribution with $\beta = 0.5$. The resulting PDF becomes as in Equation 4.8. It is noteworthy that this PDF will be most likely to be widespread across a map.

$$P_{\text{flyaway}}(\mathbf{p}) = (1-\beta)\frac{f(\mathbf{p})}{\iint_{-\infty}^{\infty} f(\mathbf{p})} + \beta g(\mathbf{p}) \tag{4.8}$$

4.2.5. Full dynamic model

The fifth category of possible descent models is a full, 6DOF, dynamic model in which descent trajectory is calculated based on the cause of the failure. For example, a loss of thrust could lead to a situation in which the control surfaces are still usable, and the other way around. Such a model is too complex to fit in any of the assumptions taken in the previous four descent models, and its characteristics differ too much per situation to give a standard formulation.

Even though a fully dynamic model can lead to the most accurate solution, the computational power needed is often beyond the scope of Risk Assessment studies.

4.2.6. Reflections on Descent Models

In this section, five descent models were discussed that are used in literature on Urban UAS risk assessments. In Table 4.2, an overview is given how these models have been used in several UAS-related works of literature in the past decade. Furthermore, their applicability to rotor vehicles and fixed-wing vehicles is discussed.

Fixed-wing aircraft are more prone to enter an uncontrolled glide, rather than a ballistic descent, provided that their wings remain intact. Much work has focused on these fixed-wing vehicles, which is visible from the 9 papers that are referenced in the table. Although [79] mentions the possibility of auto-piloted auto-rotation in rotor vehicles to be part of the category "uncontrolled glide", this seems not the most likely of scenarios. Furthermore, it is the only work of literature that considers such a scenario.

Concerning rotor vehicles, a ballistic descent is more likely. Some literature indeed analyses this possibility. One may ask why Table 4.2 provides less literature on ballistic descent than on uncontrolled glide. This is likely caused by the fact that research into fixed-wing UA has been around for longer, while research into urban risk assessments where rotor-vehicles are predominantly used, has only taken off recently. This view is supported by the fact that all works on ballistic descents in rotor vehicles stem from recent years (since 2017).

A fully dynamic model takes away some uncertainties on the flight trajectory, but the behavior of the descent still depends on how the mode of failure. As this is probabilistic, there is still some uncertainty present. A parachute descent is a promising recovery mode, as the impact velocity and thus the impact energy can be greatly reduced. Nevertheless, commercial vehicles are often not (yet) equipped with such a system.

All in all, for the risk analysis of a rotor UA, the most appropriate manner of analyzing the descent is a ballistic descent, as it is more likely to occur in real life than an uncontrolled glide (=auto-piloted auto-rotation descent), is easier to implement than a fully dynamic model, and is likelier to occur than a parachute descent (and also more conservative). The descent model can be extended to use probabilistic input values. This would especially appropriate for the wind and drag-coefficient, as both would vary in a real-life descent trajectory based on changing winds and a changing orientation of the vehicle.

	References in UAS literature	Applicability to rotor vehicles	Applicability to fixed-wing vehicles
Ballistic descent	[5, 22, 79, 106]	+	+/-
Uncontrolled glide	[1, 5, 7, 8, 17, 46, 79, 90, 106]	-	+
Parachute descent	[79, 99, 106, 144]	+	+
Flyaway	[5, 79, 106]	+	+
Full dynamic model	[48, 83, 114, 143]	+	+

Table 4.2: Comparison of different descent models and references to their application in literature on UAS risk models.

4.3. Describing the impact area

The previous section described different methods for modeling the descent trajectory from failure to impact. This section will describe how the impact area is modeled.

The purpose of modeling the impact area is to establish the consequence of the given impact. How many people are hit, and how severely are they hit? For this, risk maps are often used. These maps describe the

risk for each discrete point on the map. In literature, these maps often consist of (some of) the following four components:

4.3.1. Impact area

An important parameter in the ground risk model is the area exposed to risk caused by the crash of a UAS. Melnyk et al. [90] categorized work on the impact area in two categories:

- Hypothetical prediction models (geometry based)
- Empirical prediction models (weight, size, or category based)

The hypothetical prediction models are then subdivided into three models: planform, gliding, and vertical descent models. The planform approach assumes the impact area to be a rectangle with the length of the aircraft as length, and its wingspan as width. The gliding approach also assumes a rectangle with the wingspan as width but uses for its length the distance which the aircraft glides from the height of a man to the ground. The last hypothetical prediction model is the vertical descent model, which uses a circle based on the wingspan of the aircraft. Melnyk et al. showed that the approach for the impact area is closely related to the descent profile used [90].

In the category of empirical prediction models, use is made of historical data to predict the impact area. This data is a limited resource in the UAS industry, because of its infancy, and therefore, data from General Aviation is often used. Melnyk et al. concluded that even the models used in the General Aviation scene often underpredict the impact area when compared to accident reports [90]. So not only is the leap from a big Boeing to a quadcopter so big that it is doubtful whether relationships between weight or size can say anything on the impact area, but also is the relationship itself already poor for General Aviation accidents.

Following the discussion of the different kinds of impact area prediction models, a literature survey was conducted to investigate what approach is used by relevant papers on UAS TPR. A study by Washington et al. from 2017 analyzed 25 different UAS risk assessments and concluded that 19 of those used a gliding prediction model (often combined with a vertical prediction model). Four used a hypothetical prediction model but not a gliding model, and two used an empirical model.

Even though 2017 is not that long ago, since then a lot has happened in terms of UAS research. The methods of assessing the impact area of three relevant papers published since 2017's literature survey are included in Table 4.3.

Source	Type of vehicle	Type of descent	Modeling
Ancel et al., 2017 [1]	Rotor	Ballistic descent	Glide, Equation 4.9
Primatesta et al., 2018 [104]	Fixed-wing	Uncontrolled glide	Glide, Equation 4.11
La Cour-Harbo, 2019 [79]	Fixed-wing	Uncontrolled glide	frontal area
		1. fly-away	1. $25cm^2$
		2. ballistic descent	2. $200cm^2$
Primatesta et al., 2020 [106]	Rotor & fixed-wing	Four types of descent	Glide, Equation 4.10

Table 4.3: Impact area models in UAS TPR literature since 2017.

Table 4.3 refers to equations 4.9, 4.10 and 4.11, which describe the impact area prediction using the gliding model corresponding to their respective sources in literature (as cited in Table 4.3). In these equations, w_{span} is the wing span, R_{uav} is the UAV radius, R_p is the radius of a person, L_{uav} is the length of the UAV, H_p is the height of a person and θ and γ is the glide angle. Although these formulas are used in state-of-the-art literature, they are not new themselves, but were introduced in earlier literature [19, 88].

However, what are the (dis)advantages of each of these methods? The difference between Equation 4.10 and Equation 4.11 follows from the fact that the first formula was originally developed for analyzing the impact area following a crash of space debris. Typical to such crashes is that they are at large angles, often nearly perpendicular to the ground. Therefore, the impact area was modeled as the "shadow" of the person impacted on the ground under the impact angle. However, at low impact angles, this shadow becomes very long, while it is unlikely that all people in this shadow will be at risk because the impact energy is greatly reduced after the first impact. To correct for this, Primatesta et al. have developed Equation 4.11 [104].

Graphically, the difference between the two models is visible in 4.1(a). Here, the original model, used for space missions, is seen casting a (long) shadow on the ground on the left. The new model is shown on the right, which considers the projection of the man-cylinder to the plane normal to the speed of the vehicle. A numerical comparison for impact angles between 0 and 90 degrees is given in 4.1(b). It is seen that Equation 4.10 indeed gives a very high impact area for low angles, while both equations are similar for angles above 40 degrees.

These comparisons indicate that Equation 4.11 is in general better suitable to UAV crashes, because it better handles impacts at low angles. Then why did Primatesta et al. did not use this relation in their 2020 work? The likely explanation for this is that they only modeled high impact angles in this work, for which both equations are similar [106].

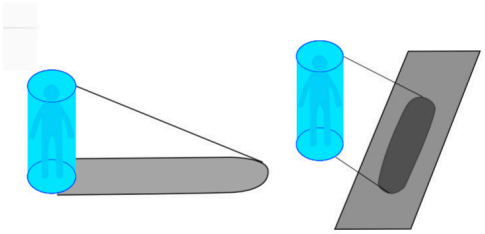

(a) Graphical comparison of two impact area models. Left is Equation 4.10, and right is Equation 4.11 [104].

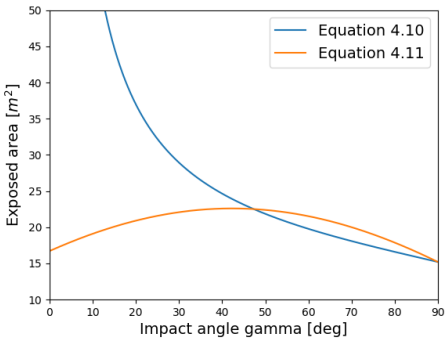

(b) Comparison of equations 4.10 and 4.11 for varying impact angles. Used parameters are $R_p = 0.2m$, $H_p = 1.8m$, $R_{uav} = 2m$.

Figure 4.1: Comparison of glide models described by equations 4.10 and 4.11.

$$A_c(\gamma) = \left(w_{\text{span}} + 2R_p\right)\left(L_{\text{uav}} + \frac{H_p}{\tan\gamma} + 2R_p\right) \tag{4.9}$$

$$A_{\text{exp}}(\theta) = 2\left(R_{\text{p}} + R_{\text{uav}}\right)\frac{H_{\text{p}}}{\tan(\theta)} + \pi\left(r_{uav} + R_{\text{p}}\right)^2 \tag{4.10}$$

$$A_{\text{exp}}(\gamma) = \pi\left(R_p + R_{\text{uav}}\right)^2\sin(\gamma) + 2\left(R_p + R_{\text{uav}}\right)\left(h_p + R_{\text{uav}}\right)\cos(\gamma) \tag{4.11}$$

It can thus be concluded that although there are many ways of modeling the impact area, two methods dominate the field. The first being an assumption based on the geometry (often the frontal impact area) of the UAV, and the other one is a hypothetical prediction model based on a gliding descent which takes into account the distance traveled descending from the height of the human head to the ground. Both models are relatively low-fidelity, in that they, for example, ignore the fact that a crash can have a secondary impact due to debris falling around. However, when the goal is not to derive an absolute risk value, but to come up with a relative value, either model is likely to provide enough accuracy, because the risk is likely to be relatively constant in the small area of a crash. (As when this would not be the case, a slightly smaller or larger impact area could lead to significant, non-linear, changes in risk).

4.3.2. Population Density

The risk upon impact at a location (x,y) very much depends on the number of people at that location. This is a factor that is considered in virtually every UAS risk analysis. A survey in 2017 by Washington et al. found that many UAS risk models assume a uniform population density, which is confirmed by analyzing earlier models [25, 137, 141]. Still, in recent years, more and more works are basing the population density on where people live [46, 105, 106]. This data is freely accessible in large parts of the world. For example, the Census Bureau in the US and the Kadaster in The Netherlands are government organizations thah provide this data on open platforms.

A limitation of this approach to population density is that it only accounts for where people live, and not where they move. Although at the time of writing this literature study, the better part of the world is home-quarantined due to the Corona-virus, under normal circumstances, people are not necessarily where they live. Melnyk et al. proposed to incorporate data on where people spend their time, such as in Table 4.4. However, still, it is not known where people are when they spend their time in the "outdoors", or where their office is located [90]. Obtaining accurate data on where people actually are is hard to obtain, and therefore, no work on UAS risk assessments has been found where anything else was used than Census-like data for modeling population density.

Information	Percentage
Time spent in Residence	68.7
Time spent Outdoors	7.6
Time spent in Vehicle	5.5
Time spent in Office/Factory	5.4
Time spent Indoors (other)	12.8

Table 4.4: Population behavior pattern data as published by Melnyk et al. [90].

4.3.3. Sheltering

Whether or not people are sheltered from impacts by objects like buildings or trees makes a very significant impact on the possible lethality of an impact, especially considering the urban area, where a lot of sheltering objects are present. Using a sheltering factor for risk assessments of UAS was made popular by Dalamagkidis in 2008 [24, 25].

Since the introduction of the sheltering factor, it has been used in recent works on risk analysis for UAS operations in Urban Environments. In such works, information on sheltering objects is typically retrieved from Open Street Map (OSM), which is an open-source map of the entire world. The advantage of this open-source approach is that it is easily reproducible, both by other researchers and for other cities. Two recent examples of research works that used OSM for a sheltering map are [104, 106].

4.4. Quantifying the risk of fatalities

Determining the probability of fatality, given a certain impact, is a question that researchers have been working on for half a century. The first describes the foundation of this field of research, as many of the state-of-the-art methods are derived from this original method. The second section describes the Blunt Criterion, and the third section described another quantitative method derived by Dalamagkidis et al. The fourth section describes a risk-assessment method based on a qualitative step-by-step analysis of the ConOps, which is used by EU lawmakers and called SORA. The last section elaborates on research that compared SORA to a High Fidelity Risk model.

4.4.1. US Army Personnel Casualty Study

The foundation for the research on the probability of fatality upon an impact was laid by Feinstein et al. in 1968, when they used injury data from the US Department of Defense to relate fragment weight and velocity with injury severity.

In their work, Feinstein et al. derived a log-normal distribution to fit the data relating impact to probability of fatality. This relation is given in Equation 4.12, in which K is the kinetic energy of the fragment, α_i and β_i are the scale and shape parameters of the log-normal. These two parameters depend on the part of the body that is impacted [38].

$$P(\text{ fatality } |K) = \int_0^K \frac{1}{x\beta_i\sqrt{2\pi}} e^{-\frac{(\ln x - \ln \alpha_i)^2}{2\beta_i^2}} dx \tag{4.12}$$

Having different parameters for different body-parts is problematic, as it is often not known a priori which part of the body will be hit during a crash. Therefore, Feinstein et al. derived an average curve based on the probability that each body part would be hit.

4.4.2. Blunt Criterion

The Blunt Criterion has been around since the 70s when it was introduced by [16]. Originally, it was an invention of the US Department of Defense which they used to analyze the result of blunt impacts to the chest. In 2009, Raymond et al. included head injuries in the model, and in 2013, it was first applied to UAS operations [108, 109].

The Blunt Criterion maps kinetic energy to injury severity, and its score can be converted to the AIS scale. It is defined as:

$$\mathrm{BC} = \ln \frac{E}{W^{1/3} T D} \tag{4.13}$$

In this equation, E is the kinetic energy, W is the mass of the impacted object, T is the thickness of the body wall (in cm) and D is the diameter of the impacting object (in cm). Sturdivan et al. shown that $T = kw^{1/3}$ with k = 0.71 for the average US male and k = 0.59 for the average US female test-subject in [126].

In [11], Bir et al. compared the BC to experimental data with velocities between 16 and 61 m/s. Using this experimental data, they related the BC to the Abbreviated Injury Scale (or AIS), which is a discrete scale ranging from 0 - 6 to categorize injuries ranging from "none" to "virtually unsurvivable" [43]. This relation is a linear fit with $R^2 = 0.78$, and is as follows:

$$AIS = 1.328BC + 0.603 \tag{4.14}$$

Hazard	Injury severity	Threshold
Blunt trauma / Crash injury	Casualty	15 J
Blunt trauma / Crash injury	Fatality	34 J

Table 4.5: Vulnerability threshold for blunt traumas [44].

It is important to note that the Blunt Criterion is derived based on missile impacts, which are fast, blunt impacts. It, therefore, does not account for energy dissipation in the impacted body. In impacts of drones, which are usually not blunt and (hopefully) impact at lower speeds than missiles, energy dissipation plays a larger role, which makes the Blunt Criterion overly conservative in such an application.

4.4.3. Kinetic energy at impact including sheltering factor

In 2008, Dalamagkidis et al. provided a novel relation between the probability of fatality and the impact, which includes a sheltering factor [24]. Such a sheltering factor accounts for the fact that humans are often protected by buildings, cars, or other structures. The relation is similar to other human vulnerability models as it is also a logistic relation. After improving the model to improve estimates at low energies, especially those close to the 34J threshold, the model is as in Equation 4.15, where k is as in Equation 4.16.

$$P(\text{ fatality} \mid \text{exposure }) = \frac{1-k}{1-2k+\sqrt{\frac{\alpha}{\beta}}\left[\frac{\beta}{E_{\text{imp}}}\right]^{\frac{3}{p_s}}} \tag{4.15}$$

$$k = \min\left[1, \left(\frac{\beta}{E_{\text{imp}}}\right)^{\frac{3}{p_s}}\right] \tag{4.16}$$

In Equation 4.15, the sheltering factor p_s determines how exposed the population is to an impact, taking values (0, inf), with higher values implying a better protection. α is the energy required for a fatality probability of 50% when $p_s = 6$, and the model approaches other models when $\alpha = 32kJ$. Furthermore, β is the impact required to cause a fatality when $p_s \to 0$, which is 34J as shown in Table 4.5.

The advantages of this model are not only the inclusion of a sheltering factor, but it has also been used in recent literature that is relevant for the field of UAS Risk operations in urban environments, most notably in [106]. Also, Dalamagkidis et al. indicate that this model is to be preferred over the Blunt Criterion for analyses at low speeds [25]. However, it is still based on Feinstein's model, which is derived from data of missile impacts, so in the absence of more relevant impact data, it is doubtful that this method performs much better than the BC on low-speed impacts.

To get a feeling for the order of magnitude of results produced using the human vulnerability model proposed in Equation 4.15, a simulation is done in which the sheltering factor is varied between 0 and 10, while the impact energy E_{imp} is chosen as $1e2$, $1e3$ or $1e4$. The result, which is shown in Figure 4.2, shows that the probability of fatality impacts primarily on the sheltering factor. It is thus imperative that it is chosen carefully. In recent work on UAS risk maps in urban environments, the approach described in this sector was applied using the sheltering factors as given in Table 4.6. Comparing these values with the probabilities shown in Figure 4.2, it is clear that sheltering factors above 10 are not used, because these values will rarely lead to a fatality.

Sheltering	Area
0	No obstacles
2.5	Sparse trees
5	Vehicles and low buildings
7.5	High buildings
10	Industrial buildings

Table 4.6: Sheltering factors as used in [106].

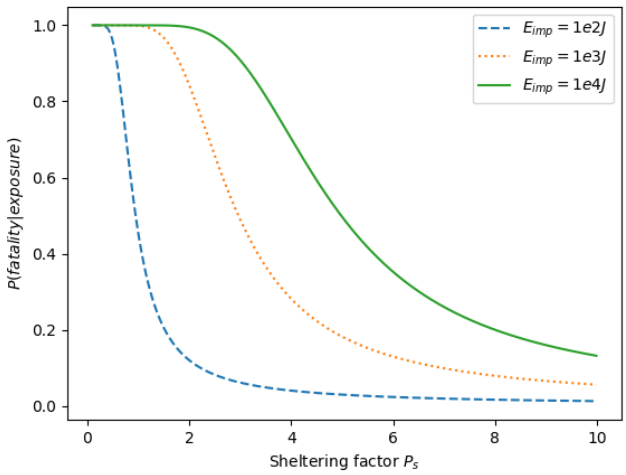

Figure 4.2: Fatality rate for different sheltering factors and impact energies using Equation 4.15.

4.4.4. Specific Operations Risk Assessment

The Specific Operations Risk Assessment, or SORA, is a guideline for assessing the risk of UAS operations. After JARUS published the first iteration of this document in June 2017, the final version was published on the 30th of January, 2019 [71]. As described in section 2.1, JARUS separates UAS operations into three categories: Open, Specific and Certified. SORA presented in this section only applies to the Specific category.

The SORA approach comes down to a step-by-step analysis for a proposed UAS ConOps, which then leads into a *SAIL value* (Specific Assurance and Integrity Level). This SAIL value then determines the required capabilities and mitigations needed for the ConOps to be approved. It is thus a *qualitative* risk assessment. Furthermore, SORA distinguishes between "ground risk" and "air risk", with the former being the risk of impacting with people or property on the ground, and the latter being the risk of impacting with another aircraft.

In order to use SORA, a 10-step process has to be followed. In the first of these steps, the ground- and air risk classes (GRC and ARC) are determined using specifics of the ConOps, such as whether flights will be operated above densely populated areas or in areas where there is a lot of air traffic. These risk classes are then tweaked using mitigation factors that have been taken, such as measures for reducing the effects of a ground impact. The GRC and ARC, together with their mitigation factors, are then combined into a SAIL value

(Specific Assurance and Integrity Level). The SAIL level dictates requirements to which the ConOps should adhere. The full 10-step process is presented in chapter 8.

All in all, the SORA provides an approach that is not overly complex to use for a wide range of UAS ConOps. In fact, Terkildsen et al. have already created an online SORA questionnaire in which UAS operators can assess the compliance of their ConOps with the JARUS Sora guidelines [128].

This simplicity, combined with the fact that it is adopted by the EU regulators, makes it a very appropriate tool for assessing the risk of UAS ConOps. However, an inherent downside to a discretization method like this, is that comparing the difference in risk between two ConOps is more difficult.

4.4.5. Comparison of SORA with a High Fidelity Risk model

The risk assessment described in the previous section (subsection 4.4.4) is based on a qualitative analysis of the ConOps to arrive at a risk score. La Cour-Harbo compared SORA with a High Fidelity Risk Model (HFRM) [23]. The author implemented the HFRM according to models which he described in two separate papers, which were published simultaneously [79, 81]. The SORA method was implemented based on proposals of the method published in 2016 and 2017 [30, 31], which were largely similar to the final version published in 2019 [71].

To ensure a meaningful result to the comparison, different scenarios were analyzed using both models. Those scenarios had in common that they were all based on an 80KM BVLOS flight operation, but were differed in terms of parachute availability, type of overflown area and type of aircraft, leading to $2^3 = 8$ different scenarios. The two different values for the three parameters are shown in Table 4.7. The *partially urban* scenario flies for 20% over an urban area, while the other 80% flies over the same rural area as the second option.

Parameter	Option 1	Option 2
Type of aircraft	HEF32 20kg rotorcraft	Cumulus One 2kg fixed-wing
Overflown area	Partly urban	Rural
Parachute	Yes	No

Table 4.7: Values for different scenarios used in comparison between SORA and HFRM [23].

The two models are compared on the fatality rate. The philosophy regarding the fatality rate of the two methods is substantially different, however. SORA uses a target fatality rate and determines operational procedures, mitigations, and additional requirements to reduce the fatality rate below the target. It thus starts at the fatality rate to arrive at the specifics of the mission. With HFRM, this process is reversed, as the specifics of the mission are used to calculate the fatality rate. The comparison was performed on the basis that when the HFRM leads to a probability of fatality (POF) of less than 10^{-6}, the flight is allowed. It is then relevant to know if the same flight would also be allowed under SORA.

The results are that both methods are largely in agreement with each other. The flights over rural areas have a POF in the order of magnitude of 10^{-7}, and have a SAIL score of II, which means that both models deem the scenario very acceptable. The flights that partially overfly an urban area arrive at a SAIL of between IV and VI, while the HFRM determines a POF of around 10^{-6}. Thus, according to both models, such a scenario is on the intersection of being allowed or disallowed. To extend the comparison, the authors also analyzed a fully urban flight, which resulted in a POF of 10^{-5}, and also the SORA method deemed such a flight as too risky to undertake.

One may wonder why such different methods, produce similar results. A probable explanation is that the SORA classifications are based on similar HFRM as used by La Cour-Harbo. While the analysis did not include a sensitivity analysis, it is likely that subtle changes in the performance of the UA, the trajectory, or the modeling of the descent and/or impact, would lead to significant changes in the POF. It is fair to conclude that both risk assessment approaches lead to the same order of magnitude, BUT concluding that they are equal alternatives would be too strong of a conclusion to be justified by the experiments performed.

4.4.6. Considerations for calculating velocity at impact

In several models, reference is made to the kinetic energy at impact. This can be established in three ways, which are described in this section. Kinetic energy is calculated as in Equation 4.17. While the mass of a UA can (usually) safely be assumed constant, the velocity at impact varies, and often varies significantly in the

trajectory leading to the crash, as this is already an off-nominal and uncontrolled trajectory.

$$E_{imp} = m v_{imp}^2 \tag{4.17}$$

Therefore, the question is how to determine the velocity at impact. This can be done in three ways. Firstly, the terminal velocity can be used, which is the velocity at which the gravity- and drag-forces are in equilibrium. The terminal velocity is to be calculated as in Equation 4.18 as derived in [92]. In this, A is the frontal area, ρ is the gas density and C_D is the drag coefficient. The application of terminal velocity in quantifying risk of fatality in UAS operations is proposed in [24, 25], however, drawbacks are that it is often very conservative and that the frontal area and drag coefficient are dependant on the orientation of the UA, and therefore subject to variation.

$$E_{imp} = \frac{m^2 g}{\rho_\alpha A C_d} \tag{4.18}$$

The second approach is derived from regular airplanes, where regulators in the EU [28] and the US [49] propose to take the velocity that is 40% above maximum operating speed as the velocity at impact. This method is also used by [47]. The third approach is simulating the velocity at impact using a physical model of the UA.

Complementary to the approaches of calculating the velocity at impact, Sturdvian et al. proposed that after impact on the ground, both the UA and the person impacted continue to move and that the impact should thus be modelled using conservation of momentum [125]. This can be done according to Equation 4.19

$$E = \frac{1}{2} m_1 v_1^2 \left(1 - \frac{m_1}{m_1 + m_2}\right) \tag{4.19}$$

4.4.7. Reflection on calculating fatality at impact

In this section, three methods of calculating the probability of fatality in UA were presented. The Blunt Criterion, presented in subsection 4.4.2, is used in several works on UAS operations, however, it is overly conservative for impacts on low speeds. Dalamagkidis et al. derived a logistic relation which is based on Feinstein's curve (presented in subsection 4.4.1), and includes a sheltering factor. Lastly, EU lawmakers linked a qualitative analysis of the ConOps to the probability of fatality, which was presented in subsection 4.4.4.

Which of the presented methods would be most suitable for analyzing urban UAS operations? That question can be answered by looking at three important factors. Firstly, one needs to have a quantitative method to be able to properly compare different approaches. Secondly, being able to take sheltering into account is important, as a large part of the urban population is shielded from impacts by some sort of sheltering. Thirdly, the ability of the model to work at low speeds is helpful, as UA impacts need not to be at high speeds.

Considering these three criteria, the logistic relation derived by Dalamagkidis et al. is the best suitable for analyzing the risk of UAS in urban environments. Another confirmation that this is an appropriate method, is that recent papers on the risk of urban UAS use this method. The most recent example is an article by Primatesta, La-Cour Harbo et al., published in 2020 [106].

4.5. Conclusions

In this chapter, four essential elements of a TPR assessment of an urban UAS ConOps have been thoroughly reviewed. In section 4.1, different methods have been discussed regarding the probability of failure. It was found that while this is a challenging field in UAS research due to the lack of reliability data, methods like FTA and especially BBN are being researched extensively to arrive at more accurate models. Still, many works on UAS research use "expert opinions", which could as well be interpreted as "ballpark" values. When the goal is to compare different risk models, an uncertain failure probability will not influence the result.

Regarding **descent trajectories**, section 4.2 presented five methods of modeling the descent after a failure has occurred. However, for rotor-vehicles, a ballistic descent is the most appropriate scenario. In the descent trajectory, there are several uncertainties, such as the orientation of the vehicle (which determines the drag coefficient) and the wind. Taking probabilistic values for these inputs, a Monte Carlo simulation can be used to arrive at a PDF of the possible descent locations.

After descent, one is concerned by **the area impacted by a crash**. Section 4.3 discussed three important parameters: the size of the area impacted by the crash, the population density at that area, and the sheltering

provided in that area. The area impacted by the crash is often modeled by either assuming a frontal surface area or assuming a gliding impact in which the distance needed to travel from the height of the average human to the ground is the area impacted. The former is most relevant for ballistic impacts, while the latter is more relevant when a gliding scenario is considered. The population density is often modeled by using Census data, which has a shortcoming that it only tells where people live and not where they are. Whether or not they are sheltered from an impact by for example buildings, can be derived from mapping software like Google Maps or Open Street Map.

The probability of fatality, given an impact, was discussed in section 4.4. Two quantitative methods that are predominantly used in the field are the Blunt Criterion and a relation found by Dalamagkidis et al. (equations 4.13 and 4.15). Both are based on missile impact data from the US Department of Defense, but the latter formula is adapted to account for possible sheltering factors, which is a clear advantage when analyzing UAS. Next to quantitative methods, JARUS developed SORA, a 10-step qualitative method to assess the risk based on a UAS ConOps. A comparison with a high-fidelity risk model showed that this is an accurate method, and it is also accepted by EU lawmakers. However, when comparing the risk of different ConOps, a quantitative method is to be preferred. Therefore, Dalamagkidis' logistic relation is the best choice when comparing the risk of different UAS ConOps.

For a drone delivery ConOps in urban areas using a rotor-vehicle, it is thus concluded that **the failure probability** can be taken as a ballpark "expert"-opinion. The descent trajectory can be modeled as a ballistic descent, employing a MC-simulation to arrive at a PDF of descent locations based on a stochastic drag-coefficient and stochastic wind. The size of the impact area can be modeled using the surface area of the drone, while the population density can be taken from Census data, and sheltering factors can be derived from open-source maps. Dalamagkidis' logistic relation (as in Equation 4.15) is a suitable way of modeling the probability of failure. Taking these steps leads to an accurate representation of the TPR, which is in line with methods applied by recent literature in the field.

5

Multi-Agent Path Finding

The Multi-Agent Path Finding problem (MAPF) considers a graph G with vertices V and edges E, such that G = (V,E). On this graph, agents need to travel from their start position $s_i \in V$ to their goal position $g_i \in V$, without conflicts. Conflicts are defined as two agents occupying the same vertex at the same time. Time is discretized into time steps. Finding the optimal solution to a MAPF problem is NP-hard, as the state space grows exponentially with the number of agents.

This chapter will commence by presenting two popular Path Finding algorithms, namely the A* algorithm in section 5.1 and the Conflict Based Search algorithm (CBS) in section 5.2. Subsequently, Risk Based Path Finding (RBPF) will be presented in section 5.3.

5.1. A* algorithm

The original A* algorithm was introduced by Peter Hart et al. in 1968 as part of a research project focused on building a robot that could plan its own actions. The main advantage of this algorithm is that it is guaranteed to find a solution if this solution exists, and that it is guaranteed that this solution is the optimal solution [50].

The disadvantage is that the state space scales exponentially, with its complexity being $O(b^k)$, where b is the branching factor, or the number of adjacent nodes per node, and k being the number of agents. Consider a 4-connected grid and 20 agents. The branching factor is thus 5, as the possible actions are movement (in 4 directions) or waiting at the current node. With 20 agents, the state space would be $5^{20} = 9.53x10^{14}$ which is computationally infeasible [119].

The procedure employed by the A* algorithm is explained in the following. This explanation is based on a blog post by Nicholas Swift, which gives a clear understanding of the procedure. The algorithm uses an "open"- and a "closed"-list, where the former includes all nodes that are yet to be analyzed, and the latter includes all nodes that are already analyzed. Both lists are empty at the start of the algorithm [50, 127].

For each algorithm, $F(x) = G(x) + H(X)$ is calculated. G(x) is the distance between the node and the start node.H(X) is a heuristic that estimates the distance from the current node to the goal node. This heuristic must be *admissible*, or in other words, that it does not overestimate the distance from the node to the goal. Mathematically, this implies that $h(x_n) \leq h*(x_n)$ for each node x_n, with $h*(x_n)$ being the lowest cost of reaching the goal from x_n. Heuristics used are the Euclidian- or the Manhattan-distance. F(X) is then the total cost of the node.

1. Add the starting node to the open list
2. Repeat the following steps:

   I Take the node with lowest cost F on the open list as current node.

   II For each of the adjacent nodes of the current node, do one of the following:

   - If it is already on the closed list or if it is not accessible, ignore it.
   - If it is not on the open list, add it to it, and note costs F, G and H.
   - If it is already on the open list, check if the path from the current node to the new node has a lower cost G, than the cost G currently associated with this node. If so, change the parent node of that node to the current node.

III As the current node is now fully analyzed, move it to the closed list.

3. Stop when either the goal node is added to the closed list, or when the open list is empty and the goal was not found. In the former case, the solution is found, and in the latter case, there is no path.

5.1.1. Improvements on A*

Since the introduction of A*, several improvements to it have been proposed. These algorithms can exponentially reduce the computational cost of A*, or have little influence at all, depending on the problem at hand. This section briefly introduces two of these improvements, independence detection (ID) and Enhanced Partial Expansion (EPEA*)

Independence Detection works by exploiting the fact that when agents have no chance of ever crossing paths, it is unnecessary to plan them simultaneously. The algorithm works by first planning individual paths for all agents, and subsequently re-planning the paths for agents that have conflicting individual paths. As computational cost increases exponentially with the number of agents, this can lead to a significant reduction of computational cost if many individual paths happen to be independent [122].

The second improvement is EPEA*, which exploits the fact that the A* generates a new node for all of its children, even though some are much worse than others. By applying heuristic- and domain-specific knowledge, it only adds those children to the open list, which will lead to an optimal solution. This avoids the generation of surplus nodes [42].

5.2. Conflict Based Search

This section discusses Conflict Based Search (CBS). This is a two-level algorithm that solves an optimal solution to the pathfinding problem. At the high level, it operates a Constraint Tree (CT), in which each node represents a set of constraints on the paths that agents are allowed to take. At the low level, it uses fast single-agent solvers to solve the paths that satisfy the constraints at each node. This section analyzes the (dis)advantages of CBS, and presents an empirical comparison with other path-finding algorithms.

Constraints in the constraint tree are tuples (a_i, v, t), where agent i is prohibited from entering vertex v at time t, and conflicts are tuples (a_i, a_j, v, t), where agents i and j are conflicting at vertex v at time t. Each node N in the Constraint Tree consists of the following:

- **A set of constraints:** The root of the CT contains no constraints, and each child node both inherits all nodes of the parent and adds one new constraint for one agent.
- **A solution:** A set of k paths for k agents, in which the path of agent i is consistent with the constraints of that agent.
- **The total cost:** The sum of the costs of all individual paths at this node.

For each node in the CT, a low-level search is performed in which the paths are found that are consistent with the constraints at that node. These paths are subsequently validated with all other paths found for this node, to find potential conflicts. If no conflicts are found, this node is declared a *goal node*, and this solution is returned. However, if conflicts are found, this node is declared a *non goal node*.

If a node is declared a non-goal-node, it is thus found that agents i and j occupy the same vertex v at the same time t. This thus implies that either agent i or agent j should be constrained from including this (v,t) in its path. This constraint should thus be added to the list of constraints in a new node. In order to ensure optimality, a new (child) node is created for both agents i and j, including respectively (a_i, v, t) or (a_j, v, t) as a constraint.

It should be noted that the structure of the CT allows for significant performance enhancements. These include that at each node, only the new constraint needs to be saved, instead of all constraints, as the list of constraints can be obtained by going up the tree. Also, in the low-level search, only the path for agent a_i that has a new constraint needs to be calculated, as all other paths remain the same.

Empirical evaluation

This section presents an empirical evaluation of CBS, A*, and EPEA* on an 8x8 4-connected open grid, with the number of agents (k) ranging between 3 and 21. In the simulation, the algorithms were run for a time limit of 5 minutes, noting a *fail* whenever a solution was not found within this time limit. The results, which are presented in Figure 5.1, show that A* is clearly inferior to EPEA*, with CBS having a slight advantage over the other [119], when one considers a high number of agents.

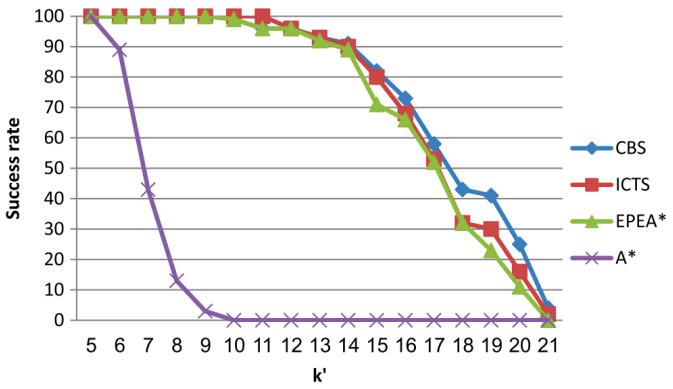

Figure 5.1: Evaluation of A*, CBS, EPEA* and ICTS (not discussed here) [119].

5.3. Risk-Based Path Finding (RBPF)

So far, this chapter has presented different methods of finding the shortest path. This section will elaborate on the situation in which not the shortest path, but the path with the least risk is to be found. This implies that the goal is to travel along a graph G = (V,E) such that agents travel from their start position, $s_i \in V$, to their goal position, $g_i \in V$, while minimizing the cost. In this problem, the cost is a metric for the *risk* of traveling across an edge.

Using Path-Finding in adversarial environments has been employed on many applications: such as operations on Mars, in Nuclear Power Plants, and for minimizing radar exposure while flying over enemy territory [72]. In general, there are two phases to this problem. The first is defining the cost of each edge, which is described in the following section, and the second is making adjustments to MAPF-algorithms to work with this new cost structure, which is described in the section after that.

5.3.1. Defining the Cost

The cost $E_{i,j}$ of traveling between (connected) nodes S_i and S_j should represent the risk. There is a multitude of ways to do this. Typically, the risk is the expected number of fatalities, which is established by identifying possible crash locations and multiplying the probability of crashing in each of those locations by the number of people at risk at that area. There are many considerations in establishing this risk measure, and these are elaborated upon at length in chapter 4.

An important distinction with the methods described in that chapter, is that for a Risk analysis of a path flown, only crash locations along this path have to be sampled. However, when one wants to analyze the risk for RBPF, the risk along *all possible paths* has to be sampled. Sampling the risk is already a costly endeavor, as the descent trajectory after failure is stochastic (Due to for example the failure mode and the presence of wind, for an elaboration on this, see section 4.2). When one has to sample the risk for all possible positions in the environment, with a range of possible initial velocities for each position, this becomes even more computationally expensive.

Because of this, it makes sense to do a more rudimentary risk analysis during Path Finding. For example, Rudnick-Cohen et al. derive a risk distribution using a Monte Carlo simulation of crashes with random initial position and speed, where less MC iterations are performed for each position-speed combination during RBPF, than during the actual risk analysis [112, 113]. Primatesta et al. simply account for a uniform distribution of crash locations in a circle with radius R from the initial position, where R is the maximum distance that can be flown from the crash location [104].

It can thus be concluded that when defining the Risk cost of traveling across the environment, in general, the same methods are applied as for when assessing the risk on a path that is actually flown. However, often, a more rudimentary analysis is done during RBPF to avoid large computational costs.

5.3.2. Modifications to Path-Planning algorithms to include Risk

In the previous section, it has been explained how the Path Finding problem can be redefined to minimize the risk, rather than the length of the path. In principle, Path Finding algorithms such as A* (as presented in

section 5.1) can be employed on this modified problem definition. However, some practicalities have to be taken into account. This section presents these practicalities for using the A*-algorithm for the RBPF problem as described in the previous section.

As described earlier, the A*-algorithm only gives an optimal solution whenever the heuristic is *admissible*. Commonly, the A* algorithm uses the Euclidian- or Manhattan distance between the current node and the goal node as a heuristic, as paths can never be shorter than that (depending on whether diagonal moves are allowed). However, the risk-cost of moving between nodes can be lower than one, depending on the way the risk is modeled. Primatesta et al. published a modified heuristic for RBPF using A*, also referred to as RA* (RiskA*) [105]. In this work, a risk-cost r_c with $r_c \in [0,1]$ is assumed and $\text{dist}(x_n, x_{n+1})$ is defined as being the Euclidian distance between the points. The modified heuristic is given in Equation 5.1.

$$\begin{aligned} h(x_\text{n}) = &\frac{r_\text{c}(x_\text{n}) + r_\text{c}\,\text{min}}{2}\,\text{dist}(x_\text{n}, x_{\text{n}+1}) + \text{dist}\left(x_{\text{n}+1}, x_{\text{goal}-1}\right) r_\text{c}\,\text{min} + \\ &+ \frac{r_\text{c}\left(x_\text{goal}\right) + r_\text{c}\,\text{min}}{2}\,\text{dist}\left(x_{\text{goal}-1}, x_\text{goal}\right) \end{aligned} \tag{5.1}$$

However, this heuristic can be simplified if the distance between all nodes is equal, thus, $\text{dist}(x_\text{n}, x_{\text{n}+1}) = \text{dist}\left(x_{\text{goal}-1}, x_\text{goal}\right) = \text{dist}_\text{min}$. The heuristic then becomes as in Equation 5.2.

$$h(x_\text{n}) = \frac{r_\text{c}(x_\text{n}) + r_\text{c}\left(x_\text{goal}\right)}{2}\text{dist}_\text{min} + \left(\text{dist}\left(x_\text{n}, x_\text{goal}\right) - \text{dist}_\text{min}\right) r_\text{c}\,\text{min} \tag{5.2}$$

Furthermore, it can be the case that one wants to minimize both the Risk and the length of the path. In such a scenario, one can extend the cost function $F(x_n)$ associated with each node in the A* algorithm with a cost $G(x_n)$, which is the moving cost, and a cost $R(x_n)$, which is the risk of moving to the same node. One could then use a scale parameter α with $\alpha \in [0,1]$ to balance the two effects as in Equation 5.3 [26].

$$F = \alpha G + (1-\alpha)R + H \tag{5.3}$$

5.4. Conclusions

In this chapter, the field of Multi-Agent Path Finding (MAPF) was discussed, and specifically how the field can be extended to minimize Risk, rather than the length of a path, using Risk Based Path Finding (RBPF). For planning individual paths of agents, the A* algorithm is an efficient solution, however, if one needs to plan multiple agents simultaneously to create collision-free paths, extensions to A* such as ID, or other algorithms such as CBS are more suitable.

In the RBPF problem, a primary concern is establishing the risk associated with traveling between two nodes, as this is a computationally expensive process. It is seen that authors of literature on risk for UAS applications often use a less precise, but more efficient approach for estimating this factor. Regular path-finding algorithms such as A* can be used on the RBPF problem, however, the heuristic usually should be tweaked to guarantee optimality. The A* algorithm with a heuristic based on risk, is also referred to as the RiskA* algorithm.

6

Conclusion

In chapters 2-5, a literature survey was conducted on (urban) UAS Operations. Research gaps have been identified that could be starting points for future research, and essential components were described that are needed for such research projects. This chapter summarizes the most important conclusions following from the literature survey.

UAS Operations

Current regulations for Urban UAS Operations restrict beyond-VLOS traffic. However, regulatory bodies are working intensively on new regulations, and it is likely that in the future, BVLOS traffic over densely populated areas will be allowed. This prompts the analysis of ConOps that would operate in such a scenario.

When analyzing such ConOps, deliveries of small packages through drones is a good application due to its relatively low complexity and high expected traffic density. Furthermore, it is not far from reality, since companies such as Amazon and UPS are conducting fly-tests with vehicles for such applications [85, 97, 130].

An important part of a ConOps is the UTM architecture, because it specifies the interactions between the actors in the system. UTM architectures are proposed in which a distinction is made between the Air Navigation Service Provider, the UAS operator and its systems, and third parties such as people on the ground. Essential to such a UTM architecture is whether UAS (operators) are allowed to conduct flight-plans without having them pre-approved by ANSP. Even though this matter is still being debated, it was concluded that it makes the most sense if UAS operators are indeed allowed to do this.

Lastly, an evaluation of possible ConOps and relevant literature showed that Agent-Based Modeling is a suitable paradigm for modeling UAS Operations.

Airspace Structures

On this topic, it was concluded that the current airspace structure, in which the airspace is categorized in Classes A-G, is not suitable for Urban UAS operations. Instead, it is likely that the urban airspace structure will consist of a custom network made of tubes (3D-roads), layers (specific traffic or headings per altitude band) and zones (no-fly zones, or specific zones for specific traffic). Furthermore, three research projects were discussed that analyzed the influence of airspace structures on (urban) traffic management.

Hoekstra et al. showed that an airspace structured with tubes results in the least complex traffic situation, and reached the more general conclusion that the performance of an airspace is very sensitive not only to the type of airspace structures used but also to the topology of these structures [58]. Also, work by Salleh et al. showed that in extremely high traffic densities, the less structure, the higher the capacity, but also concluded that a structured airspace in general leads an urban airspace that is safer and easier to manage. Furthermore, it proposed *adaptive airspace structures*, in which the airspace structure changes based on simple metrics such as the traffic density [9, 10]. Following these results, two research gaps were identified:

1. Little is known about the influence of airspace structures on TPR, and on the influence of the topology of such airspace structures on TPR.

2. *Adaptive airspace structures* are proposed for managing the capacity of an airspace, but can they also be applied to minimize TPR of an airspace?

Furthermore, Chung et al. analyzed whether agents that adjust the cost of traveling over an edge dynamically based on only local information can reduce the travel time in a congested network. In their research, agents used evolutionary algorithms to learn a successful policy. Their results indicated that such an approach can indeed improve system performance on a global level, even if the policies are based on local information and are only applied locally [15]. This leads to the third research gap:

3. Can the result of Chung et al. be applied to the topology of airspace structures, such that local agents learn to adjust the topology in such a way that the global system performance is improved?

Risk

It is established that third-party risk is the most important factor to consider in risk assessments of urban UAS operations. Four building blocks have been discussed that together form a coherent and accurate TPR-asessment. These four building blocks and their respective conclusions are:

- **Failure rate:** The probability of failure of a UAS is hard to analyze due to the absence of reliability data. State-of-the-art methods that attempt to work around this problem are Bayesian Belief Networks. However, when one is not interested in an absolute risk value but merely wants to compare different ConOps, the absolute failure rate is not so important, and one can use guestimates.
- **Descent trajectories:** One needs an accurate description of the trajectory between failure and an inevitable crash. A literature survey was conducted summarizing the different descent models in recent third-party-risk research, concluding that a ballistic descent is the most appropriate scenario when analyzing a multi-rotor vehicle.
- **Modeling the impact area:** A risk map can be used to analyze the impact at a certain location. Typical elements of such a map include the population density and a layer containing sheltering-factors, which indicate whether the population in that area is sheltered from an impact by structures such as buildings or trees. Furthermore, the size of the impact area can be estimated using geometrical methods based on the dimensions of the vehicle.
- **Expected fatalities given an impact:** Two quantitative methods used for this are the Blunt Criterion, and a relation found by Dalamagkidis et al. Another assessment is SORA, which analyzes the risk of a ConOps using a 10-step qualitative assessment. Although the latter method is adopted by EU lawmakers, a quantitative method is better suited for a comparison involving different scenarios. Of the two quantitative methods, the one proposed by Dalamagkidis is most appropriate for risk models of Urban UAS, because it includes a correction for people sheltered by buildings or other structures [25].

Multi-agent Path Finding

Two algorithms were presented that are commonly used in the field of Path Finding. These methods are the *A* algorithm* and *Conflict Based Search*. For planning individual paths of agents, the former is a good method, however, if one needs to plan collision-free paths for multiple agents simultaneously, algorithms such as CBS are more suitable.

In Risk-Based Path Finding, the goal is not to find the shortest path, but the path with the least risk. Existing algorithms can regularly be implemented to the RBPF problem by changing the cost structure of a graph. This cost structure should represent the risk of traveling along each edge and can be established using methods described in the chapter on Risk. Sometimes, slight modifications are made to MAPF algorithms for RBPF. For example, RiskA* is a variant of A*, in which the heuristic is changed to be compatible with the risk-based problem. This method is also recommended for minimizing the risk of individual paths.

7

Research Plan

Building on the conclusions of the literature study, a research plan is established that aims to add value to the academic field of UAS Risk in urban areas through an MSc Thesis. In this chapter, the first section describes the research objective of this project. In the second section, a work plan is made which should aid in the process of fulfilling the research objective.

7.1. Research Objective

Large scale urban UAS operations are not yet permitted. One of the primary reasons for this is the risk that Urban UAS operations would pose to the large number of people living, working, and traveling through a city. A scenario in which people not involved in the operations of a UAS (also called *third parties*) would suffer a fatal accident because of a crashing UAS should be avoided at all costs.

A significant amount of work has been done on finding ways to make UAS safer by reducing their rate of failure, and by creating mechanisms that reduce crash severity, such as deploying parachutes and airbags [144]. However, the field of operations can also contribute to this effort by contributions in two areas: Firstly, by developing models that allow regulators to better understand the TPR following from UAS Operations, and secondly, by inventing new ways of mitigating the TPR through smart operational procedures.

Researching the influence of operational concepts on UAS safety is the focus of this MSc thesis. Following the analysis provided in chapter 2, it is chosen to analyze drone delivery operations, as these operations are expected to be one of the first applications to be introduced in the urban environment, and their operations are relatively simple to model.

When developing methods to enhance safety of these operations, not only should one create models that have good results in simulations, one should also look at methods that are feasible in the real world. For example, accurate risk-based path-finding algorithms (such as described in chapter 5) use a high-fidelity representation of the risk, in which full information on the whereabouts of people is needed. However, it is not feasible to constantly monitor, collect and distribute such data, and even if it were possible, it would not be desirable to share such data with UAS Operators as Google and Amazon due to privacy concerns.

Several researchers, including Hoekstra et al. and Salleh et al. have proposed structuring the Urban Airspace with elements such as tubes, layers and no-fly zones in order to reduce complexity [9, 10, 56–59]. However, it has not yet been analyzed how those structures influence the TPR compared to an unstructured airspace. Also, the development of an Agent-Based Model of Urban UAS operations serves as a foundation of not only this research objective, but of many research objectives in the field of (Urban) UAS Safety analysis, and is therefore seen as a valuable addition to the academic field. Both the development of this ABM and studying the effect of no-fly zones and tubes on this model are the focus of this MSc thesis. This is defined in the following research objective:

> The objective is to model the Third Party Risk of drone delivery operations in urban areas, and to investigate how the urban airspace can be structured using tubes and/or no-fly zones in order to minimize third-party-risk following from a drone delivery ConOps.

In order to work on this objective, it is split in several sub-objectives, such that when these sub-objectives are met, the full objective is met. These sub-objectives are:

1. **To create an agent-based model of package delivery using UAS in an urban environment**
   As established in chapter 2, package delivery is a suitable application of UAS in urban areas due to the relatively low complexity of the concept, and the high expected density compared to other applications. The same chapter also concluded that agent-based modeling is the most suitable paradigm for such a model. Within this sub-objective, the following objectives are identified:

   (a) **Create a 2D-environment that approximates a typical city or cities (including its buildings and population), such that it can be used to simulate UAS package-delivery operations.**
   In chapter 3, previous work on Urban UAS operations showed that the results are sensitive to the chosen city. However, in chapter 2, research is presented that mathematically showed that all cities in the world can be categorized in one of four categories, based on their geographical structure. By creating a model of cities belonging to different categories, it becomes more likely that the work of this MSc Thesis can be generalized to different cities.

   (b) **Create a scenario for the logistics of the operations in this environment**

      i. Investigate where to place the hubs and the delivery points.

      ii. Create a model that represents the expected demand for deliveries.

2. **Create a model that analyzes the Third Party Risk following from this ABM**
   After having implemented the delivery operations in an ABM, the TPR following from these operations should be analyzed. As discussed in chapter 4, the following four elements need to be included to obtain an accurate model of the TPR:

   (a) The probability of failure of a UAS
   (b) The descent trajectory after failure
   (c) A representation of the impact area
   (d) An quantification of the fatalities following the impact

3. **Analyze the effect of different airspace structures on the TPR**
   In chapter 3, it is shown that structuring the urban airspace with no-fly zones, tubes or layers could reduce complexity, and it is a common opinion that those structures will be essential building blocks of the Traffic Management of Urban UAS Operations (UTM). It was found that the influence of those structures on TPR is still uncharted territory. This objective will attempt to fill this void.

4. **Develop a method that minimizes the TPR by structuring the urban airspace**
   While the previous sub-objective investigated the relation between different airspace elements and the TPR, this sub-objective aims to find ways to reduce the TPR by using airspace structures. The relevancy of such a method would be that if successful, it implies that UAS operators would only need to include the airspace structures when planning, rather than having to consider a full representation of the whereabouts of the population in the environment. This would both be computationally more efficient, and more desirable regarding privacy.

   The success can be evaluated by comparing three scenarios: 1) a MAPF-algorithm such as A* that only attempts to find the shortest path, 2) a MAPF-algorithm that is extended to minimize the risk, such as RiskA*, and 3) a MAPF-algorithm such as A* that finds the shortest path, but respects the constraints set by the airspace structures. Sub-sub-objectives would include:

   (a) **Investigate whether simple heuristics can be used to place airspace structures that minimize TPR**
   Salleh et al. proposed simple criteria, such as the number of UAS in the air, that can be used to adapt the structure of the airspace dynamically to make a trade-off between throughput and complexity [10]. Could heuristics of similar simplicity be employed to place airspace structures in such a way that the TPR is reduced?

(b) **Investigate whether controller-agents can learn to adjust the airspace structures to minimize TPR**
Chung et al. proposed a traffic-management system where agents with local information adjust the cost of traveling along their corresponding edges, and showed that such a system can improve system-wide performance. It is interesting to see if a similar optimization approach can be applied to the problem of minimizing TPR through urban airspace structures.

7.2. Work Plan

Based on the research objectives set out for this MSc thesis, five work packages are defined that should lead to the fulfillment of the research objective. It should be noted that the progress of a project such as this, spanning several months, is hard to predict. Findings early on in the project will result in modifications and tweaks to other work-plans. Therefore, these work-plans are not to be considered as set in stone, but rather as a fluid document that will be subject to an iterative development process.

In WP1, assumptions scattered throughout the report will be summarized in an initial research scope. WP2 describes how an existing model of Urban UAS operations can be modified to a model that is suitable for this research, and WP3 mentions the additions that are needed for the TPR-assessment. Subsequently, WP4 implements several scenarios including airspace structures such as tubes and no-fly zones and analyzes the results of these scenarios. WP5 attempts to generalize the results of WP4 into methods that can generate such airspace structures automatically, and investigates whether these models can be improved by applying learning. It is noted that WP5 is rather a welcome addition to the research objective, rather than being the core of the research. Therefore, it is indicated as being optional and not included in the planning.

WP1: Setting the research scope

A good research scope is essential to the success of the following work packages. A wise man once said *"saying no to one thing means saying yes to something else"*. The same holds for the research scope. A prioritization decides which elements are most required to fulfill the research objective, and which ones can be ignored without sacrificing the added value to the academic field.

Within the field of Urban UAS, many components can be added to make the result as realistic as possible. For example, one can make a simulation either 2D or 3D, a fully dynamic model of a drone can be added, and one could consider one of the 50+ advanced CD&R concepts of which the applicability on UAS is evaluated by Jenie et al. [64, 65]. Throughout chapters 2-5, several considerations for such assumptions were presented. These assumptions are summarized in the remainder of this section.

Assumptions regarding third-party risk

The third-party risk-model is essential to the research. Therefore, a high-fidelity model will be used based on state-of-the-art methods, such as presented in chapter 4. Regarding risk, it is assumed that:

- Only third-party risk following an all-engines out failure resulting in a ballistic descent is considered. No other failure modes are considered, such as collisions or a-few-engines-out failure.
- A high-fidelity (static) risk-map, including population density and sheltering factors for urban structures.
- In a ballistic descent, neither control of the vehicle (such as an auto-rotation maneuver) or deployment of any risk-mitigating objects (such as parachutes) is possible
- The UAS fails with a constant failure rate based on expert values.
- A crash only results in a primary impact, there are no secondary impacts due to debris.

Assumptions regarding the operational concept

Creating a representative full dynamic model of a UA including its interactions with the environment and other vehicles is a field of research on its own. The operational concept should be realistic enough to derive findings that apply to the real world, but not more complex than that. Furthermore, the objective of this research is to find the relative differences between specific concepts of operations, not to find any absolute values that would be representative of reality. Therefore, regarding operations, the following is assumed:

- One or a few cities that are proven to have a topography that is similar to other cities.
- Package ($\leq 2.5kg$) delivery operations in which drones fly between delivery points and hubs.
- Multi-rotor vehicles with VTOL capabilities.
- Vehicles fly at the same speed and same altitude which is reached directly after take-off (So 2D).
- Low traffic densities are used, such that the effect of conflicts can be ignored.

The research scope will remain fluid throughout the project. It is likely to be broadened in some areas and narrowed in other areas following preliminary results. Nevertheless, the assumptions presented in this work-plan will prove a good foundation for initial model development.

WP2: Creating a working simulation of the Concept of Operations

The author of this proposal participated in the creation of a model of a drone delivery ConOps in the city of Delft, which is a good foundation for the research in the MSc thesis. A visual representation of the used model is given in Figure 7.1. At least the following changes should be made to this model:

- Use a bigger city to model, such as Amsterdam or New York, and create a model of that city.
- Generate locations of hubs, delivery-points, and demand for package delivery between them.
- Choose and model a vehicle with representative speed and mass for real-world operations.
- Verification, validation, and optimization of the computational efficiency of the model.

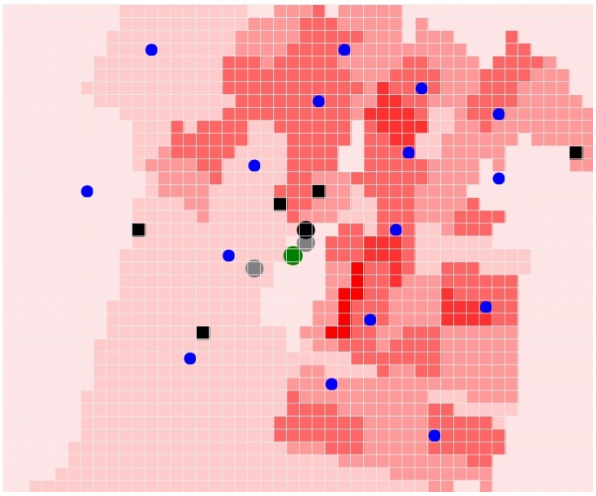

Figure 7.1: Setup of UAS delivery model created for AE4429 course [67].

WP3: Modeling the third-party risk

The model presented in WP2 already has some TPR capabilities. However, these should be expanded with at least the following steps:

- Include stochastic wind and drag in the ballistic descent model.
- Generate an initial risk-map using an MC simulation that can be used for risk-based path-finding.
- Optimize the heuristic of the risk-A* algorithm for risk-based path-finding.
- Pre-generate risk-maps for increased efficiency during large calculations.

Using this model, the following analyses are to be performed:

- Compare shortest-path A* and risk-A* using the following metrics:
  - Distance "as-the-crow-flies" vs traveled distance
  - Travel-time
  - Packages delivered
  - Total risk (and risk per delivery and/or risk per time)
- Perform a sensitivity analysis using the most important components.

WP4: Model airspace structures and their effect on TPR

The goal of this work-package is to implement the airspace structures, such as described in chapter 3 and in research objective 3. No-fly zones and tubes should be implemented in a range of scenarios and topologies, and the effect on the metrics mentioned under WP3 of each of these scenarios should be analyzed.

WP5: Generalize the results of WP4 to a formalized method (optional)

As mentioned before, WP5 is an addition rather than the core of the research project, and its relevancy depends on results following from earlier research. If executed, the goal of WP5 is to answer the objectives as explained in objective 4. The following steps can be taken:

- Based on the scenarios that performed well in WP4, determine to set up simple heuristics that would generate such airspace structures automatically.
- Investigate the possibility of applying multi-agent learning techniques to improve the found result.

An important note is that when evaluating the performance of the generated methods, this should be done on new scenarios to avoid over-fitting. The number of scenarios should be increased to the point that it is possible to derive some statistical assurances about the validity of the results. Therefore, it is likely that new scenarios should be created according to the steps described in WP1.

8

Appendix A: SORA

Specific Operations Risk Assessment, or SORA, is a qualitative way of assessing the risk of a UAS ConOps. It consissts of ten steps, which are presented in the following.

1. **Description of Concept of Operations (ConOps)**

2. **Determination of UAS Intrinsic Ground Risk Class**
   Based on the dimensions of the UA, its typical expected energy and the operational scenarios it will operate in, an Intrinsic UAS Ground Risk Class is attributed to the proposed scenario on a scale of 1-10. For the full table, please be referred to Table 8.2.

3. **Final Ground Risk Class determination** Based on mitigation factors, the intrinsic GRC proposed in the previous step is corrected. Possible mitigation factors include strategic mitigations for ground risk (such as risk-based path-planning), the reduction of the effect of a ground impact and whether a validated emergency response plan is in place. How exactly the GRC is adjusted for these mitigation factors can be found in table Table 8.3.

4. **Determination of Air Risk Class**
   The ARC is categorized in levels ARC-A, -B, -C or -D. For this, an extensive flow-chart is used. Factors such as operating height, whether or not the UAS operations take place in controlled airspace, and whether or not the operations take place in urban areas are determining factors. Generally, ARC-A is used for airspace where the risk of collission between UAS and manned aircraft is already extremely low without any mitigation, as is the case in operations in some parts of Alaska, for example. ARC-B, -C and -D are used for operations in airspace of increasing risk of collission between UA and manned a/c. For the full flowchart of determining the ARC, please be referred to the original SORA guide by JARUS in [71].

5. **Application of strategic mitigations to determine final ARC**
   The fifth step is used for tweaking the ARC based on any mitigations taken. Contrary to the mitigations for the GRC such as described in step three, no clear guide for judging mitigations for ARC is given. Whether or not mitigation factors for ARC are sufficient is left to the consideration of the regulatory body approving a specific ConOps. This process is described in the SORA documentation in [71].

6. **Tactical Mitigation Performance Requirement**
   The sixth step involves an analysis whether the TMPR is sufficiently capable and robust for the ARC determined in steps four and five. For increasing ARCs, these capabilities have to meet stronger requirements. DAA is an example of a functionality that is specified in this category.

7. **SAIL Determination** Based on the ARC and GRC found in the previous steps, the SAIL can be determined, using the conversion table in Table 8.1.

	Residual ARC			
Final GRC	**a**	**b**	**c**	**d**
<2	*I*	*II*	*IV*	*VI*
3	*II*	*II*	*IV*	*VI*
4	*III*	*III*	*IV*	*VI*
5	*IV*	*IV*	*IV*	*VI*
6	*V*	*V*	*V*	*VI*
7	*VI*	*VI*	*VI*	*VI*
>7	*Category C operation*			

Table 8.1: Determination of SAIL based on GRC and ARC, adapted from [71]

8. **Operational Safety Objectives**
   Based on the SAIL level determined in the previous step, an operator can determine which OSOs have to be taken into account. In total there are 24 OSOs, to which an operator should demonstrate no-, low-, medium- or high- compliance based on the SAIL factor. OSOs include for example controller certifications, the communication mechanism, and maintenance procedures [71].

9. **Adjacent Area/Airspace Considerations**
   Includes analyzing whether the infringement of adjacent airspace, resulting from a loss of control, changes the conclusions regarding the SAIL level.

10. **Comprehensive Safety Portfolio**
    The last step is a formal one, in which the findings from step 1-9 are combined into the Combined Safety Portfolio.

8.1. Determination of UAS Ground Risk Class

Intrinsic UAS Ground Risk Class				
Max UAS characteristics dimension	*1m*	*3m*	*8m*	*>8m*
Typical Kinetic Energy Expected	<700 J	<34 KJ	<1084 KJ	>1084 KJ
Operational Scenarios				
VLOS/BVLOS over controlled ground area	*1*	*2*	*3*	*4*
VLOS in sparsely populated environment	*2*	*3*	*4*	*5*
BVLOS in sparsely poplated environment		*4*	*5*	*6*
VLOS in populated environment	*4*	*5*	*6*	*8*
BVLOS in populated environment	*5*	*6*	*8*	*10*
VLOS over gathering of people	*7*	*n.a.*	*n.a.*	*n.a.*
BVLOS over gathering of people	*8*	*n.a.*	*n.a.*	*n.a.*

Table 8.2: Determination of Intrinsic Ground Risk Class, adapted from [71]

		Robustness		
Mitigation Sequence	**Mitigations for ground risk**	Low/None	Medium	High
1	M1 - Strategic mitigations for ground risk	*0: None* *-1: Low*	*-2*	*-4*
2	M2 - Effects of ground impact are reduced	*0*	*-1*	*-2*
3	M3 - An Emergency Response Plan (ERP) is in place, operator validated and effective	*1*	*0*	*-1*

Table 8.3: Incorporation of mitigation factors in Ground Risk Class, adapted from [71]

III

Supporting work

1

Model Elaboration

This chapter elaborates on the proposed model. Section 1.1 elaborates on all assumptions underlying the model. Furthermore, section 1.2 discusses the process of generating the environment. Subsequently, section 1.3 outlines the bottum-up estimation of the population density. Lastly, section 1.4 presents the standard values of all relevant model parameters.

1.1. Model Assumptions

The model is based on the following main assumptions:

Limited to Third-Party Risk: As regulators consider TPR to be the most relevant dimension of risk, this research project limits itself to it.

Limited to an all-engines-out-failure resulting in a ballistic descent: Many failure modes are conceivable, including partial systems failure or a loss of control, and these failure modes can result in a multitude of descent trajectories. However, this project is limited to an all-engines-out failure resulting in a ballistic descent.

Unrestricted airspace: It is assumed that the airspace is not constricted in any way by elements such as no-fly zones. It is assumed that UASs can fly any route they desire, without any need for approval.

No physical obstacles: It is assumed that there are no physical obstacles limiting the route of the UAS, such as buildings or trees.

Fixed cruise altitude: UASs fly at a constant cruise-altitude all throughout the flight.

No diagonal movements: It is assumed that vehicles fly only in the vertical or horizontal direction. This implies that the vehicle conducts a vertical take-off directly after the start of a mission, and conducts a vertical descent at the end of it.

No interactions between vehicles: No interactions at all between UASs are incorporated in the model. This also implies that UASs are never in conflict with each-other, and never take any measures to detect and avoid a conflict.

Homogeneous vehicles: It is assumed that all vehicles are identical, and perform identically.

Instant re-charging: It is assumed that vehicles recharge their batteries instantly upon arriving at hubs and delivery-points.

No weather-/wind-influence during the flight: Flight-plans are not adjusted for the presence of adverse weather and/or wind. It is also assumed that wind does not influence the ground speed of the vehicle. Wind

is only included in the crash-descent-trajectory, after a failure has occurred.

Constant probability of failure: Initially, it is assumed that the probability of failure is constant throughout the flight.

Constant mass: The mass of the vehicle is assumed to be constant, irrespective of whether or not it is carrying a package.

No recovery- and/or other risk mitigation measures: It is assumed that no measures mitigate the risk after a failure has occurred, such as parachute recovery systems or airbags. It is also assumed that the population density layer does not change after failure, which implies that people do not flee the crash location, after they have identified a crashing UAS.

Small packages: Only the delivery of small packages (<2.5kg) is assumed.

1.2. Generation of the environment

This section elaborates on the process followed to generate the environment. This process consists of three steps, leading to the maps in figures 1.1 - 1.3.

1. Obtaining the input data: For the data on population density, data provided by the government authorities of the respective countries was accessed ([14] for Delft, [2] for Paris, [134] for New York). This data is based on a census, in which the government surveys the amount of people living in each place. OpenStreetMap data was accessed as input data for the sheltering factors.

2. Cleaning the input data: The input data should be cleaned to allow for better processing in the next steps. Usually, images consist of hundreds or thousands unique RGB-values, because two different colors always fade into eachother. However, those in-between colors cannot be mapped to a discrete value for sheltering or population density. Using a tool [60], those in-between values are changed to disrete values, such that the figures only consist of a low amount of unique RGB-values. Also, the sheltering- and population-density map must be exactly overlayed, such that each pixel on one map corresponds with the same pixel-location on the other map. This process is done manually.

3. Mapping rgb-values to respectively density or sheltering values: Steps one and two have provided us with two maps that have unique, discrete RGB-values for all pixels, and of which the location of pixels exactly match. Now, unique colors can be matched to values for sheltering data or population density.

4. Modify low-/high-buildings: The sheltering data does not distinguish for low- or high buildings. Therefore, an additional map is generated, on which a black pixel indicates that a building is high, and a blue pixel indicates that a building is low. When pixel P is classified as a building by the sheltering map, this new map is then accessed to determine whether this buildings should be classified as high or low. Using Google Streetview, it was determined whether buildings are low- or high in practice. Only for Paris, it was determined that buildings all around the city are quite similar, and therefore, this step is skipped.

1.3. Bottum-up estimation of population density

Unfortunately, population density data is only provided for where people live, which does not reflect where people are. Therefore, a bottum-up estimation was made of where people spend their time during the day. Using such a method is inspired by Melnyk et al. [90], who estimted the average population density per location, irrespective of the time of the day. A serious limitation to this method is that there are likely significant differences in population density based on the time of the day, because behavior of people differs if it's rush hour our working hours. In our work, we decided to correct for this, by creating separate scenarios for office-time (9am - 5pm), rush-hour (7am - 9am and 5pm - 7pm), and night-time (7pm - 9am).

In total, our research categorizes people in one of the six categories mentioned below. Initially, according to the population density data, 100% of the people are in category one, at home.

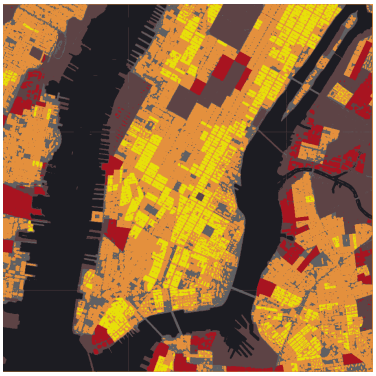

(a) Population Density map

(b) Sheltering Factor map

(c) Low- vs high-buildings (black=high, blue=low)

Figure 1.1: The environmental representation of New York.

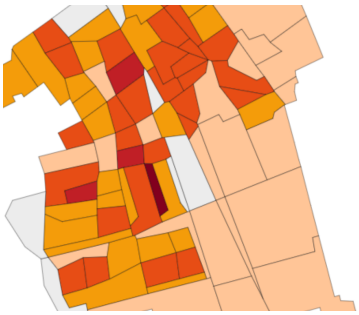

(a) Population Density map

(b) Sheltering Factor map

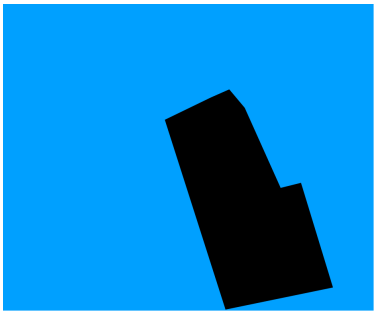

(c) Low- vs high-buildings (black=high, blue=low)

Figure 1.2: The environmental representation of Delft.

(a) Population Density map

(b) Sheltering Factor map

Figure 1.3: The environmental representation of Paris.

1. Home
2. Office
3. School
4. Outside
5. Transportation
6. Water

Based on various data-sources provided by the Dutch bureau of Statistics [13], the where-abouts of people have been estimated. Table 1.1 presents the where-abouts of the people who are registered as working. In

respectively tables 1.2, 1.3 and 1.4, the where-abouts of people during daytime, nighttime and rush-hour are calculated. It should be noted that albeit this data is much more advanced than the current state-of-the-art, which doesn't discriminate between phases of the day, it is still limited. For example, this data is assumed to be homogeneous for Paris, New York and Delft, while that certainly will not be the case. Also, all tables sum to 100%, while it is likely that the population of a city during the day is not the same as its population during the night. Furthermore, many of the individual numbers are assumptions and could be subject to critique. For example, it is estimated that people who travel to school, have a trip of the same duration as the people traveling to work, while these trips are often of entirely different characteristics and use different transportation mechanisms. Another limitation is that the data used, are averages for The Netherlands. This data is not representative of the population of Delft, as Delft has many more young people due its university.

ID	Category	Percentage
1A	Percentage of population holding a job	47.7%
1B	People only work only 29 hours per week on average (=73% of possible time)	73% * 1A = 34.6%
1C	10% works from home	10% * 1B = 3.5%
1D	10% works at night (= people working from the office during night-time)	10% * 1B = 3.5%
1E	15% works outside (=people working outside during daytime)	15% * 1B = 5.2%
1F	15% people who work on the road	15% * 1B = 5.2%
1G	People working at the office during daytime	1B - 1C - 1D - 1E - 1F = 16.4%

Table 1.1: Where-abouts of the 47.7% of people holding a job, during the day.

ID	Category	Percentage
2A	At the office	1G = 17.4%
2B	At school (based on CBS data)	16.3%
2C	Transportation	1F = 5.2%
2D	People outside	1E = 5.2%
2E	People on water	0%
2F	People at home	100% - 2A - 2B - 2C - 2D = 55.9%

Table 1.2: Where-abouts of people during day-time

ID	Category	Percentage
3A	At the office	1D = 3.5%
3B	At school (based on CBS data)	0%
3C	Transportation	15% * 1D ≈ 0%
3D	People outside	0%
3E	People on water	0%
3F	People at home	96.5%

Table 1.3: Where-abouts of people during night-time

ID	Category	Percentage
4A	At the office = 25% of people at the office	25% * 2A = 4.4%
4B	At school = 25 % of people at school	25% * 2B = 4%
4C	Transportation = 25% of people at school + at office	25% * (2A + 2B) = 16.8%
4D	People outside = 25% of people at school + at office	25% * (2A + 2B) = 16.8%
4E	People on water	0%
4F	People at home = all other people	100% - 3A - 3B - 3C - 3D = 58.0%

Table 1.4: Where-abouts of people during rush-hour

1.4. Standard values of parameters

In Table 1.5, the standard values of all relevant model parameters are specified. Unless specified otherwise, the model was set-up using these values for all of the experiments.

Variable name	Description	Standard value
add_value	Value for the heuristic in path-finding	0.001
air_density	density of the air	1.225 kg/m^3
gravitational_constant	Gravitational constant	9.81 m/s^2
cruise_altitude	Cruise altitude of the UAS agent	100 m
demand_per_hour	Number of tasks generated per hour	1000
max_steps	Number of model steps per simulation	28800 (8 hours)
scaling_factor	Factor with which the maps are scaled	5
num_agents	Number of UAS agents	20
failure_rate_cruise	Probability of failing in one step	0.0005
failure_rate_climb_descent	Probability of failing in one step	0.0005
city	City to model	Delft
n_random_hubs	Number of random hubs to generate	4
n_random_delivery_locations	Number of delivery points to generate	100
random_CD	Whether or not to generate random drag coefficient	False
include_wind	Whether or not to include (KNMI) wind data	False
modify_dens_scenario	Modify population density according to specified scenario	1
vertical_speed	Vertical speed of the drone	5 m/s
cruise_speed	Horizontal speed of the drone	10 m/s
frontal_area	Frontal area of the drone	0.04 m^2
mass	Drone mass	9 kg
α	Variable of the fatality model	32000 J
β	Variable of the fatality model	34 J

Table 1.5: Standard values of all model parameters.

2

Additional Experiments

This chapter presents three experiments that have not been discussed in the main paper.

2.1. Varying the cruise altitude

This experiment compares the scenario that the cruise altitude is 500 meters instead of 100 meters. It should be noted that it is doubtful whether operations at such a high altitude would be allowed. For example, US regulators indicate a maximum operating altitude of 500 feet (about 150 meters) for urban UAS traffic [35].

We hypothesise that operations at an altitude of 500 meters would be more dangerous, as crashing from such an altitude increases the kinetic energy upon impact, which increases the risk. Inspecting the results, which are presented in Figure 2.1 and Table 2.1 confirms this hypothesis, indicating that the fatalities per 1m flight-hours increases by 14%, with a p-value of $<10^{-10}$ that these scenarios are different.

Two effects cause this difference. Firstly, the risk is increased because of an increased altitude. Secondly, more time is spent in the Climb and Descent phase (C&D), where the risk is lower. This has a mitigating effect on the overall risk. Figure 2.1(a) shows that the fatalities per 1m flight-hours increase by 29.5% and 17.5% in the C&D- and cruise-phases, respectively, but only with 14.2% in the total simulation. At first, this may seem erroneous. However, despite the large increase in risk in the C&D-phase, it is still lower than the risk in the cruise-phase. In Table 2.1 it is noted that the time in the C&D-phase increases almost fivefold. Spending more time in a safer flight-phase, is the reason that the overall number of fatalities is still lower. Using Equation 2.1, it can be concluded that the numbers are indeed correct.

Furthermore, it can be concluded that the risk per delivery is 47% higher when the cruise altitude is 500 meters. Because both the flight time and the risk per flight-hour increases, the risk per delivery increases even more.

$$FAT-1M_{total} = \frac{FAT-1M_{cruise} * t_{cruise} + FAT-1M_{C\&D} * t_{C\&D}}{t_{total}} \tag{2.1}$$

Summarizing, it is concluded that flying at a higher altitude increases the risk per flight-hour by 14%, and the risk per delivery by 47%. Even though the primary driver for limiting the maximum altitude of UAS operations, is to avoid conflicts with General Aviation, this analysis adds another reason to the discussion. Nevertheless, flying at a higher altitude may open up additional possibilities for steering the UAS away from risky areas during the descent trajectory, following a failure, however, this effect is neglected in this analysis.

	Fat. per 1m flight-h [fat/h]			Total collective risk [fat]			Flighttime [seconds]		
Flight-phase	C&D	Cruise	Total	C&D	Cruise	Total	C&D	Cruise	Total
Altitude = 100m	4.6	5.7	5.8	0.05e-3	0.87e-3	0.92e-3	37e3	534e3	573e3
Altitude = 500m	5.9	6.9	6.6	0.25e-3	0.81e-3	1.06e-3	150e3	423e3	573e3

Table 2.1: Results of varying the cruise altitude on peformance indicators.

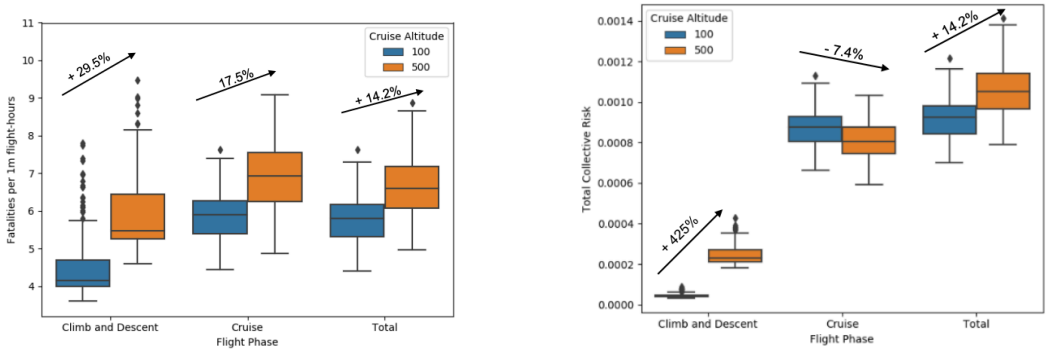

(a) Fatalities per 1m flight-hours per flight-phase.

(b) Total collective risk per flight-phase.

Figure 2.1: Change in fatalities per 1m flight-hours and total collective risk because of changing the cruise altitude.

2.2. Bi-objective optimization

Not only is the objective to minimise risk to third-parties, the objective is also to maximise the efficiency of the operations. This section investigates the hypothesis that an optimal solution in terms of safety is not optimal in efficiency. Even though the process of bi-objective opitimization is also discussed in the main paper, this section summarises this approach. Subsequently, we present the results of an experiment carried out in the New York environment.

In the model, this bi-objective implementation is implemented as a sub-optimal A* solver [50]. The cost function of the A* algorithm is in Equation 2.2, where $G(x, y)$ is the cost of moving from node x to y, and $H(x, y)$ is the heuristic associated with that movement. In this analysis, the heuristic is set in according to Equation 2.3, with dist(x,y) being the distance between nodes x and y.

$$F(x, y) = G(x, y) + H(x, y) \tag{2.2}$$

$$h(x, y) = \text{add\_value} * dist(x, y) \tag{2.3}$$

A sensitivity study is done in which *add_value* is varied for the city of New York. The hypothesis is that increasing add_value will increase both risk and efficiency. Studying the results, which are in Figure 2.2, this hypothesis is confirmed. The results take a logistic shape, with the largest variation visible between 0.001 and 1. Note that even though an add_value of 0 has been used in the experiments, however, this gave undesirable results, because the UAS sometimes got stuck in areas with zero risk. Based on this analysis, we selected an add_value of 0.001 as initialisation of the model.

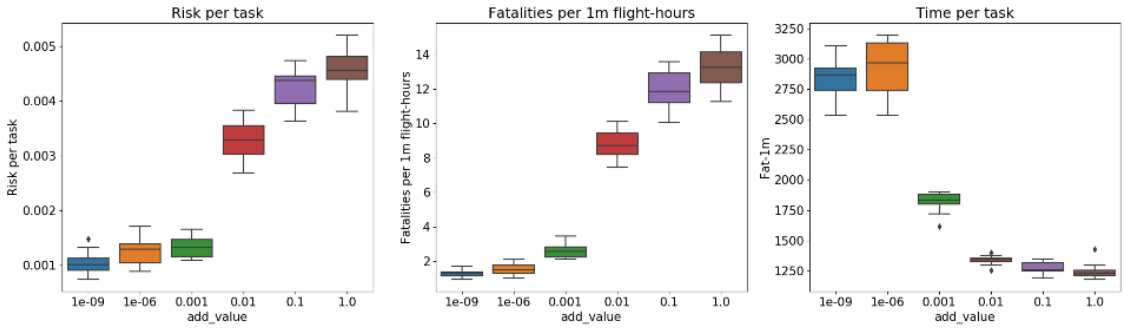

Figure 2.2: Effect on three parameters following from the variation of add_value.

2.3. Place Delivery Points according to population density

In the model, Delivery Points (DPs) are placed according to population density. This implies that there are more DPs in areas where the population density is higher. This section presents the experiment in which this scenario is compared with the scenario where DPs are placed randomly. As can be seen in fig. 2.3(a), there is a difference in FAT-1M of approximately 6%. However, when measuring the average risk around DPs, there is a difference of about 54% visible, as can be seen in fig. 2.3(b).

The absence of a difference in risk is thought to be caused by two reasons. Firstly, as presented in the paper, the climb&descent-phase only contributes to about 6.5% of the total risk. Taking into account that the hubs remain at the same position, the DPs only account for about 3% of the overall risk. Furthermore,

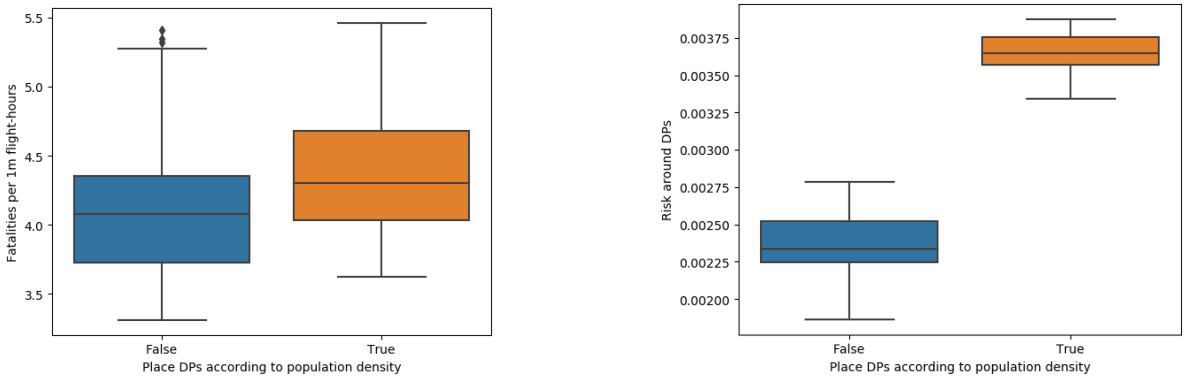

(a) Fatalities per 1m flight-hours per scenario.

(b) Average risk around DPs.

Figure 2.3: Change in fatalities per 1m flight-hours and total collective risk because of changing the placement of DPs.

2.4. Sensitivity study on threshold parameter of experiment 1

Experiment one analysed the influence of environmental parameters on model behaviour. Two of the five used parameters consider low-risk areas. These areas are defined as the areas in the T%-percentile. By default, T is set to 20%. In this section, the results are presented of for different values of T. These results yield slightly different results, but identical conclusions regarding the direction- and strength of the observed correlations. In Table 2.2, the results are presented using T = 10%. In Table 2.2, the results are presented using T = 25%.

Correlation coefficient (T = 10%)	y = fatalities per 1m flight-hours (FAT-1M)					
	Delft		New York		Paris	
x parameter	*conclusion*	*r-value*	*conclusion*	*r-value*	*conclusion*	*r-value*
avg_risk_env	↑	0.17	↑↑	0.30	–	0.02
ann_lowrisk	↓	-0.17	↓↓	-0.44	↓↓↓	-0.64
avg_risk_lowrisk	↑↑↑	0.82	↑↑↑	0.71	↑↑↑	0.81
avg_risk_DP	–	-0.03	–	0.06	–	0.00
avg_risk_hubs	↑	0.23	↑↑	0.44	–	0.17

Table 2.2: Spearman's rank correlation coefficient of environmental parameters with respect to number of fatalities per task. Arrow indicates direction and strength of correlation, and when an arrow is indicated, correlation is statistically significant after applying Bonferroni's correction. One arrow = weak correlation, two arrows = moderate correlation, three arrows = strong correlation. – implies no (statistically significant) correlation.

Correlation coefficient (T = 25%)	y = fatalities per 1m flight-hours (FAT-1M)					
	Delft		New York		Paris	
x parameter	*conclusion*	*r-value*	*conclusion*	*r-value*	*conclusion*	*r-value*
avg_risk_env	↑	0.44	↑↑	0.50	–	-0.02
ann_lowrisk	↓↓	-0.50	↓↓	-0.40	↓↓	-0.69
avg_risk_lowrisk	↑↑	0.44	↑↑↑	0.71	↑↑↑	0.81
avg_risk_DP	↓	-0.22	–	0.05	–	0.00
avg_risk_hubs	↑	0.25	↑↑	0.46	–	0.16

Table 2.3: Spearman's rank correlation coefficient of environmental parameters with respect to number of fatalities per task. Arrow indicates direction and strength of correlation, and when an arrow is indicated, correlation is statistically significant after applying Bonferroni's correction. One arrow = weak correlation, two arrows = moderate correlation, three arrows = strong correlation. – implies no (statistically significant) correlation.

2.5. Sensitivity study on grid size

As explained in the paper, the grid-size of the environment had to be increased because of constraints on the availability of computational power. The hypothesis is that this has little impact on the results, because the input data (population density and sheltering factor) is only defined on a relatively coarse grid (population density often only per neighborhood). In our work, we used a grid-size of 25 meters, but we carried out a sensitivity study where we varied the size of the edge of a grid between 10 meters and 50 meters. This study was simulated for a standard scenario of the city of Delft.

The results are presented in Figure 2.4. No difference in the results can be found per grid-size. A t-test is also used to analyse the difference between grid-sizes, which also gave no reason to reject the null-hypothesis that the results are from a different distribution. This reaffirms our hypothesis that the input data is the limiting factor to the accuracy of the results, and not the grid-size that was chosen because of computational limitations.

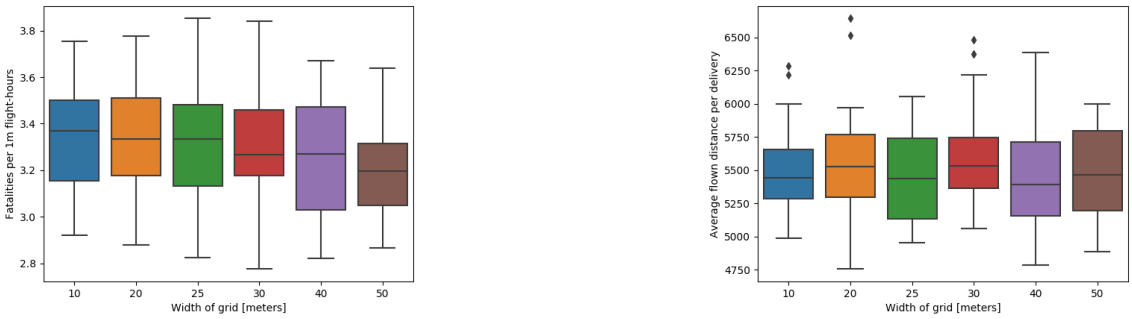

(a) Fatalities per 1m flight-hours.

(b) Average distance flown per delivery (meters).

Figure 2.4: Change in fatalities per 1m flight-hours and total collective risk because of changing the size of grids.

3

Statistical Elaboration on Experiments

This chapter elaborates on the experiments discussed in the paper. The first section elaborates on the statistical methods used in this research. The second section presents three experiments that have not been discussed in the main paper. The last section presents several statistical evaluations that were used to assure the integrity and reproduceability of the results.

3.1. Statistical elaboration

This section elaborates on the methodology behind three statistical methods, as used in this work. The first is the methodology for testing the hypothesis that the data follow a normal distribution. The second is how to determine the required number of experiments. Thirdly, the methodology for conducting a Global Sensitivity Analysis is conducted.

3.1.1. Test of normality

Testing if the data follows a normal distribution is an important step, as the T-test can only be used when the data follows a normal distribution. In this research, two methods are used, being the quantile-quantile-plot and the shapiro-wilk test.

A QQ-plot plots the data of a distribution against the normal distribution. If data is normally distributed, it should follow a straight line. An example is given in Figure 3.1.

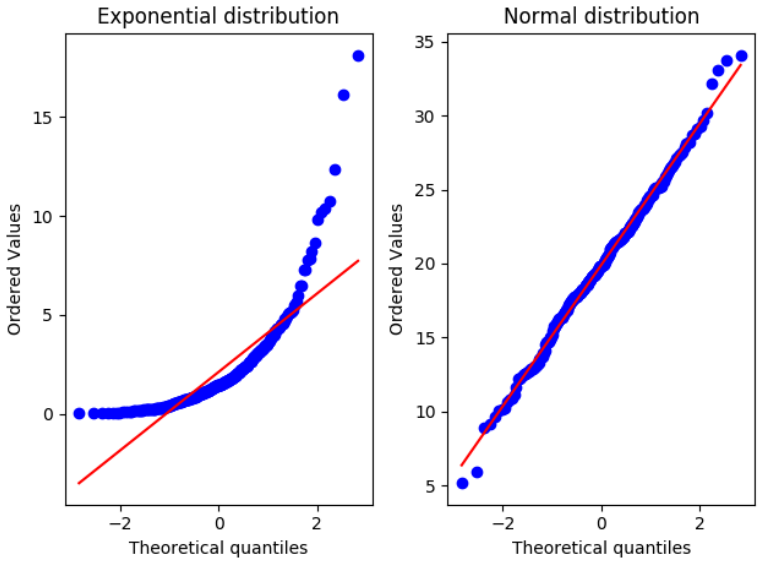

Figure 3.1: Example of one QQ-plot following a not-normal distribution (left), and one QQ-plot following a normal distribution (right).

The Sapiro-Wilk test is a statistic that can be used to test the null-hypothesis that a dataset follows the normal distribution [118]. It outputs a statistic W, and a p-value. If this p-value $<\alpha$, our confidence level, this means that the null-hypothesis should be rejected, i.e., that it is to be assumed that the data doesn't follow a normal distribution. If, however, the p-value is larger than α, it is not possible to say anything at all about the data, only that the null-hypothesis cannot be rejected.

3.1.2. Determining the required number of experiments

Assuming that the model output can have any distribution, one can look at the coefficient of variation (CV) to judge whether enough datapoints are gathered. This is defined as Equation 3.1. Alternatively, when the data follows a normal distribution, one can analyze the confidence interval and see whether it is small enough.

$$c_v = \frac{\sigma(o)}{\mu(o)} \tag{3.1}$$

3.1.3. Global Sensitivity Analysis

A famous saying is "Garbage in, garbage out". This saying implies that the underlying data- and assumptions of a model are poor, the results will be poor, as well. On this, Leamer said the following [82]:

> I have proposed a form of organized sensitivity analysis that I call 'global sensitivity analysis' in which a neighborhood of alternative assumptions is selected and the corresponding interval of inferences is identified. Conclusions are judged to be sturdy only if the neighborhood of assumptions is wide enough to be credible and the corresponding interval of inferences is narrow enough to be useful.

Next to using a local sensitivity analysis, a global sensitivity analysis (GSA) can be employed as well. In a GSA, multiple assumptions are varied simultaneously [116]. When performing such an analysis, we used partial correlation to investigate the relation between variables [4].

After having selected the bounds that input parameters can take, one should sample input parameters between those bounds. This is a computationally expensive process. Several techniques exist to sample variables more efficiently, reducing the variance of the Monte-Carlo simulations and thus reducing the required number of simulations. In this work, we use Latin-Hypercube Sampling. It derives its name from a Latin square, which is a square where there is only one sample on each axis. In LHS, the number N of samples needs to be determined beforehand. Each input parameter is then stratified in N components. This stratification process can be done using any appropriate distribution, however, in this work, the uniform distribution was taken. Having N components for M parameters, one can create the samples by taking one random component for each of the parameters. All parameters should be drawed exactly once to arrive at N samples of M components [89].

3.2. Statistical Evaluation

This section presents the steps taken to ensure that the data follows a normal distribution, and that enough experiments were run.

3.2.1. Proof of normality

Four experiments were analysed using the T-Test. For these experiments, both the QQ-plots were analysed, and the Shapiro-Wilk values were verified. These results of the Shapiro-Wilk tests, as well as the references to the QQ-plots, can be found in Table 3.1.

Experiment	Dataset	Shapiro-Wilk p-value
Vary P(failure) (Figure 3.3)	Base	0.31
	Higher P(Failure) in cruise&descent	0.08
Wind (Figure 3.5)	No wind	0.33
	Stochastic wind	0.25
Cruise altitude (Figure 3.2)	altitude = 100 meters	0.57
	altitude = 500 meters	0.55
V_impact (Figure 3.4)	base case	0.73
	Terminal	0.58
	Terminal + 40%	0.37

Table 3.1: Table showing validations of normality using Shapiro-Wilk test, as well as QQ-plots.

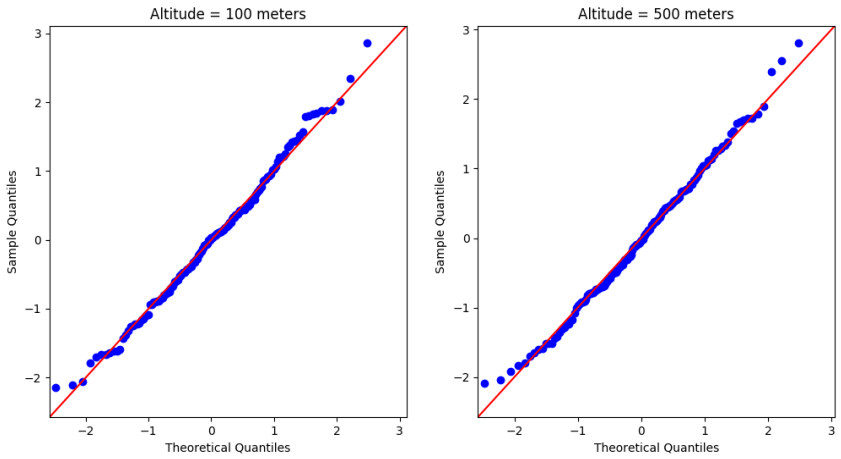

Figure 3.2: QQ-Plots of experimenting with the cruise altitude.

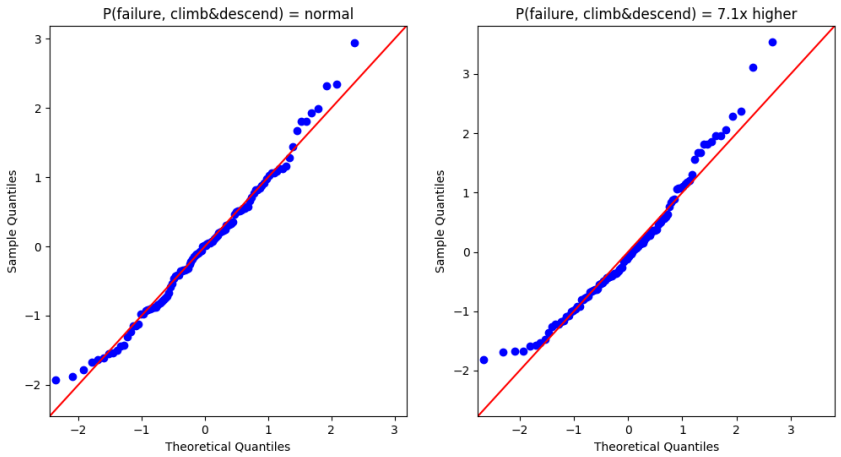

Figure 3.3: QQ-Plots of experimenting with the failure rate.

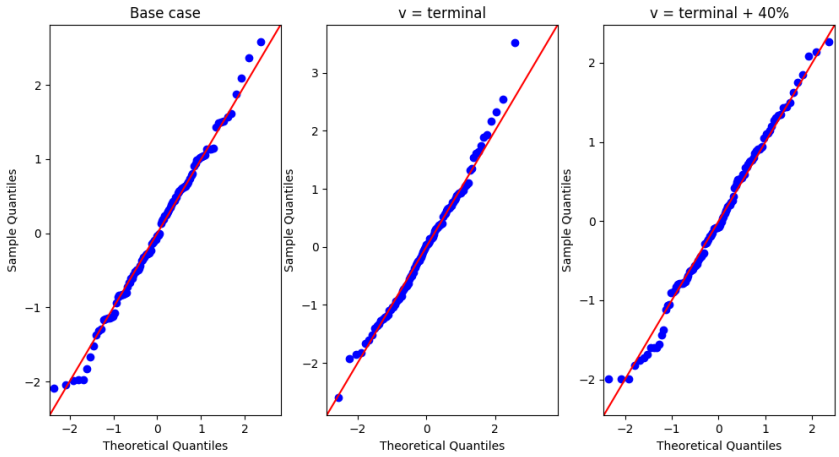

Figure 3.4: QQ-Plots of experimenting with the speed upon impact.

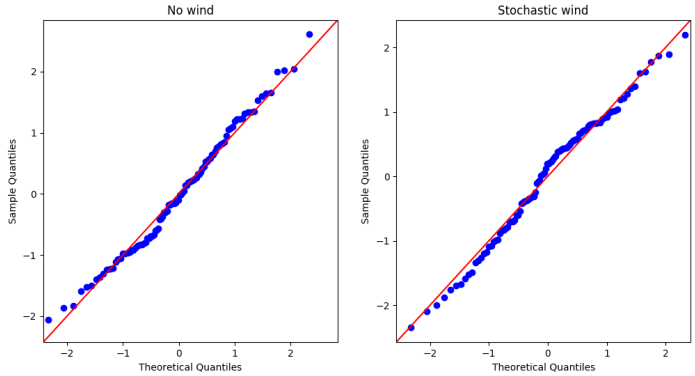

Figure 3.5: QQ-Plots of experimenting with the wind in the descent trajectory.

3.2.2. Required number of experiments

As described in subsection 3.1.2, the coefficient of variation is used to determine the required number of variations. This section shows the coefficients of variation for the four experiments where different distributions were compared.

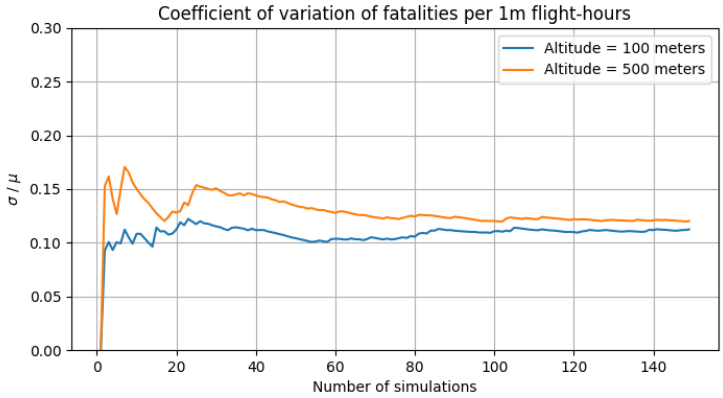

Figure 3.6: Coefficient of Variation of experimenting with the cruise altitude.

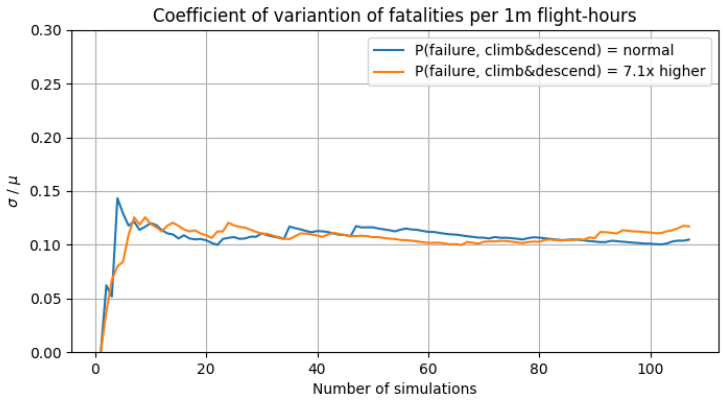

Figure 3.7: Coefficient of Variation of experimenting with the failure rate.

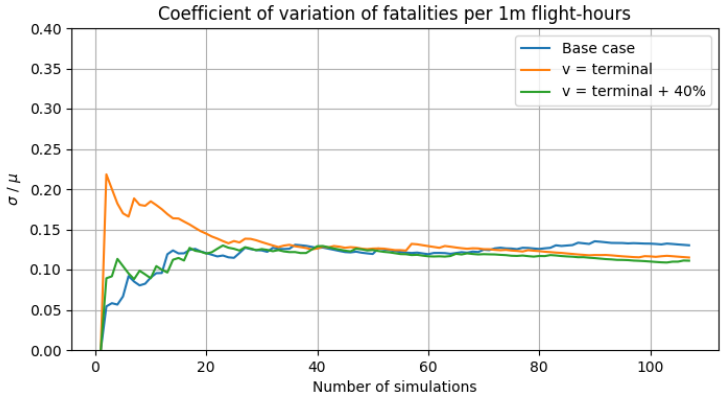

Figure 3.8: Coefficient of Variation of experimenting with the speed upon impact.

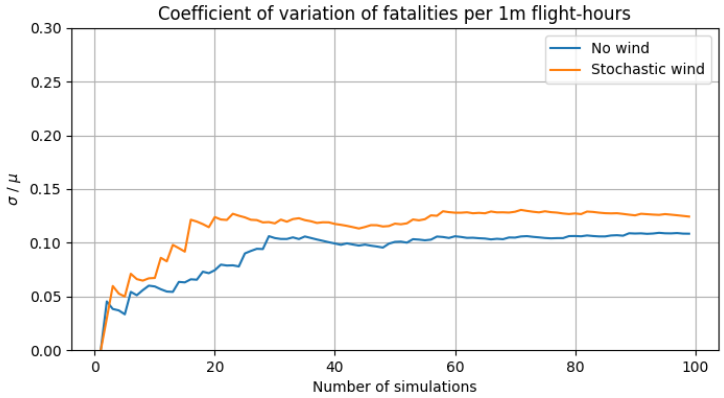

Figure 3.9: Coefficient of Variationof experimenting with the wind in the descent trajectory.

4

Recommendations for Future Work

This research has answered some questions regarding the safety of Urban UAS applications. However, many more questions remain. This chapter presents an overview of the recommendations for future work in this direction.

4.1. General future work

More accurate representation of population density: In the current model, the population density is based on Census data, which is based on where people live. This data is corrected using estimates of the whereabouts of people. However, this is only a coarse approximation of where people actually are in the environment. Secondly, this does not account for the fact that people may move during operations. Therefore, future work should validate whether our estimations of the whereabouts of people are reliable, and if not, develop better methods. Also, future work should research the effects of dynamic population density layers on the operations.

More accurate representation of sheltering data: We propose two improvements to the use of sheltering data. Firstly, in our work, only four distinct sheltering factors are used. Increasing the domain of possibilities could increase accuracy. Secondly, more data sources could be accessed to improve the coupling of location data with sheltering data. For example, the age and height of buildings could serve as additional data-points, increasing the reliability of the sheltering data.

Research effect of missing or wrong environmental data: Our research assumes that information on sheltering data and population density is fully available, without any errors. In reality, it is feasible that this data is not fully available, and it may be erroneous. The effect this has on the system behaviour should be researched, and also, possible mitigation methods could be researched.

Better drone dynamics: The dynamics of the drone are modeled simplistically. For example, instantaneous acceleration in any direction is assumed. It should be researched how the implementation of a realistic 6-degrees-of-freedom vehicle would affect the model behaviour. Alternatively, approximations of this behaviour could be made, for example through Dubins smoothening of paths [26].

Research sub-optimal heuristics: Path-finding proved to be computationally expensive in this research. One of the recommendations is to increase the resolution of the map. However, this will also increase the computational effort required for path-finding. The search-space of possible heuristics is unbounded [86], and several heuristics specifically for risk-based path-finding have been developed (e.g. [26, 104, 113]). Future work could analyze which heuristics increase the efficiency of path-finding specifically for the operations considered in this research.

Online and dynamic path-finding: Not only can the general efficiency of path-finding be improved, one should also consider the practicalities of implementing these path-finding methods into UASs. It is likely that these vehicles will have limited computational power, which can be limiting to path-finding. Recently,

Levasseur et al. researched the possibilities of training Neural Networks to approximate the paths found by path-finders. After initial training, these networks were shown to be able to find paths with less computational effort than regular path-finders [83]. Combining this direction with a dynamic environment and limited information would be a relevant future direction of research.

Task Allocation: Random task allocation was assumed in this research. Ways to perform a more efficient task allocation could include the bundling of tasks, scheduling tasks so they are performed when the risk along the path of this task is low, or scheduling a task from another hub whenever that could reduce risk.

Airspace structures: Several researchers (e.g. [58] and [9]) investigated how airspace structures (such as 3D highways through the sky) could be used to reduce the complexity of an airspace. Future research could investigate how such structures would influence risk.

Interactions with other vehicles: This research did not consider losses of separations between vehicles, and the avoidance-manoeuvres that may be required following such a loss of separation. Research may consider the effect of such manoeuvres on the risk, and could investigate ways to make such manoeuvres with the minimum risk as possible. Jenie et al. researched many methods of conflict detection and resolution for UAVs, and this research could serve as a starting point [66].

Other societal implications: Motivation for this research was the demand from regulators to better understand the safety of urban UAS operations. This demand follows from the task of regulators to protect safety from adverse effects of (new) technologies. However, there are more societal implications to be considered next to safety. A literature review of research on Unmanned Aerial Mobility lists noise as another barrier to the roll-out of urban UAS operations [140]. Noise is an example of another societal implication of these operations, and should also be researched.

4.2. Future work regarding risk computation

Several future directions, specifically concerning the risk computation, are identified.

Research other failure modes: This research only considered a ballistic descent following an all-engines-out failure. However, many more failure modes are conceivable. The system may fail in different ways, for example, because only some engines fail but not all, it may catch fire, or it may lose its payload. Also, the UAS could fail due to adverse weather or crashes with obstacles, birds or other vehicles. The last direction, concerning crashes with other vehicles, has been researched by Kim et al., who concluded that the crash absorbs a significant part of the kinetic energy [77]. These failure modes all have their own associated descent trajectory and probability of failure. Boolean trees [83], Bayesian Belief Networks [139], or Barrier Bowtie Models [18], can be used to research the probabilities and implications of all these failure modes.

research the probability of fatality: The current implementation of the function mapping an impact to a probability of fatality, is based on missile impact data from the 1960s [38]. Not only is an impact by a UAS very different from an impact by a missile, but also did this original data not consider sheltering from impacts. Both the effect sheltering has on reducing the impact energy, as the effect drone impacts have on human fatalities, should be researched in future work.

Other types of risk: This research considered Third-Party Risk. Even though this is indicated by many as the most important risk factor in analysing the safety of these operations, it doesn't cover the full picture. There is also a risk to operators of the vehicle, as well as to people who receive packages. Future research should investigate the magnitude of first- and second-party risk compared to third-party risk.

Risk mitigation: After failure, there is still much that can be done to mitigate risk. A UAS might be capable of changing its descent trajectory to reduce the risk, or recovery systems, such as parachutes (as proposed by [144], can be employed.

Bibliography

[1] Ersin Ancel, Francisco M. Capristan, John V. Foster, and Ryan C. Condottax. Real-time risk assessment framework for unmanned aircraft system (UAS) traffic management (UTM). *17th AIAA Aviation Technology, Integration, and Operations Conference, 2017*, pages 1–17, 2017. doi: 10.2514/6.2017-3273.

[2] Atelier Parisien d'Urbanisme (APUR). Population Density Paris, 2019. URL https://www.apur.org/en/geo-data/population-density.

[3] Autoriteit Consument & Markt. Post- en Pakkettenmonitor 2018. *Rapport*, page 31, 2018. URL https://d2aye3ggtndtn5.cloudfront.net/app/uploads/2019/08/post-en-pakkettenmonitor-2018-1.pdf.

[4] Kunihiro Baba, Ritei Shibata, and Masaaki Sibuya. PARTIAL CORRELATION AND CONDITIONAL CORRELATION AS MEASURES OF CONDITIONAL INDEPENDENCE. *Australian <html_ent glyph="@amp;" ascii="&"/> New Zealand Journal of Statistics*, 46(4):657–664, 12 2004. ISSN 1369-1473. doi: 10.1111/j.1467-842X.2004.00360.x. URL http://doi.wiley.com/10.1111/j.1467-842X.2004.00360.x.

[5] Lawrence C. Barr, Richard L. Newman, Ersin Ancel, Christine M. Belcastro, John V. Foster, Joni K. Evans, and David H. Klyde. Preliminary risk assessment for small unmanned aircraft systems. *17th AIAA Aviation Technology, Integration, and Operations Conference, 2017*, (June), 2017. doi: 10.2514/6.2017-3272.

[6] Mayank V. Bendarkar, Ameya Behere, Simon I. Briceno, and Dimitri N. Mavris. A Bayesian Safety Assessment Methodology for Novel Aircraft Architectures and Technologies Using Continuous FHA. (June), 2019. doi: 10.2514/6.2019-3123.

[7] S. Bertrand, N. Raballand, F. Viguier, and F. Muller. Ground risk assessment for long-range inspection missions of railways by UAVs. *2017 International Conference on Unmanned Aircraft Systems, ICUAS 2017*, pages 1343–1351, 2017. doi: 10.1109/ICUAS.2017.7991331.

[8] S. Bertrand, N. Raballand, and F. Viguier. Evaluating Ground Risk for Road Networks Induced by UAV Operations. *2018 International Conference on Unmanned Aircraft Systems, ICUAS 2018*, pages 168–176, 2018. doi: 10.1109/ICUAS.2018.8453441.

[9] Mohammed Faisal Bin Mohammed Salleh, Da Yang Tan, Choon Hian Koh, and K. H. Low. Preliminary concept of operations (ConOps) for traffic management of unmanned aircraft systems (TM-UAS) in urban environment. *AIAA Information Systems-AIAA Infotech at Aerospace, 2017*, (January):1–13, 2017. doi: 10.2514/6.2017-0223.

[10] Mohammed Faisal Bin Mohammed Salleh, Wanchao Chi, Zhenkun Wang, Shuangyao Huang, Da Yang Tan, Tingting Huang, and K. H. Low. Preliminary concept of adaptive urban airspace management for unmanned aircraft operations. *AIAA Information Systems-AIAA Infotech at Aerospace, 2018*, (209989): 1–12, 2018. doi: 10.2514/6.2018-2260.

[11] Cynthia Bir and David C. Viano. Design and injury assessment criteria for blunt ballistic impacts. *Journal of Trauma - Injury, Infection and Critical Care*, 57(6):1218–1224, 2004. ISSN 00225282. doi: 10.1097/01.TA.0000114066.77967.DE.

[12] Michal Cáp, Peter Novák, Jirí Vokrínek, and Michal Pechoucek. Multi-agent RRT*: Sampling-based Cooperative Pathfinding (Extended Abstract). *CoRR*, abs/1302.2, 2013. URL http://arxiv.org/abs/1302.2828.

[13] Centraal Bureau voor Statistiek. StatLine (OpenData). URL https://opendata.cbs.nl/#/CBS/nl/.

[14] Centraal Bureau voor Statistiek. Kadastrale Kaart Delft, 2018. URL https://kadastralekaart.com/gemeenten/delft-GM0503.

[15] Jen Jen Chung, Carrie Rebhuhn, Connor Yates, Geoffrey A. Hollinger, and Kagan Tumer. A multiagent framework for learning dynamic traffic management strategies. *Autonomous Robots*, 43(6):1375–1391, 2019. ISSN 15737527. doi: 10.1007/s10514-018-9800-z. URL https://doi.org/10.1007/s10514-018-9800-z.

[16] Victor R Clare, James H Lewis, Alexander P Mickiewicz, and Larry M Sturdivan. Blunt Trauma Data Correlation. (May):54, 1975. URL https://apps.dtic.mil/docs/citations/ADA012761.

[17] Reece Clothier, Rodney Walker, Neale Fulton, and Duncan Campbell. A casualty risk analysis for unmanned aerial system (UAS) operations over inhabited areas. *Second Australasian Unmanned Air Vehicle Conference*, pages 1–15, 2007. ISSN 0980321506.

[18] Reece Clothier, Brendan Williams, and Achim Washington. Development of a Template Safety Case for Unmanned Aircraft Operations over Populous Areas. *SAE Technical Papers*, 2015-Septe(September), 2015. ISSN 01487191. doi: 10.4271/2015-01-2469.

[19] Reece A. Clothier, Jennifer L. Palmer, Rodney A. Walker, and Neale L. Fulton. Definition of an airworthiness certification framework for civil unmanned aircraft systems. *Safety Science*, 49(6):871–885, 2011. ISSN 09257535. doi: 10.1016/j.ssci.2011.02.004. URL http://dx.doi.org/10.1016/j.ssci.2011.02.004.

[20] Reece A. Clothier, Brendan P. Williams, and Kelly J. Hayhurst. Modelling the risks remotely piloted aircraft pose to people on the ground. *Safety Science*, 101:33–47, 2018. ISSN 18791042. doi: 10.1016/j.ssci.2017.08.008.

[21] Compendium voor de Leefomgeving. Bevolkingsomvang en aantal huishoudens, 2018. URL https://www.clo.nl/indicatoren/nl000118-bevolkingsomvang-en-huishoudens.

[22] Anders La Cour-Harbo. Ground Impact Probability Distribution for Small Unmanned Aircraft in Ballistic Descent. *Reliability Engineering & System Safety*, (June), 2017. URL https://pdfs.semanticscholar.org/fdbe/5904b1265e80f1adbee8d30e88e1c0f7e128.pdf?_ga=2.195665260.1791246636.1560413169-249321586.1557012606.

[23] Anders La Cour-Harbo. The Value of Step-by-Step Risk Assessment for Unmanned Aircraft. *2018 International Conference on Unmanned Aircraft Systems, ICUAS 2018*, pages 149–157, 2018. doi: 10.1109/ICUAS.2018.8453411.

[24] K. Dalamagkidis, K. P. Valavanis, and L. A. Piegl. On unmanned aircraft systems issues, challenges and operational restrictions preventing integration into the National Airspace System. *Progress in Aerospace Sciences*, 44(7-8):503–519, 2008. ISSN 03760421. doi: 10.1016/j.paerosci.2008.08.001.

[25] Konstantinos Dalamagkidis, Kimon P. Valavanis, and Les A. Piegl. *On Integrating Unmanned Aircraft Systems into the National Airspace System*. Springer Netherlands, Dordrecht, 2012. ISBN 978-94-007-2478-5. doi: 10.1007/978-94-007-2479-2. URL http://link.springer.com/10.1007/978-94-007-2479-2.

[26] Luca De Filippis, Giorgio Guglieri, and Fulvia Quagliotti. A minimum risk approach for path planning of UAVs. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 61(1-4):203–219, 2011. ISSN 09210296. doi: 10.1007/s10846-010-9493-9.

[27] Malik Doole, Joost Ellerbroek, and Jacco Hoekstra. Drone delivery: Urban airspace traffic density estimation. *SESAR Innovation Days*, (December), 2018. ISSN 07701268.

[28] EASA. Policy for unmanned aerial vehicle (UAV) certification. *Management*, (16):1–42, 2005.

[29] EASA. Concept of Operations for Drones A risk based approach to regulation of unmanned aircraft. pages 1–8, 2015. URL http://www.sei.cmu.edu/cmmi/background/conops.html.

[30] EASA. NPA 2017-05 (A): Introduction of a Regulatory Framework for the Operation of Drones. 05:1–128, 2017. doi: RelatedA-NPA:2015-10RMT.023018.12.2015. URL https://www.easa.europa.eu/sites/default/files/dfu/NPA2017-05(A)_0.pdf.

[31] EASA. Notice of Proposed Amendment 2017-05 (B) - Introduction of a regulatory framework for the operation of drones. 05:128, 2017. URL https://www.easa.europa.eu/system/files/dfu/NPA2017-05%28B%29.pdf.

[32] EASA (European Aviation Safety Agency). European Aviation Safety Agency Introduction of a regulatory framework for the operation of unmanned aircraft. pages 1–50, 2015.

[33] European Aviation Safety Agency. Airworthiness Certification of Unmanned Aircraft Systems Policy Statement. pages 1–17, 2009.

[34] European Commission. Commission Delegated Regulation (EU) 2019/945 of 12 March 2019 on unmanned aircraft systems and on third-country operators of unmanned aircraft systems. (May 2019): C/2019/1821, 2019. URL http://data.europa.eu/eli/reg_del/2019/945/oj.

[35] Federal Aviation Administration. Pilot's Handbook of Aeronautical Knowledge. pages 74–81, 2016. doi: 10.4324/9780429030772-15. URL https://www.faa.gov/regulations_policies/handbooks_manuals/aviation/phak/media/17_phak_ch15.pdf.

[36] Federal Aviation Administration (FAA). FAR 107 - Code of Federal Regulations, 2016. URL http://www.ecfr.gov/cgi-bin/text-idx?c=ecfr&tpl=/ecfrbrowse/Title14/14tab_02.tpl.

[37] Federal Aviation Administration (FAA). Unmanned Aircraft System Traffic Management Concept of Operations. Technical report, FAA, 2020.

[38] D.I. Feinstein, W.F. Heugel, and M.L. Kardatzke. *Personnel Casualty Study*. Defense Technical Information Center, 1968.

[39] Rita Fontanella, Amedeo Rodi Vetrella, Giancarmine Fasano, Domenico Accardo, Rosario Schiano Lo Moriello, and Leopoldo Angrisani. Requirements, platform specifications, and system architectures for future unmanned traffic management systems. *AIAA Information Systems-AIAA Infotech at Aerospace, 2017*, (January):1–10, 2017. doi: 10.2514/6.2017-0225.

[40] Andrew T. Ford and Kevin J. Mcentee. Assessment of the risk to ground population due to an unmanned aircraft in-flight failure. *10th AIAA Aviation Technology, Integration and Operations Conference 2010, ATIO 2010*, 1(September):1–13, 2010. doi: 10.2514/6.2010-9056.

[41] Nigel Gilbert. Agent-Based Models. *Encyclopedia of GIS*, (September), 2008. doi: 10.1007/978-0-387-35973-1.

[42] Meir Goldenberg, Ariel Felner, Roni Stern, Guni Sharon, Nathan Sturtevant, Robert C. Holte, and Jonathan Schaeffer. Enhanced partial expansion A*. *Journal of Artificial Intelligence Research*, 50:141–187, 2014. ISSN 10769757. doi: 10.1613/jair.4171.

[43] LORNE GREENSPAN, BARRY A. McLELLAN, and HELEN GREIG. Abbreviated Injury Scale and Injury Severity Score. *The Journal of Trauma: Injury, Infection, and Critical Care*, 25(1):60–64, 1 1985. ISSN 0022-5282. doi: 10.1097/00005373-198501000-00010. URL http://journals.lww.com/00005373-198501000-00010.

[44] Dugway Proving Ground, Reagan Test Site, Yuma Proving Ground, White Sands, Missile Range, Naval Air, Aircraft Division, Naval Air, Weapons Division, Naval Undersea, Naval Undersea, Pacific Missile, Range Facility, Arnold Engineering, Development Complex, and National Aeronautics. Common Risk Criteria Standards for National Test Ranges Distribution a : Approved for Public Release. 2017.

[45] Wen-Lin Guan and Kay Yong. Ballistic Trajectory Analysis for the CI611 Accident Investigation. 2004. URL https://www.tailstrike.com/250502ballistic.pdf.

[46] Giorgio Guglieri and Gianluca Ristorto. Safety Assessment for Light Remotely Piloted Aircraft Systems. *2016 INAIR - International Conference on Air Transport*, 1(December):1–7, 2016. URL http://www.mavtech.eu/site/assets/files/1261/safety_assessment_for_light_remotely_piloted_aircraft_systems_inair_template_2016_final.pdf.

[47] Giorgio Guglieri, Alessandro Lombardi, and Gianluca Ristorto. Operation Oriented Path Planning Strategies for Rpas. *Amer. J. Sci. Technol.*, 2(6):321–328, 2015. ISSN 2375-3846. URL http://www.aascit.org/journal/ajst.

[48] Y. Haartsen, R. Aalmoes, and Y. S. Cheung. Simulation of unmanned aerial vehicles in the determination of accident locations. *2016 International Conference on Unmanned Aircraft Systems, ICUAS 2016*, pages 993–1002, 2016. doi: 10.1109/ICUAS.2016.7502548.

[49] D R Haddon and C J Whittaker. AIRCRAFT AIRWORTHINESS CERTIFICATION STANDARDS FOR CIVIL UAVs. D.R.Haddon, C.J.Whittaker - Civil Aviation Authority, UK. *Aviation*, 107(August):79–86, 2002. URL http://www.caa.co.uk/docs/1416/srg_acp_00016-01-120203.pdf.

[50] Peter E Hart, Nils J. Nilsson, and Bertram Raphael. Formal Basis for the Heuristic Determination eijj ,. *Systems Science and Cybernetics*, 4(2):100–107, 1968.

[51] Ludwig Hausmann, John Murnane, Mudra Neelesh, and Florian Neuhaus. The future of parcel delivery: Drones and disruption The Next Normal. *McKinsey & Company*, 2019. URL https://www.mckinsey.com/~/media/McKinsey/FeaturedInsights/TheNextNormal/The-Next-Normal-Thefuture-of-parcel-delivery-vF.

[52] Andrew J. Hawkins. Uber unveils a new look for its food delivery drones, 2019. URL https://www.theverge.com/2019/10/28/20936410/uber-eats-food-delivery-drone-design.

[53] Kit Heren. Surrey Police use drone to break up gatherings during coronavirus lockdown, 4 2020. URL https://www.standard.co.uk/news/uk/drones-surrey-police-gatherings-coronavirus-lockdown-a4412486.html.

[54] Carsten Hirschberg, Alexander Rajko, Thomas Schumacher, and Martin Wrulich. The Changing Market for Food Delivery. *McKinsey & Company*, 2016. URL https://www.mckinsey.com/industries/high-tech/our-insights/the-changing-market-for-food-delivery.

[55] J.M. Hoekstra, Stefan Kern, Olivier Schneider, Franz Knabe, and Bruno Lamiscarre. Metropolis Scenario Definition Report, 2015.

[56] J.M. Hoekstra, Stefan Kern, Olivier Schneider, Franz Knabe, and Bruno Lamiscarre. Metropolis Development & Metrics Definition. Technical report, 2015.

[57] J.M. Hoekstra, Stefan Kern, Olivier Schneider, Franz Knabe, and Bruno Lamiscarre. Metropolis Concept design, 2015.

[58] J.M. Hoekstra, Stefan Kern, Olivier Schneider, Franz Knabe, and Bruno Lamiscarre. Metropolis Simulation Results and Analysis Report. Technical report, 2015.

[59] J.M. Hoekstra, Stefan Kern, Olivier Schneider, Franz Knabe, and Bruno Lamiscarre. Metropolis Urban Airspace Design. Technical report, 2015.

[60] https://opendata.cbs.nl/#/CBS/nl/Imagemagick. ImageMagick. URL https://imagemagick.org/index.php.

[61] ICAO. ICAO Annex 11: Air Traffic Services, Chapter 2, Section 2.6, 1990. URL https://web.archive.org/web/20140407130450/http://www.icao.int/safety/ism/ICAOAnnexes/Forms/AllItems.aspx.

[62] Frank Jackman. Nearly Half of Commercial Jet Accidents Occur During Final Approach, Landing, 2014. URL https://flightsafety.org/asw-article/nearly-half-of-commercial-jet-accidents-occur-during-final-approach-landing/.

[63] JARUS. What's JARUS. URL http://jarus-rpas.org/.

[64] Yazdi I. Jenie, Erik Jan Van Kampen, Coen C. De Visser, Joost Ellerbroek, and Jacco M. Hoekstra. Selective velocity obstacle method for deconflicting maneuvers applied to unmanned aerial vehicles. *Journal of Guidance, Control, and Dynamics*, 38(6):1140–1145, 2015. ISSN 07315090. doi: 10.2514/1.G000737.

[65] Yazdi I. Jenie, Erik Jan Van Kampen, Joost Ellerbroek, and Jacco M. Hoekstra. Taxonomy of Conflict Detection and Resolution Approaches for Unmanned Aerial Vehicle in an Integrated Airspace. *IEEE Transactions on Intelligent Transportation Systems*, 18(3):558–567, 2017. ISSN 15249050. doi: 10.1109/TITS.2016.2580219.

[66] Yazdi Ibrahim Jenie. *AUTONOMOUS CONFLICT DETECTION AND RESOLUTION FOR UNMANNED AERIAL VEHICLES*. 2019. ISBN 9789055841745. doi: 10.4233/uuid.

[67] Chengpeng Jiang, Mathijs Post, Pieter Griffioen, Sieglinde Goossenaerts, and Bastiaan Zwanenburg. Report group 3 - AE4429 Advanced Agent-Based Modeling and Simulation in Air Transport. Technical report, 2020.

[68] Tao Jiang, Jared Geller, Daiheng Ni, and John Collura. Unmanned Aircraft System traffic management: Concept of operation and system architecture. *International Journal of Transportation Science and Technology*, 5(3):123–135, 2016. ISSN 20460449. doi: 10.1016/j.ijtst.2017.01.004. URL http://dx.doi.org/10.1016/j.ijtst.2017.01.004.

[69] Martin Joerss, Jürgen Schröder, Florian Neuhaus, Christoph Klink, and Florian Mann. Parcel delivery: The future of last mile. *McKinsey & Company*, (September):1–32, 2016. URL https://www.mckinsey.com/~/media/mckinsey/industries/traveltransportandlogistics/ourinsights/howcustomerdemandsarereshapinglastmiledelivery/parcel_delivery_the_future_of_last_mile.ashx.

[70] Marcus Johnson, Jaewoo Jung, Joseph Rios, Joey Mercer, Jeffrey Homola, Thomas Prevot, Daniel Mulfinger, and Parimal Kopardekar. Flight test evaluation of an unmanned aircraft system traffic management (UTM) concept for multiple beyond-visual-line-of-sight operations. *12th USA/Europe Air Traffic Management R and D Seminar*, (June), 2017.

[71] Joint Authorities for Rulemaking of Unmanned Systems (JARUS). JARUS guidelines on Specific Operations Risk Assessment (SORA), EDITION 2.0. 2019. URL http://jarus-uas.org.

[72] Myungsoo Jun and Raffaello DAndrea. Path Planning for Unmanned Aerial Vehicles in Uncertain and Adversarial Environments. pages 95–110, 2003. doi: 10.1007/978-1-4757-3758-5{_}6.

[73] Stan Kaplan. The words of risk analysis. *Risk Analysis*, 17(4):407–417, 1997. ISSN 02724332. doi: 10.1111/j.1539-6924.1997.tb00881.x.

[74] Uluhan Cem Kaya, Atilla Dogan, and Manfred Huber. A probabilistic risk assessment framework for the path planning of safe task-aware uas operations. *AIAA Scitech 2019 Forum*, (January):1–17, 2019. doi: 10.2514/6.2019-2079.

[75] Alicia Kelso. Food Delivery Via Drones May Be Closer Than You Think, 2019. URL https://www.forbes.com/sites/aliciakelso/2019/07/12/food-delivery-via-drones-may-be-closer-than-you-think/#652cb6be7438.

[76] Christopher G Kevorkian. *UAS Risk Analysis using Bayesian Belief Networks : An Application to the Virginia Tech ESPAARO*. PhD thesis, Virginia Polytechnic Institute and State University, 2016.

[77] Sang Hyun Kim. Third-party risk analysis of small unmanned aircraft systems operations. *Journal of Aerospace Information Systems*, 17(1):24–35, 2020. ISSN 23273097. doi: 10.2514/1.I010763.

[78] Parimal Kopardekar and Steve Bradford. UTM Research Transition Team Plan. Technical report, NASA Ames Research Center, 2017.

[79] Anders la Cour-Harbo. Quantifying Risk of Ground Impact Fatalities for Small Unmanned Aircraft. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 93(1-2):367–384, 2019. ISSN 15730409. doi: 10.1007/s10846-018-0853-1.

[80] Anders la Cour-Harbo. Ground impact probability distribution for small unmanned aircraft in ballistic descent. *2020 International Conference on Unmanned Aircraft Systems, ICUAS 2020*, (February):1442–1451, 2020. doi: 10.1109/ICUAS48674.2020.9213990.

[81] Anders la Cour-Harbo and Henrik Schiøler. Probability of Low-Altitude Midair Collision Between General Aviation and Unmanned Aircraft. *Risk Analysis*, 39(11):2499–2513, 2019. ISSN 15396924. doi: 10.1111/risa.13368.

[82] Edward Leamer. Let's Take the Con Out of Econometrics. *The American Economic Review*, 73(1):31–43, 1983. URL https://www.jstor.org/stable/1803924%0A.

[83] Baptiste Levasseur, Sylvain Bertrand, Nicolas Raballand, Flavien Viguier, and Gregoire Goussu. Accurate Ground Impact Footprints and Probabilistic Maps for Risk Analysis of UAV Missions. *IEEE Aerospace Conference Proceedings*, 2019-March:1–10, 2019. ISSN 1095323X. doi: 10.1109/AERO.2019.8741718.

[84] Chin E. Lin, Tsung Pao Chen, Pei Chi Shao, Ying Chih Lai, Tzung Cheng Chen, and Yun Chao Yeh. Prototype Hierarchical UAS Traffic Management System in Taiwan. *Integrated Communications, Navigation and Surveillance Conference, ICNS*, 2019-April:1–13, 2019. ISSN 21554951. doi: 10.1109/ICNSURV.2019.8735380.

[85] David Locascio, Mason Levy, Kiran Ravikumar, Brian German, Simon I. Briceno, and Dimitri N. Mavris. Evaluation of Concepts of Operations for sUAS Package Delivery. *16th AIAA Aviation Technology, Integration, and Operations Conference, 2016*, (June):1–16, 2016. doi: 10.2514/6.2016-4371.

[86] Michael Lones. *Sean Luke: essentials of metaheuristics*, volume 12. 2015. ISBN 9781300549628. doi: 10.1007/s10710-011-9139-0.

[87] Rémi Louf and Marc Barthelemy. A typology of street patterns. *Journal of the Royal Society Interface*, 11 (101):1–7, 2014. ISSN 17425662. doi: 10.1098/rsif.2014.0924.

[88] Christopher W. Lum and Blake Waggonery. A risk based paradigm and model for unmanned aerial systems in the national airspace. *AIAA Infotech at Aerospace Conference and Exhibit 2011*, pages 1–31, 2011.

[89] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 1979. ISSN 15372723. doi: 10.1080/00401706.2000.10485979.

[90] Richard Melnyk, Daniel Schrage, Vitali Volovoi, and Hernando Jimenez. A third-party casualty risk model for unmanned aircraft system operations. *Reliability Engineering and System Safety*, 124:105–116, 2014. ISSN 09518320. doi: 10.1016/j.ress.2013.11.016. URL http://dx.doi.org/10.1016/j.ress.2013.11.016.

[91] Alec Momont and TU Delft. Ambulance Drone. URL https://www.tudelft.nl/en/ide/research/research-labs/applied-labs/ambulance-drone/.

[92] NASA. Terminal Velocity (gravity and drag). URL https://www.grc.nasa.gov/WWW/K-12/rocket/termvr.html.

[93] NASA. NASA UTM in 2019: Highlights. Technical report, NASA, 2019. URL https://utm.arc.nasa.gov/utm2019.shtml.

[94] National Transportation Safety Board. Aircraft Accident Report Trajectory study - N93119. 1996. URL http://www.ntsb.gov/doclib/reports/2001/aar0102.pdf.

[95] Victoria Chibuogu Nneji, Alexander Stimpson, Mary Missy Cummings, and Kenneth H. Goodrich. Exploring concepts of operations for on-demand passenger air transportation. *17th AIAA Aviation Technology, Integration, and Operations Conference, 2017*, (June):1–13, 2017. doi: 10.2514/6.2017-3085.

[96] Oberlo.com. Global Ecommerce Sales, 2017. URL https://www.oberlo.com/statistics/global-ecommerce-sales.

[97] Maarten Oonk. Using Cargo Drones in Last Mile Delivery. *Deloitte*, 2018. URL https://www2.deloitte.com/nl/nl/pages/consumer-industrial-products/articles/using-cargo-drones-in-last-mile-delivery.html.

[98] oozo.nl. Wetenswaardigheden, cijfers en statistieken over Amsterdam, 2018. URL https://www.oozo.nl/cijfers/amsterdam.

[99] Ashim Panta, Simon Watkins, and Reece Clothier. Dynamics of a small unmanned aircraft parachute system. *Journal of Aerospace Technology and Management*, 10:1–14, 2018. ISSN 21759146. doi: 10.5028/jatm.v10.752.

[100] Pitney Bowes. Shipping Index, 2019. URL https://www.pitneybowes.com/us/shipping-index.html.

[101] Port Technology International. Rotterdam hosts first ever drone delivery, 5 2020. URL https://www.porttechnology.org/news/rotterdam-hosts-first-ever-drone-delivery/.

[102] PostEurop. Facts & Figures, 2018. URL https://deliver4europe.eu/facts-figures/.

[103] Thomas Prevot, Joseph Rios, Parimal Kopardekar, John E. Robinson III, Marcus Johnson, and Jaewoo Jung. UAS Traffic Management (UTM) Concept of Operations to Safely Enable Low Altitude Flight Operations. (June):1–16, 2016. doi: 10.2514/6.2016-3292.

[104] Stefano Primatesta, Luca Spanò Cuomo, Giorgio Guglieri, and Alessandro Rizzo. An Innovative Algorithm to Estimate Risk Optimum Path for Unmanned Aerial Vehicles in Urban Environments. *Transportation Research Procedia*, 35:44–53, 2018. ISSN 23521465. doi: 10.1016/j.trpro.2018.12.006. URL https://doi.org/10.1016/j.trpro.2018.12.006.

[105] Stefano Primatesta, Giorgio Guglieri, and Alessandro Rizzo. A Risk-Aware Path Planning Strategy for UAVs in Urban Environments. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 95 (2):629–643, 2019. ISSN 15730409. doi: 10.1007/s10846-018-0924-3. URL http://dx.doi.org/10.1007/s10846-018-0924-3.

[106] Stefano Primatesta, Alessandro Rizzo, and Anders la Cour-Harbo. Ground Risk Map for Unmanned Aircraft in Urban Environments. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 97(3-4):489–509, 2020. ISSN 15730409. doi: 10.1007/s10846-019-01015-z.

[107] Laurent Probst, Bertrand Pedersen, Lauriane Dakkak-Arnoux, and PwC. Digital Transformation Monitor The: Drones in Agriculture. *European Union*, (January), 2018. URL https://ec.europa.eu/growth/tools-databases/dem/monitor/sites/default/files/Drones_vf.pdf.

[108] Alexander Radi. Human Injury Model for Small Unmanned Aircraft Impacts. pages 1–31, 2013. URL https://www.casa.gov.au/sites/g/files/net351/f/_assets/main/airworth/papers/human-injury-model-small-unmanned-aircraft-impacts.pdf.

[109] David Raymond, Chris Van Ee, Gregory Crawford, and Cynthia Bir. Tolerance of the skull to blunt ballistic temporo-parietal impact. *Journal of Biomechanics*, 42(15):2479–2485, 2009. ISSN 00219290. doi: 10.1016/j.jbiomech.2009.07.018. URL http://dx.doi.org/10.1016/j.jbiomech.2009.07.018.

[110] Joseph L Rios, Irene S Smith, David R Smith, Sheryl Jurcak, and Randy Strauss. UTM UAS Service Supplier Development. *NASA Technical Memorandum*, (December):1–21, 2018.

[111] Robotics & Automation. UPS Foundation supports Ghana's vaccine drone delivery network, 2019. URL http://roboticsandautomationnews.com/2019/05/03/ups-foundation-supports-ghanas-vaccine-drone-delivery-network/22162/.

[112] Eliot Rudnick-Cohen, Jeffrey W. Herrmann, and Shapour Azarm. Risk-based path planning optimization methods for unmanned aerial vehicles over inhabited areas. *Journal of Computing and Information Science in Engineering*, 16(2):1–7, 2016. ISSN 15309827. doi: 10.1115/1.4033235.

[113] Eliot Rudnick-Cohen, Shapour Azarm, and Jeffrey W. Herrmann. Planning unmanned aerial system (UAS) takeoff trajectories to minimize third-party risk. *2019 International Conference on Unmanned Aircraft Systems, ICUAS 2019*, pages 1306–1315, 2019. doi: 10.1109/ICUAS.2019.8798149.

[114] Eliot Rudnick-Cohen, Jeffrey W. Herrmann, and Shapour Azarm. Modeling unmanned aerial system (UAS) risks via monte carlo simulation. *2019 International Conference on Unmanned Aircraft Systems, ICUAS 2019*, pages 1296–1305, 2019. doi: 10.1109/ICUAS.2019.8798313.

[115] Peter Sachs, Marc Baumgartner, Ed Diaz-gomez, Steve Weidner, Isabel Del-pozo-de poza, Robert Hoffman, Bruno Ley, Philippe Masson, and Joerg P Mueller. For the sky. 2018.

[116] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global Sensitivity Analysis. The Primer*. 2008. ISBN 9780470059975. doi: 10.1002/9780470725184.

[117] SESAR JU. U-space blueprint - SESAR Joint Undertaking. 2017. doi: 10.2829/614891. URL https://www.sesarju.eu/u-space-blueprint.

[118] S. S. Shapiro and M. B. Wilk. An Analysis of Variance Test for Normality (Complete Samples). *Biometrika*, 52(3/4):591, 1965. ISSN 00063444. doi: 10.2307/2333709.

[119] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R. Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015. ISSN 00043702. doi: 10.1016/j.artint.2014.11.006. URL http://dx.doi.org/10.1016/j.artint.2014.11.006.

[120] A Sharpanskykh. AE4422-19 Agent-based Modelling and Simulation in Air Transport (Lecture 1), 2019.

[121] Alexei Sharpanskykh. Towards an Agent-based Approach for Safety and Risk Analysis of Safety-Critical Organizations. In *International Conference on Modelling and Analysis of Safety and Risk in Complex Systems*, 2014.

[122] Trevor Standley. Finding optimal solutions to cooperative pathfinding problems. *Proceedings of the National Conference on Artificial Intelligence*, 1:173–178, 2010.

[123] Statista. Online Food Delivery Market, 2020. URL https://www.statista.com/outlook/243/100/ecommerce/worldwide#market-revenue.

[124] Jonathan D. Stevenson, Siu OYoung, and Luc Rolland. Estimated levels of safety for small unmanned aerial vehicles and risk mitigation strategies. *Journal of Unmanned Vehicle Systems*, 3(4):205–221, 2015. ISSN 2291-3467. doi: 10.1139/juvs-2014-0016.

[125] Anna Straubinger, Raoul Rothfeld, Michael Shamiyeh, Kai Daniel Büchter, Jochen Kaiser, and Kay Olaf Plötner. An overview of current research and developments in urban air mobility Setting the scene for UAM introduction. *Journal of Air Transport Management*, 87(August 2019), 2020. ISSN 09696997. doi: 10.1016/j.jairtraman.2020.101852.

[126] Larry M. Sturdivan, David C. Viano, and Howard R. Champion. Analysis of Injury Criteria to Assess Chest and Abdominal Injury Risks in Blunt and Ballistic Impacts. *Journal of Trauma - Injury, Infection and Critical Care*, 56(3):651–663, 2004. ISSN 00225282. doi: 10.1097/01.TA.0000074108.36517.D4.

[127] Nicholas Swift. easy-a-star-pathfinding-7e6689c7f7b2 @ medium.com, 2017. URL https://medium.com/@nicholas.w.swift/easy-a-star-pathfinding-7e6689c7f7b2.

[128] Kristian Husum Terkildsen and Kjeld Jensen. Towards a tool for assessing UAS compliance with the JARUS SORA guidelines. *2019 International Conference on Unmanned Aircraft Systems, ICUAS 2019*, pages 460–466, 2019. doi: 10.1109/ICUAS.2019.8798236.

[129] Texas A&M Transport Institute. Urban mobility Scorecard. (August), 2015. URL https://static.tti.tamu.edu/tti.tamu.edu/documents/umr/archive/mobility-scorecard-2015.pdf.

[130] The Anonymous Widower (Blog). Parcel Pick-up Point 2. URL https://anonw.com/2017/04/06/between-exeter-st-davids-and-exeter-central-stations/dscn0327-4/.

[131] Trust for London. London's geography and population, 2016. URL https://www.trustforlondon.org.uk/data/londons-geography/.

[132] Uber. Uber Copter infopage, 2020. URL https://www.uber.com/nl/nl/ride/uber-copter/.

[133] Uber. Uber Elevate, 2020.

[134] US Census Bureau. New York City population density, 2019. URL https://www.arcgis.com/home/webmap/viewer.html?url=https%3A%2F%2Fportal.sgm.gob.mx%2Farcgis%2Frest%2Fservices%2FDatosAbiertos%2FDatosAbiertos%2FMapServer&source=sd.

[135] Stefanie Valentic. Cardiac Arrest Survival Rates Higher with AED Use, 2018. URL https://www.ehstoday.com/health/article/21919523/cardiac-arrest-survival-rates-higher-with-aed-use.

[136] Voom. Company website, 2020. URL https://www.voom.flights/en.

[137] Achim Washington, Reece A. Clothier, and Jose Silva. A review of unmanned aircraft system ground risk models. *Progress in Aerospace Sciences*, 95(August):24–44, 2017. ISSN 03760421. doi: 10.1016/j.paerosci.2017.10.001. URL https://doi.org/10.1016/j.paerosci.2017.10.001.

[138] Achim Washington, Reece A. Clothier, and Brendan P. Williams. A Bayesian approach to system safety assessment and compliance assessment for Unmanned Aircraft Systems. *Journal of Air Transport Management*, 62:18–33, 2017. ISSN 09696997. doi: 10.1016/j.jairtraman.2017.02.003. URL http://dx.doi.org/10.1016/j.jairtraman.2017.02.003.

[139] Achim Washington, Reece Clothier, Natasha Neogi, Jose Silva, Kelly Hayhurst, and Brendan Williams. Adoption of a Bayesian Belief Network for the System Safety Assessment of Remotely Piloted Aircraft Systems. *Safety Science*, 118(March):654–673, 2019. ISSN 18791042. doi: 10.1016/j.ssci.2019.04.040. URL https://doi.org/10.1016/j.ssci.2019.04.040.

[140] S. Watkins, J. Burry, A. Mohamed, M. Marino, S. Prudden, A. Fisher, N. Kloet, T. Jakobi, and R. Clothier. Ten questions concerning the use of drones in urban environments. *Building and Environment*, 167 (August 2019):106458, 2020. ISSN 03601323. doi: 10.1016/j.buildenv.2019.106458. URL https://doi.org/10.1016/j.buildenv.2019.106458.

[141] Roland E Weibel and R John Hansman. Safety Considerations for Operation of Unmanned Aerial Vehicles in the National Airspace System. *Transportation*, 37(March):26–30, 2005.

[142] Gerhard Weiss. *Multiagent Systems*. MIT Press, 2 edition, 2013. ISBN 9780262018890. URL https://mitpress-mit-edu.tudelft.idm.oclc.org/books/multiagent-systems-second-edition.

[143] Paul Wu and Reece A. Clothier. The development of ground impact models for the analysis of the risks associated with unmanned aircraft operations over inhabited areas. *11th International Probabilistic Safety Assessment and Management Conference and the Annual European Safety and Reliability Conference 2012, PSAM11 ESREL 2012*, 7:5222–5234, 2012.

[144] Tim Wyllie. Parachute recovery for UAV systems. *Aircraft Engineering and Aerospace Technology*, 73(6): 542–551, 2001.

[145] Zipline. Company Website, 2020. URL https://flyzipline.com/.