Algorithms and Visualizations to Support Airborne Detection of Vertical Obstacles

by

Paul Flanigen

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Robotics)
in the University of Michigan
2023

Doctoral Committee:

Professor Ella M. Atkins, Co-Chair
Professor Nadine B. Sarter, Co-Chair
Associate Professor Dimitra Panagou
Associate Professor X. Jessie Yang

Paul Flanigen

pflani@umich.edu

ORCID iD: 0000-0002-4009-8490

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

APPENDIX

ABSTRACT

Slow or failed detection of low salience vertical obstacles and associated wires is one of today's leading causes of fatal helicopter accidents. The risk of collisions with such obstacles is likely to increase as Advanced Aerial Mobility and broadening drone activity promises to increase the density of air traffic at low altitudes, while growing demand for electricity and communication will expand the number of vertical structures. The current 'see-and-avoid' detection paradigm relies on pilots to spend much of their visual attention looking outside for obstacles. This method is inadequate in low visibility conditions, cluttered environments and given the need for pilots to engage in multiple competing visual tasks. With the expected growing number of hazards and an increased traffic volume, the current approach to collision avoidance will become even less tenable. This dissertation provides methods for automatic detection and correlation of vertical obstacles and evaluates the effectiveness of sensor visualizations and graphic augmentations for supporting flight crews in noticing hazards.

The first contribution of this line of research is a modular set of algorithms which extract towers from raw point clouds. Vertical structures compose less than 0.2% of real world tiles containing over 100 million points. A mesh filter quickly removes large, flat surfaces. Next, a sphere detector finds and eliminates vegetation protrusions. Dense point clouds undergo clustering and a proportional height filter which increases the density of vertical structures over 2,000%. Sparse and cluttered point clouds pass through an overlap filter which effectively identifies vertical structures amid clutter.

The second contribution is an exploration of the current challenges and mitigations for obstacle detection, followed by a simulator study that compared tower detection times for combinations of sensor visualizations and graphic augmentations. A set of focus groups revealed that detecting obstacles remains a significant challenge and that current mitigation strategies are not sufficient to prevent collisions. A subsequent human-in-the-loop simulator study revealed that graphic augmentations led to faster tower detection time when ambient visibility and illumination was reduced close to the limit for visual flight. Bounding boxes around towers were detected first in all conditions but tended to mask the obstacle they were meant to highlight. Sensor visualization affected tower detection time only at night, where night vision goggles were more effective than the infrared thermal sensor.

The third contribution of this dissertation is a method to efficiently correlate vertical structure observations with existing databases and infer the presence of power lines. The method uses a spatial hash key which compares an observed tower to existing towers and updates similar objects based on height and position. When applied to Delaware's Digital Obstacle File, average horizontal uncertainty decreased from 206 to 56 ft. Power line presence is inferred by automatically examining the arrangement of towers in the more accurate database. Over 87% of electrical transmission towers were correctly identified with no false negatives.

In summary, this thesis contributes to a better understanding of the current limitations of vertical obstacle detection and avoidance. It proposes and assesses modular methods to automatically detect, catalog, and categorize hazardous obstacles that are currently neglected, and it evaluates the effectiveness of current visualization technologies and sensor- and database-informed graphic augmentations for supporting pilots in the timely and reliable detection of towers. Taken together, this research will contribute to enhanced aviation safety in the low altitude environment.

CHAPTER 1

Introduction

1.1 Motivation

Close surveillance of the surrounding airspace is a fundamental task of all aircrews. They must quickly and accurately locate a wide variety of expected and unexpected hazards. Aircrews are cautioned to remain on the lookout for unknown hazards, especially at low altitudes. Searching for expected and noticing unexpected, low salience vertical obstacles (such as communication towers shown in Figure 1.1) is particularly challenging [4]. Inadequate detection and perception of these obstacles is behind only loss of situation awareness such as controlled flight into terrain (CFIT), as the most common known cause of fatal civilian helicopter accidents [5]. Obstacle strikes are also a leading cause of military rotor craft losses [6].

Figure 1.1: A helicopter pilot's perspective approaching a ridge line (left) and the same view with eight towers overlaid with red symbols (right).

The situation is likely to get worse as the same stakeholders that want convenient transportation and delivery services also have an increasing appetite for information and energy. The quantity of communication and electrical transmission towers continue to grow to meet this need. Over 67,000 cell sites were built in the US between 2018-2020, which is more than the total built in the preceding seven years [7]. There were more than 418,000 operational cell sites at the end of 2021 [8]. The number of electrical transmission towers will also increase as they adapt to increasingly decentralized renewable energy sources and rising energy needs. To fill these gaps, electricity transmission systems need to expand 60% by 2030 [9]. This burgeoning infrastructure will proliferate in the same low altitudes coveted by next generation aircraft.

There are few recent examples of efforts [10] [11] [12] to address the need for supporting enroute obstacle detection and avoidance. These studies equipped helicopters with onboard sensors, additional displays and automation that interfaced directly with the aircraft flight controls to reduce pilot workload and allow nearly autonomous approaches to helicopter landing zones. However, the systems required experienced pilots familiar with the intricacies of each technology, supplemented by ground and air control stations to monitor each flight. The above efforts did not prioritize enroute obstacle detection at low altitudes. And they did not address the need to highlight obstacles for human operators who are charged with performing multiple sometimes competing tasks and supervisory responsibilities. An algorithm or display which presents an obstacle but does not actively support the operator in noticing the obstacle's presence is at best a partial solution.

Helicopter pilots spend most of their time multitasking. They scan the environment continuously to identify and react to potential hazards. They must also maintain awareness of their position relative to the intended flight path and navigation aids. This involves monitoring navigation instruments, charts, and landmarks out the window to ensure they stay on course and reach their destination safely. Pilots monitor and transmit on radios, sometimes supplemented by Traffic Collision Avoidance Systems (TCAS). Flight crews must continually assess weather conditions, especially during low altitude flights where visibility can be affected by fog, low clouds, or precipitation. Monitoring weather radar, assessing cloud formations, and observing changes in visibility are crucial tasks to ensure safe flight operations. Various helicopter parameters and performance indicators, such as engine and transmission gauges, fuel levels, and payloads also need to be monitored. Dividing attention between all these visually demanding tasks is a challenge and calls for interventions that capture and guide pilots' attention to critical objects and events in a timely manner.

To be effective, the design of such interventions should be driven by insights from models of human attention, such as Guided Search [13] and NT-SEEV (Noticing Time-Salience Effort Expectancy Value) [14]. These models will be described in some detail in Chapter 3 which reports on a simulator study on the effectiveness of graphic augmentations of obstacles. Attention guidance often relies on increasing the salience of an object and needs to guard against the risk of invoking inattentional blindness (e.g., [15]). This phenomenon tends to be experienced when it becomes impossible for a human to attend to all stimuli in a given situation, and the individual then fails to see unexpected objects that are fully visible because of a lack of attention. For example, a vertical obstacle may be highlighted for a pilot who then misses another critical event, such as an intruding aircraft. Our automatic vertical structure detection method (Chapter 2), visualization options (Chapter 3), and correlation approach (Chapter 4) add options for finding and presenting obstacles among overwhelming information in the low altitude flight environment

1.2 Problem Statement

In this dissertation, we aim to address the following questions:

1. How can we efficiently find vertical structures in raw sensor data?
   There are a variety of offline and online approaches to discern vertical obstacles [1]. Offline processes [16] [17] [18] [19] take substantial time to segment large amounts of data. Leading online approaches, such as Density-Based Clustering of Applications with Noise (DBSCAN) clustering [20], are simpler but also vulnerable to sparsity and clutter. Ground-based Mobile LiDAR Systems (MLS) [21] [22] [23] [24] [25] [26] take advantage of a relatively uniform and predictable environment that adjoins roadways. Existing airborne sensing approaches [27] [28] [29] [30] [31] [32] require manual offline manipulation of sensor data for accurate segmentation.

2. What types of sensor visualization and graphic augmentation are most effective at decreasing obstacle notice time?
   Additional sources of visual information, such as Synthetic Vision Systems (SVS), are designed to decrease noticing time by providing a supplemental view of the external environment. However, previous work [33] has shown that pilots tend to miss unexpected obstacles only visible out the window when they fixate on a panel mounted SVS display. Augmenting visualizations with graphics designed to highlight obstacles has the potential to decrease misses and time needed to notice obstacles [34, 35]. Pilots initiated avoidance maneuvers for obstacles placed along a simulated helicopter 1,000 or 200 ft Above Ground Level (AGL) flight route sooner when the object was flashing or brightened on a moving map display [36]. Helicopter pilots in a low altitude (100 ft AGL) study noticed obstacles farther away when they were augmented with graphics on a heads-up display [34]. One risk associated with increasing the salience of a target is that it can cause inattentional blindness (failing to notice a target even though it is within foveal vision) towards un-augmented hazards [36, 37, 15, 38]. Varying the cueing precision also affects search time. Less precise cueing increases the time required to find targets and this tendency was exacerbated with lower salience targets [15]. More detailed obstacle cueing is not always better. In several other helicopter flight simulations [39, 40, 41, 42], pilots preferred the most basic obstacle position information without suggested maneuver cueing and other symbology. Other visual cueing schemes that include 3D conformal symbology show potential to improve obstacle detection, but have been inconclusive or tended to over-saturate pilot participants [43, 44, 45].

In addition, sensor playback studies have revealed pilots' preference to turn off information layers (including sensor visualizations) to decrease information overload [46]. The sensor image from an SVS thus has the potential to, ironically, decrease awareness of unexpected obstacles.

3. How can vertical structure position information be efficiently consolidated?
   The Federal Aviation Administration's (FAA's) Digital Obstacle File (DOF) is the definitive, publicly available source for vertical structures that could be a hazard to flight operations. Adding and revising man-made obstacles that are far away from airports with instrument approaches largely relies on voluntary reporting from infrastructure builders. Obstructions greater than three miles from designated airports that are less than 499 ft above ground level are not considered "obstructions to air navigation," [47] making low-altitude obstacles especially prone to oversight.

   Due to the ever changing nature and sheer quantity of information [48] and electrical [49] infrastructure, there have been several efforts to automate the mapping process. Recent efforts use night-time lighting patterns in satellite imagery to predict infrastructure position to within 1,000 m 70% [50] to 75% [51] of the time. Other efforts at mapping vertical structures do not update the high resolution map [52, 53]. Refs [54, 55] rely on repeated encounters at close range and refer to a small 30 m square map. Recent approaches that use semantic labelling for localization and mapping rely on continuous, dense surfaces that are associated with rich imagery [56].

4. How can vertical structure position information be used to find power lines?
   Wire strikes are an even more widespread hazard to low-altitude flight operations than collisions with vertical obstacles. Current wire finding methods depend on detecting wires directly. Airborne methods that automatically segment power lines rely on continuous contact [57] [58], known location [59] [60], and/or very close range [61] [62] [63].

1.3 Research Approach and Dissertation Outline

This dissertation explores algorithms designed to find vertical structures, along with power lines and other associated hazards, within raw and voluminous point cloud data. Labelled point cloud databases provide a baseline to test our methods. Various obstacle databases provide partial lists of obstacles with assorted levels of inaccuracy. We supplement these databases with simulated point clouds that are realistically sparse to check that our methods have the potential for online implementation. Although LiDAR sensors are the source for this data, the approach is designed to be sensor agnostic, allowing for automatic distillation of co-registered point clouds that might come from different perspectives and sensor modalities.

Automatically detected towers are valuable when they are integrated into existing databases. More recent, and typically higher accuracy, observations are compared to the existing Digital Obstacle File using cross referenced hash tables. This improved accuracy is exploited to infer the presence of power lines across the entire database.

After presenting an approach for finding and cataloging towers, we investigate the effect of sensor-informed graphics and visualizations on tower detection time. A human-in-the-loop study examines the challenges human pilots face when trying to locate hazardous vertical structures in cluttered or low visibility surroundings. Specifically, a high-fidelity helicopter flight simulator is coupled with a detailed model of the San Francisco peninsula to provide an immersive flight experience. The San Francisco model is supplemented by real world vertical structures to provide additional realism. Sensor visualizations are adjusted to imitate actual thermal and image intensification sensors in weather conditions which are just above the regulatory limit for visual flight. Two types of graphic augmentations (bounding boxes and circles around tower bases) are evaluated for their effectiveness at improving detection rate and time for towers.

The dissertation research pipeline is shown below in Figure 1.2.

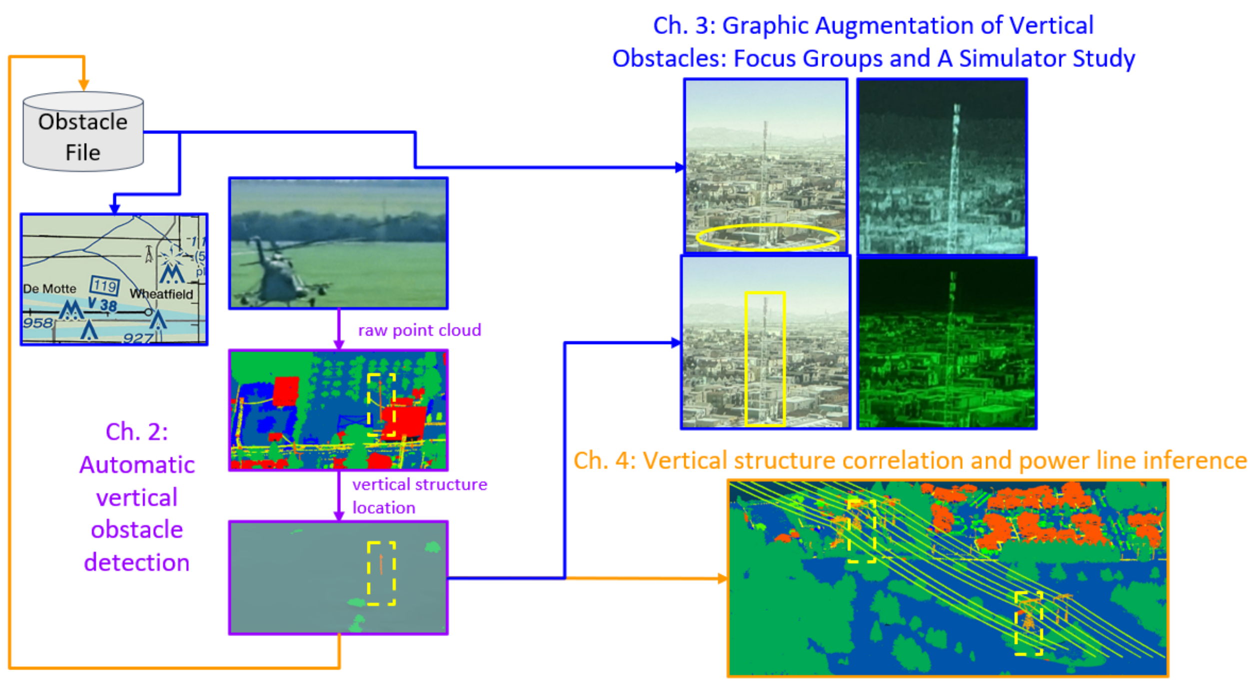

Figure 1.2: Research approach and dissertation outline.

This dissertation accomplishes the following three objectives:

1. (Chapter 2) Automatic Vertical Obstacle Detection

2. (Chapter 3) Graphic Augmentation of Vertical Obstacles: Focus Groups and A Simulator Study

3. (Chapter 4) Efficient Vertical Structure Correlation and Power Line Inference

1.3.1 Automatic Vertical Obstacle Detection

Our method starts with a novel mesh filter process. This step removes the vast majority of lower altitude points from consideration. Next, we find embedded spherical point arrangements generally associated with remaining tall vegetation. The remaining points are analyzed with two parallel processes. The first process finds structures with nearly continuous returns using a proportional height filter. A second parallel process is designed to find structures from sparse returns using our overlap filter. The objects of interest from each process are then consolidated.

Our approach is modular, fast, explainable, and inclusive. Each discrete filtering step accepts and outputs raw and unorganized point clouds that are sensor agnostic. The overall approach is designed for online implementation that requires minimal tuning or adjustments based on the environment. The logic within each module explains how a feature is segmented (or not) to provide a sense of limitations. Finally, the modules process rather than arbitrarily discarding sparse point cloud data that are often necessary to deduce the presence of vertical obstacles.

1.3.2 Studying the effect of sensor visualization and augmentation on obstacle notice time in the low-altitude flight environment

A series of focus groups were conducted to gain current perspective on the obstacle detection challenge [4]. Twelve participants from industry, military and government discussed the difficulty of detection and merits of various mitigations for avoiding vertical obstacles. Active military and civilian helicopter pilots were paired with a variety of engineers in four 2-hour online focus groups. Descriptive codes were manually assigned to the transcripts. Mitigation themes and sentiments emerged.

A pilot-in-the-loop simulator study then evaluated the effectiveness of various visual cueing schemes for supporting fast and reliable obstacle detection. Specifically, a within-subjects design was employed to compare the efficacy of various visualization types (unaided, image intensification, or thermal imaging) and obstacle augmentation techniques. Obstacle augmentation consisted of supplemental graphics (none, a priori circle, or more precise bounding box) shown in an immersive flight simulation. Eye tracking data, audio recordings, and qualitative surveys from participants were used to investigate ways of balancing vital aircrew duties with timely perception of obstacles in the low-altitude flight environment.

1.3.3 Efficient Vertical Structure Correlation and Power Line Inference

The last chapter proposes a method for automatically correlating vertical structures. Current databases are saddled with manual processes that result in incomplete and inaccurate data. We create a hash table to enable lookup and correlation with $O(1)$ complexity. Our hash key resolution is within DOF's most accurate horizontal uncertainty, thus avoiding repeated entries for the same coordinates while retaining relevance throughout a hemisphere. Next, an Index Hash Table (IHT) assigns this spatial key to existing entry values (such as position and height uncertainty). Simply searching for the nearest neighbor is not sufficient due to numerous vertical obstacles and their associated large, uneven, and overlapping horizontal uncertainties. The Uncertainty Hash Table (UHT) builds a list of vertical structure hashes whose uncertainty encompasses a given spatial hash. Observed towers are efficiently correlated and used to update existing entries.

We propose a novel power line inference method which, instead of searching for the virtually invisible wire, predicts a wire's presence based on the configuration and geometric arrangement of the supporting transmission towers or other vertical structure. This section of work leverages accurate tower positioning (Chapter 2) and the aforementioned improved database accuracy. Pairs of towers that are within a distance proportional to their height and have a similar height are associated. The angle between continuous sets of three towers is also taken into account to reduce false positives.

1.4 Contributions and Innovations

Specific contributions of this thesis are:

- Definition of the vertical obstacle challenge. Although obstacle collisions are a leading cause of accidents, no previous work has investigated why it remains a challenge with insight from pilots and other experts.

- Simulation of airborne returns in a virtual environment. Vertical structures make up a small portion of existing point cloud databases and are typically not uniquely labelled.

- Benchmark of existing DBSCAN algorithm effectiveness on classifying vertical obstacles. This widely used clustering approach in used in many other applications.

- Definition of metrics to quantitatively describe the prevalence and prominence of under-represented vertical objects. These metrics help to more accurately portray the limited tower returns (prevalence) and the significance of towers based on their surroundings (prominence).

- Evaluation of the effectiveness of sensor-informed graphic augmentations on obstacle notice time. A flight simulator study compared tower detection time between towers without graphic argumentation to detection times when towers were augmented with increasing precision.
- Evaluation of the effectiveness of high fidelity sensor visualizations on obstacle notice time. A flight simulator study compared tower detection time between leading unaided and sensor visualization types.

Specific innovations of this thesis are:

- A new overlap algorithm to find vertical structures represented by sparse and cluttered point clouds. Vertical structures small cross section usually does not create dense returns.
- A novel mesh filter process that quickly removes extraneous points from a scene. Vertical structure returns are normally a very small subset of real world scenes.
- A new clustering approach which uses vegetation's inherent sphericity to isolate tall protrusions. Vegetation, like vertical structures, can protrude above surrounding terrain, but is not as insidious.
- A robust proportional height filter to rapidly distill vertical structures. Given dense returns, this filter recognizes clusters that are likely to contain towers.
- A new method to efficiently correlate vertical structures. This method accommodates accuracy details for each entry over an entire hemisphere.
- A novel approach to reliably finding potentially hazardous wires. Instead of detecting power lines and other hazards supported by vertical structures, this method infers the presence of power lines by examining the arrangement of the support structure.

CHAPTER 2

Automatic Vertical Obstacle Detection

This chapter surveys and assesses previous methods of vertical obstacle detection, including remote sensing, photogrammetry, and LIDAR sensors. Next, it examines the suitability of current point cloud segmentation methods, including clustering methods. Two new algorithms overcome unique aerial detection challenges. They broaden the safe volume by considering sparse data available from small obstacles at distances that will allow high flight speeds. Presented and existing clustering methodologies are evaluated against a real world point cloud dataset or, to simulate sparse and cluttered airborne returns, a variety of vertical obstacles in an Unreal Engine [64] environment using Microsoft AirSim [65] LIDAR simulation.

Section 2 summarizes related work. Section 3 provides an overview of the three modules of our method: the mesh filter, the proportional height filter for continuous structures, and the overlap filter for sparse structures. Section 4 introduces and describes the DALES data set used for our experimentation. Section 5 presents results of our methods organized by vertical structure type. Section 6 discusses overall performance of our method and applicability to the aerial vehicle use case.

2.1 Background

There are a variety of offline and online approaches to discern vertical obstacles. Offline processes take substantial time to segment large amounts of data. Faster online approaches are simpler but also vulnerable to sparsity and clutter. Ground-based Mobile LiDAR Systems (MLS) take advantage of a relatively uniform and predictable environment that adjoins roadways. Existing airborne sensors require manual offline manipulation of sensor data for accurate segmentation. Related work is summarized below.

2.1.1 Offline Segmentation

Accurately classifying details in large amounts of data is a significant research area in the geospatial information sciences and agricultural fields. However, existing techniques are not designed for online utilization. Approaches that do not ignore thin objects require tuning of multiple parameters, depend on regular point spacing, continuous point clouds, or specific characteristics (such as cross arms) to find vertical structures.

Remote Sensing Ref. [16] proposed ways of extracting aviation obstacle information from satellite-based Synthetic Aperture Radar (SAR). SAR uses the moving perspective of the satellite to produce a three-dimensional image from multiple range-bearing returns. Some SAR processes place objects within $10m$ of their actual location. The low cross sections of towers and other vertical obstacles from a satellite's perspective makes photogrammetry, combining sets of visual images of the same location from different perspectives, challenging [17]. Satellite-based LiDAR systems have the capability to resolve one meter height differences for thin objects, but there are few orbiting systems due to high cost.

Segmentation from Mobile LiDAR Point Cloud Mobile LiDAR Systems (MLS) provide dense point clouds that contain poles along roadways. Previous works sought to precisely locate poles for navigation or surveying electrical infrastructure. Ref. [21] divided a point cloud into columns, which are further divided into stacks of blocks. A continuous stack of occupied blocks indicated a pole-like object. Ref. [22] also used a block stacking approach. Ref. [23] extracted poles from a point cloud by stacking doughnut-shaped clusters. Occupied points were allowed inside of the inner radius but not allowed between the "hole" and outer radius. Stacking pucks was also used by Ref. [24]. Ref. [25] used regular scan line spacings that were not continuous with other surfaces. Their process allowed for 2D segmentation of scan lines and only analyzed neighboring points for the sake of computational efficiency. Ref. [26] found poles from MLS by reducing 3D point clouds to a vertical plane then used Principal Component Analysis (PCA) to compare eigenvalues and find a principal direction. Ref. [25] also used PCA to associate vertically clustered point clouds that shared the same axis. These approaches used PCA to find the dominant axis given the area for a single vertical object. All these MLS pole segmentation approaches require some amount of prior knowledge about pole geometry or position in the environment (along the side of the road, for example). Also, these rule-based approaches do not propose an automatic process for handling large data sets in real time.

Inference from Airborne LiDAR Point Cloud The primary challenge in finding vertical obstructions from an airborne platform is to detect small objects within huge point clouds encompassing a variety of scenes. More than 99% of sensor data from a typical scene is composed of returns from objects other than towers [27]. Ref. [28] examined how LiDAR tilt angles, divergence, and flying height could be optimized to obtain points within a search cylinder around known vertical obstacles. Low reflectance combined with small cross sectional area makes vertical obstacles especially difficult to detect. Steep look down angles were necessary (0 to 40 degrees from vertical) for sufficient returns. Most approaches to segmentation of Airborne LiDAR System (ALS) data require a user to manually remove outliers or other points that are assumed to be less relevant. In addition to manually removing outliers, many recent methods for segmenting trees require the creation of a Canopy Height Model ([29], [30], [31]). In reference [27], the lowest elevation points which generally correspond to vegetation, from three to eight meters above ground level (AGL), were removed. Ref. [32] successfully found 28 electrical pylons within a dense, raw point cloud, but their approach required specific settings to segment nearly identical vertical structures from uniform power line corridors.

Machine learning also struggles to find vertical obstacles in sparse and cluttered point clouds. Classes which represent the minority of data sets generally have higher error rates[66]. One recent Convolutional Neural Network (CNN) [18] had an F_1 score of 54.7% for the poles class of the Dayton Annotated LiDAR Earth Scan (DALES) [67] ALS data set. Another neural network achieved an F_1 score of 91.8% on DALES' pole class [19], exceeding the performance of other recent networks with F_1 scores that ranged from 20 to 75%. A machine learning approach specifically designed for finding towers required at least 20 points per tower object while leveraging radiometric properties that are specific to advanced LiDAR systems. Reference [27] relied heavily on coincidentally collected color and infrared data to differentiate vegetation by electromagnetic radiation. Although it is tempting to use these radiometric properties, waveform qualities vary due to the season (e.g. leaf and grass color along with leaf presence for deciduous vegetation) and weather (e.g. snow cover). Radiometric filtering also increases a system's hardware and software cost and complexity. These tuned approaches depend on accurate and comprehensive training data and similar evaluation sets. Also, these approaches do not prioritize finding pole points that are likely to be most significant, such as those at a higher altitude above the ground.

2.1.2 Online Passive Sensing

Passive sensing options (including cameras) are ubiquitous due to their low cost and simplicity. However, it is challenging to detect objects with a low cross section that are small or far away.

Database Informed Terrain Awareness and Warning Systems (TAWS) and Synthetic Vision Systems (SVS) provide warning of impending collisions with obstacles or terrain based on databases. TAWS use terrain and obstacle databases to advise the pilot of impending collisions. Low database accuracy prevents them from being used as a source of navigation, however. The authors propose an algorithm in Ref. [68] that updates the existing Digital Elevation Model (DEM) with higher resolution obstacle information. However, their approach extends the no fly zone imparted by the obstacle height across the entire area between the DEM's post spacing.

Photogrammetry Camera sensors have the advantage of being light, relatively simple, and cheap. However, a camera's ability to portray small objects decreases rapidly as range to the object increases. Detection range decreases further when an object adjoins non-uniform textures or when ambient lighting is low. Photographs from Google's Street View database were used in Ref. [69] to map utility poles with cross arms. Pole-like objects were assumed to originate from the same plane as the automobile. This approach required multiple perspectives for each pole which is not a tenable approach for an air vehicle. In Ref. [57], an Unmanned Air Vehicle (UAV) used edge detection and Hough Transformation to discern power lines from various backgrounds. They required nearly continuous visual contact with the straight line under consideration, which is not feasible for aerial detection of vertical obstacles. In another effort to automatically track power lines, Ref. [61] used corner detection from an airborne camera to detect the cross beams of poles. Detecting these corners required a cross beam image width of approximately ten pixels. This resolution of obstacle data is not available across the field of action and distance required to allow efficient ground speed. An autonomous helicopter seeking to avoid obstacles in Ref. [70] hosted a stereo camera pair and 2D scanning laser with effective sensing ranges of 35 and 15 m, respectively. The camera-only setup was only 42% effective. Ref. [62] used stereo cameras to detect edges of thin obstacles in an office environment including power cords and networking cables. Image sequences were used to deduce the obstacle's 3D location from a quadcopter. However, the sensor was within five feet of the obstacle, and the thin obstacle's background was uniform.

2.1.3 Online Active Sensing

Airspace Occupancy Mapping Airborne Laser Scanners (ALS) have been used for obstacle warning [71]. Previous work relied on point cloud post-processing to detect vertical obstacles with motion planners reacting to nearly perpendicular angles of incidence. Ref. [72] sensed and avoided vertical obstacles including single trees, groups of trees, wire fences, and sheds at a cruise speed of $1m/s$. Ref. [73] used local planning to avoid obstacles sensed by LiDAR to achieve cruise speeds up to $10m/s$. Ref. [70] checked for occupied voxels along a line of spheres that advanced ahead of a UAV's projected path, achieving a cruise speed of $1m/s$. Ref. [60] extracted power line primitives by applying a Voxel-based Piece-wise Line Detector [74] but assumed power line posts had previously been identified. All these techniques avoided airspace that coincided with potentially occupied voxels and are vulnerable to false returns.

Table 2.1: Existing Segmentation Algorithm Comparison from [1].

Algorithm	Description	Pros	Cons
Euclidean Clustering [20][75]	Associates points that are closer than a given distance from one another	Multiple quick options Few (2-3) parameters to adjust	Builds according to point spacing Run time increases quickly with larger point spacing
Principal Component Analysis [25][26]	Determines variation along each axis of a point set; principal axis has most variance	One parameter to adjust	Requires dominant axis of variation Struggles with irregular data Only derives one axis per point set
Connected Components [76]	Clusters points based on angle of incidence between two adjoining rays	Two parameters Distinguishes close objects	Published online implementation relies on simplified 2D range image Struggles with discontinuous point clouds
Ring Stacking [23][24][25]	Rings with an occupied inner radius and unoccupied outer radius are stacked to indicate the presence of vertical object	Infers presence of thin objects despite clutter Tolerates some leaning	Five parameters to adjust Depends on uniform cross section
Block Stacking [21][22]	Block with minimum concentration of points are stacked	Compensates for clutter above or below a block stack	Six parameters to adjust Requires uniform vertical spacing between blocks Assumes poles are isolated

2.2 Methods

Our approach is modular, fast, explainable, and inclusive. Each discrete filtering step accepts and outputs raw and unorganized point clouds that are sensor agnostic. The overall approach is designed for online implementation that requires minimal tuning or adjustments based on the environment. The automatic and dynamic tuning for module can explain how a feature is segmented (or not) to provide a sense of limitations. Finally, the modules do not arbitrarily discard sparse point cloud data that are often necessary to deduce the presence of vertical obstacles.

Starting with a raw, unorganized point cloud, we first apply a mesh filter from above the terrain. This step removes the vast majority of lower altitude points from consideration. Next, we apply a RANdom SAmple Consensus (RANSAC) sphere detector which finds embedded spherical point arrangements within remaining tall vegetation. The remaining points are analyzed with two parallel processes. To find structures with nearly continuous returns, points are consolidated with 3D connected component clustering. Next, a filter checks the proportional height of each connected component cluster to see whether the bounding box is sufficiently tall and narrow. The second parallel process is designed to find structures from sparse returns. The same remaining points are clustered by the Density-Based Clustering of Applications with Noise (DBSCAN). Then the overlap filter checks among the DBSCAN clusters to determine if disparate returns belong to the same vertical structure. The objects of interest from each process are then consolidated.

2.2.1 Coarse Mesh Filter

The Cloth Simulation Filter (CSF) [77] is a widely used algorithm used to extract ground points from point clouds. We substantially modified the CSF process to efficiently remove ground and

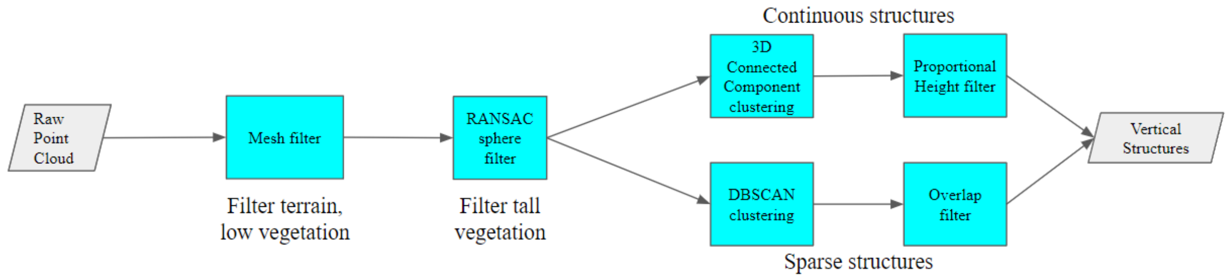

Figure 2.1: Overview of methods.

other low-lying points. First, all points are retained without any manual or automatic outlier removal. Second, the cloth simulation starts above (not below) the point cloud surface. Third, the distance between cloth particles is increased 2- to 24-times larger than the default value. This coarse mesh allows vertical obstacles to protrude while discarding innocuous ground points. The coarse mesh size also drastically decreases computation time. We determine the optimal grid spacing from our experimental results. Fourth, each point higher than the nearest coarse mesh point is classified as a non-ground point, regardless of proximity to the mesh.

Our coarse mesh process accepts an unfiltered set of 3D sensor points (gray boxes in Figure 2.2) to represent the terrain. This set includes outliers that are not clearly associated with a ground object. The blue mesh particles are initialized with uniform wide spacing above the raw point cloud. This mesh of blue dots descends directly downwards, bypassing any gray 3D point which is not directly below it. Mesh particles stop descending when they collide with a 3D point or are constrained by an immovable neighbor. This deformation constraint is symbolized by the lines between the orange particles in the final mesh position. Red points that protrude above the final coarse mesh are identified as points of interest.

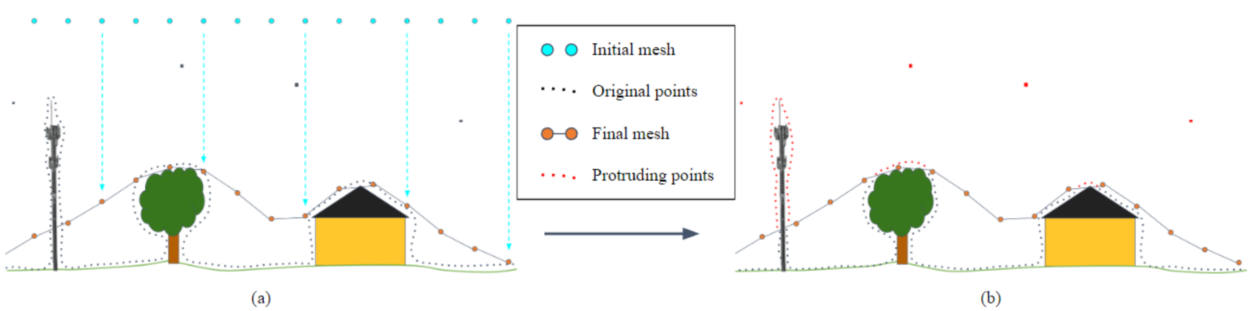

Figure 2.2: Coarse mesh process. Mesh particles (blue dots) are initialized above the 3D points (gray rectangles). The mesh particles descend until colliding with corresponding 3D points or until their movement is constrained due to immovable neighbors (orange dots in a). The points above the settled coarse mesh are shown in red (b).

As in [77], we consider external and internal forces acting on each particle. The external forces are descent due to gravity opposed by contact with a LiDAR point. The internal forces are those caused by the simulated connections between mesh particles. The following preliminaries present how we adapt the existing descent and deformation approach to our purpose.

2.2.1.1 Mesh descent

The mesh with user-defined particle spacing s starts in a 2D horizontal layer above the terrain. Each point on the mesh is projected downwards onto the LiDAR points. The altitude of each grid particle's corresponding point is recorded as the lowest point the particle can descend. Equation 2.1 is a simplified version of this process from [77]:

$$z_{t+1} = 2 \cdot z_t - z_{t-1} + g \cdot \Delta t^2 \tag{2.1}$$

where z_t is the current particle altitude, z_{t-1} is the particle's previous altitude, g is acceleration due to gravity, and Δt is the change in time.

2.2.1.2 Mesh deformation

For each time step with movable and immovable grid particles, the additional motion of movable particles that adjoin immovable particles is described by Equation 2.2 simplified from [77].

$$m_{t+1} = \frac{1}{2} \cdot (i_t - m_t) \tag{2.2}$$

where i_t is the height of the immovable grid particle and m_t is the height of the movable grid particle. This has the effect of dampening the deformation. Algorithm 1 presents a simplified time step of the CSF approach. Particles fall according to Equation 2.1. A particle's descent is constrained by immovable neighbors. This constraint decreases the vertical displacement for each time step, eventually halting the mesh descent and deformation prior to colliding with lower LiDAR points.

Increasing the mesh particle spacing significantly decreases the number of particles in the grid. With default CSF 0.5 m grid resolution for a 500 m^2 tile, each loop must traverse 1 million particles. Increasing spacing to 2 or 12 m decreases the order of magnitude of particles under consideration by 2- to 3-times, respectively. Additionally, the number of iterations decreases as the widely-spaced mesh more rapidly reaches the break point where the internal force constraint (within the `Deform` function) reduces maximum descent difference ($Diff$) to less than 0.005 m. Our unique top-down mesh application, exploits this deformation constraint to stop the mesh from settling onto the ground.

Algorithm 1 Mesh algorithm from the source code referenced in [2] with our modifications highlighted in blue. Note that while the original CSF method and source code are from [2] the CSF algorithm was not formally defined in previous work.

procedure MESH
  $m_M \leftarrow True$ ▷ all mesh particles initially movable
  while $Iterations \leq 500$ **do**
    for $m_t \in G$ **do** ▷ for each mesh particle in grid
      $m_{t+1} \leftarrow Fall(m_t)$ ▷ free fall
    for $m_{t+1} \in G$ **do**
      if $neighbor_{t+1} \neq m_M$ **then** ▷ if a neighboring mesh particle is immovable
        $m_{t+1} \leftarrow Deform(m_{t+1})$
    $maxDiff \leftarrow 0$
    for $m_t, m_{t+1} \in G$ **do**
      $Diff \leftarrow m_t - m_{t+1}$
      if $Diff > maxDiff$ **then**
        $maxDiff \leftarrow Diff$
    for $m_{t+1} \in G$ **do**
      $m_{t+1}, m_M \leftarrow Collide(m_{t+1})$
    if $Diff \leq 0.005$ **then** *break*
  for $l_j \in \mathcal{P}$ **do** ▷ for each LiDAR point in original point cloud
    if $l_j > m_C$ **then** ▷ if LiDAR point is higher than corresponding mesh point
      $l_j \leftarrow AboveMesh$
function FALL($m_t, m_{t-1}, g, \Delta t$)
  if m_M **then**
    $m_{t+1} \leftarrow 2 \cdot (m_t - m_{t-1}) + g \cdot \Delta t^2$
function DEFORM(m_{t+1})
  $Constrain \leftarrow 0.00678$ ▷ displacement restriction based on 15 neighbors
  $m_{t+1} \leftarrow m_{t+1} \cdot Constrain$ ▷ mesh particle descent constrained by neighbors
function COLLIDE(m_{t+1})
  $l_i \leftarrow CorrespondingLiDARPoint$
  if $m_{t+1} \leq l_i$ **then**
    $m_{t+1} \leftarrow l_i$ ▷ mesh particle altitude is set to LiDAR point altitude
    $m_M \leftarrow False$

2.2.1.3 Protrusion identification

Algorithm 1 also reveals (in blue) how we modify the previous approach to identify *all* points above the mesh as points of interest. After the cloth reaches stopping criteria in the previous CSF approach, only LiDAR points that are within a certain range (default 0.5 m) of the cloth are identified as ground points. The ultimate output of this coarse mesh module is a very small (generally less than 1%) subset of the original raw point cloud that has an increased density of returns from vertical objects.

2.2.2 Clustering

After applying the coarse mesh, the much smaller quantity of points above the mesh include tops of building and trees. This smaller quantity makes it feasible to quickly consider every remaining point. For most points on a shape there exists a neighborhood such that all of the points belong to that shape. To remove these short and wide distractions, we cluster the points based on the inherent spherical shape of vegetation tops. Next, we use the insight on their height in the context of grid spacing to isolate the low profile vertical objects.

Both of the following clustering steps use octrees to batch the unorganized point cloud. The original three dimensional space is subdivided into eight octants. Each octant that contains a point continues to be subdivided until reaching a specified depth or number of levels. For eight levels in a 500 by 500 m point cloud, this results in a grid step size of approximately 2 m.

2.2.2.1 Clustering by sphericity

The tree top protrusions tend to be uniquely bulbous after the mesh. Most other protrusions are more linear (such as towers) or planar (roof tops). We use sphere detection from the RANSAC shape detection approach [3] to identify tall vegetation. RANSAC checks random samples of the given point cloud until it finds an object primitive that satisfies specified criteria.

Algorithm 2 presents modified pseudo-code from the RANSAC point-cloud shape detection approach [3]. While searching for spheres, the criteria are maximum radius, r_{max}, and minimum number of LiDAR points, $|l|_{min}$, per sphere. We select r_{max} based on reasonably expected widths of protrusions from the previous mesh filtering module. An $r_{max} = 10m$ was suitable for the selected data set. The minimum number of points is largely dependent on the expected density of points in the scene. We set $|l|_{min} = 500$ for our dense data set with 400% coverage. The minimum number of points may decrease for a use case with more sparse point clouds.

The sphere detection algorithm starts by finding sphere candidates from the point cloud $\mathcal{P}$ left over from the mesh filter. The random selection of the first point is followed by more careful selection of a second point within the same octree level. If the angular difference between the

Algorithm 2 RANSAC sphere extraction derived from [3], added

procedure VEGETATION EXTRACTION($\mathcal{P}, r_{max}, |l|_{min}, \alpha$)
  while $P(|C|, r_{max}) > 99\%$ **do** ▷ probability of overlooking number of candidates with maximum radius
    $C \leftarrow C \cup newCandidates(\mathcal{P}, \alpha)$
    $l \leftarrow bestCandidate(C)$
    if $P(|l|, |C|) > 99\%$ **then** ▷ probability of overlooking candidate with minimum number of points
      $\mathcal{P} \leftarrow \mathcal{P} \setminus C_l$
      $\mathcal{V} \leftarrow C \cup l$
      $\mathcal{S} \leftarrow c, r$
  for $l_i \in \mathcal{P}$ **do**
    for $\mathcal{S}_j \in C$ **do**
      if $Dist(l_i, c_j) \leq 2 \cdot r_j$ **then**
        $\mathcal{V} \leftarrow l_i$
  return $\mathcal{V}$
function NEWCANDIDATE($\mathcal{P}, \alpha$)
  $l_1 \leftarrow random(\mathcal{P})$ ▷ randomly selected first point
  $l_2 \leftarrow sameLevel(\mathcal{P})$ ▷ second point from same octree level
  $c \leftarrow midpoint(l_1, l_2)$
  $r \leftarrow \frac{||l_1 - c|| + ||l_2 - c||}{2}$
  $\Delta norm \leftarrow |normal(l_1) - normal(l_2)|$
  if $\Delta norm \leq \alpha$ **then**
    $C \leftarrow candidate$ ▷ candidate sphere with center c and radius r

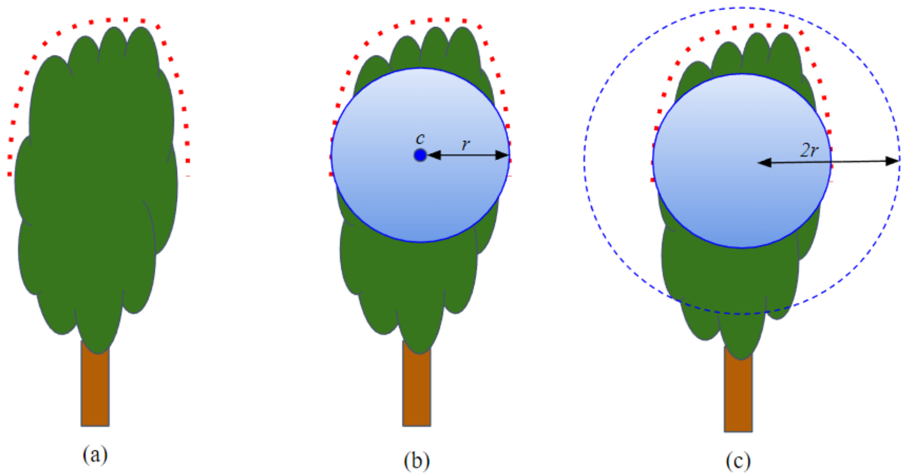

Figure 2.3: Identifying vegetation LiDAR points. Red LiDAR points associated with tall vegetation remain after mesh filtering in (a). RANSAC shape detection finds a blue sphere with center c and radius r that contain some of the LiDAR points (b). Expanding the radius 1.5 times includes more vegetation points (c).

normals of these points is less than or equal to a threshold, α, the candidate sphere gets a score. The score depends on the density and connectivity of points along the surface of the candidate sphere. The candidate sphere is accepted if there is a 99% probability that it is the best candidate based on number of points it contains and the running quantity of candidates. Points, l, associated with this sphere candidate, C, are removed from the point cloud, $\mathcal{P}$, and added to the set of vegetation points, $\mathcal{V}$. The center and radius define a sphere $\mathcal{S}$. Searching for candidate spheres stops when the probability of missing a sphere based on the number and size of sphere candidates is less than 1%.

As shown in Figure 2.3, this sphere does not necessarily contain all points associated with the tall vegetation. To capture points from these higher branches, we expand the vegetation set, $\mathcal{V}$ to include remaining LiDAR points, l, that are within 2 times the radius of each sphere's center (line 11 of Algorithm 2).

2.2.2.2 Filtering by proportional height

After identifying spherically aligned protrusions, low and wide point groups remain. We adapt a 3D connected components approach [78] to conglomerate these remaining points. The approach starts with the points arrayed across the previously described octree structure. The octree is considered as a stack of 2D slices. The cells at each height of the octree are checked for connectivity in 2D using a top down pass followed by a bottom up pass and equivalence table labelling [79]. Next, these labelled octree slices are compared to orthogonal neighbors. Points that are connected across slices are grouped into the same clouds. Finally, we add a filtering step (shown in 3) that evaluates the proportional height of each bounding box b in the set of connected component bounding boxes

$\mathcal{C}$. If the bounding box around the connected component's height b_z is greater than the mesh particle spacing, s, or 8 m, whichever is less, and if the hypotenuse of the depth b_y and width b_x is less than the height, the associated points are added to the set of potential vertical structures V.

Algorithm 3 Proportional height algorithm

```
procedure PROPORTIONAL HEIGHT FILTER(C, s)
    for b ∈ C do                    ▷ for every connected component's bounding box
        if b_z ≥ s OR b_z ≥ 8 then
            e ← sqrt(b_x^2 + b_y^2)
            if b_z > e then
                V ← b
    return V
```

2.2.3 Overlap Algorithm

Each tile of the DALES data set has uniform, 400% coverage. An aerial vehicle flying among the terrain cannot expect this repeated and overlapping exposure from different perspectives. Aerial vehicles would also have few returns from objects that are further away. A raw point cloud would also contain false returns and other outliers that were manually removed from this data set. For cases where the obstacle points are realistically sparse, we cannot depend only on continuous clustering to find vertical structure candidates. We adapt our previous work [1] to automatically detect vertical obstacles that remain in a realistically sparse and cluttered points.

To process sparse and cluttered point clouds, we adapted aspects of existing segmentation approaches to build an Overlap algorithm and a layered Density-Based Spatial Clustering of Applications with Noise (DBSCAN) Neighbor clustering algorithm. The Overlap algorithm starts by initially clustering points based on their spacing. Then we characterize the basic characteristics of each initial cluster. Next, we determine the amount of overlap between clusters, determine whether overlapping clusters are adjacent, and return these correlated points. This approach is shown in Figure 2.4. Unlike previous approaches, these algorithms do not require prior knowledge of the area of a vertical object, recognize a variety of shapes and textures, and are resilient (require minimal tuning).

2.2.4 Overlap Correlation

Our initial clustering step judiciously uses a simple and fast algorithm to filter isolated returns and reduce the number of objects that will be considered by subsequent processes. We use the DBSCAN algorithm [20] due to its relatively fast run time[75], two simple parameters (minimum number of points and distance between points, eps), and ability to cluster points associated with odd

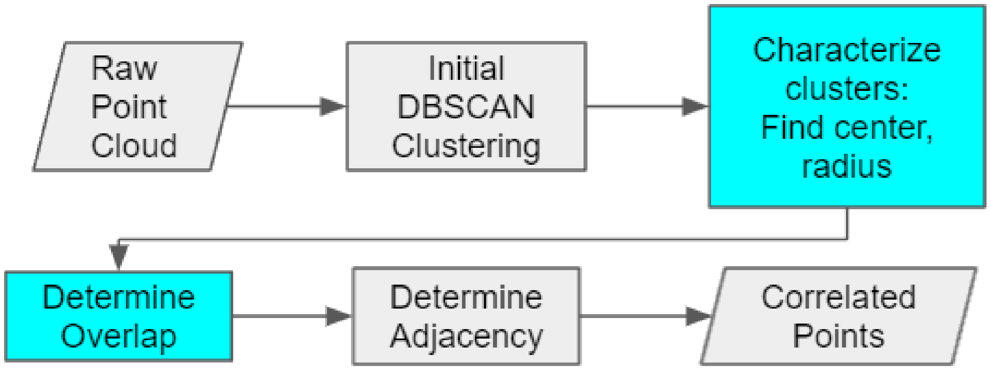

Figure 2.4: Overlap Algorithm Flowchart.

shapes. DBSCAN requires the definition of the distance between points in a neighborhood (eps) and minimum number of points in that neighborhood to declare that it is a cluster. To address the known sparsity of points in our situation, we set the minimum points to two. This decision is intended to remove singular range wrap points while still allowing consideration of occasional structural returns. The eps distance was set to 0.5 meters. We chose this value since it was a reasonable vertical footprint per Ref. [28] from our simulated aerial LIDAR at its maximum range of 100 meters, based on multiple simulations. This distance setting conservatively associated structural returns on the same raster line. The primary benefit of initial clustering is quickly reducing the number of points under consideration by three orders of magnitude (from 43k to 73 for one case).

In preparation for our rule-based Overlap algorithm, we next calculate the three-dimensional center of each cluster. We also find the standard deviation of cluster points with respect to each axis. We use the largest standard deviation to establish a radius from the center of each cluster. Since a cluster associated with the narrow portion of a vertical object may contain as few as two very close points, we add 0.5 meters to the radius to encourage a reasonable overlap. This analysis captures the key information about each cluster needed by subsequent algorithms that are more nuanced.

The overlap correlation algorithm associates clusters defined in the initial clustering step. It was inspired by the rule-based ring- and block-stacking approaches described previously in Refs. [22], [21], [24], [25] and [23]. Our approach reduces tuning parameters to a single overlap threshold variable. Instead, it is informed by the inherent cluster characteristics found in the previous step. This approach seeks to associate clusters with radii that overlap from the top-down view. The methodology for determining the area of overlapping circles in lines 14-17 was adapted from an example in Ref. [80].

The LabelOverlap procedure populates a dictionary for every cluster permutation in cluster list L_c with a corresponding intersection area a_i. The cluster list was populated with the (x, y) coordinate of the center and the radius r calculated in the previous cluster characterization step. The OLArea function determines the intersection area a_i between each cluster permutation based on

Algorithm 4 Overlap algorithm

procedure LABELOVERLAP(L_c)
  $Dict \leftarrow zeros(L_c \times L_c)$
  for $Pt_0 \in L_c$ **do**
    $Dict(Pt_0) \leftarrow L_c(Pt_0)$
    for $Pt1 \in L_c$ **do**
      $Dict(Pt_0, Pt_1) \leftarrow OLArea(Pt_0, Pt_1)$
function OLAREA($x_0, y_0, r_0, x_1, y_1, r_1$)
  $d \leftarrow \sqrt{(x_0 - x_1)^2 + (y_0 - y_1)^2}$
  if $d \leq |r_0 - r_1|$ **then** ▷ circle is contained
    $a_i \leftarrow 1$
  else if $d \geq r_0 + r_1$ **then** ▷ no overlap
    $a_i \leftarrow 0$
  else
    $\alpha \leftarrow \cos \frac{d^2+r_1^2-r_0^2}{2 \cdot d \cdot r_1}$
    $\beta \leftarrow \cos \frac{d^2+r_0^2-r_1^2}{2 \cdot d \cdot r_0}$
    $a \leftarrow r_1^2 \cdot \alpha + r_0^2 \cdot \beta - 0.5 * (r_1^2 \cdot \sin(2 \cdot \alpha) + r_0 \cdot \sin(2 \cdot \beta))$
    $a_i \leftarrow \frac{a}{\pi \cdot \min(r_0, r_1)^2}$ ▷ normalize
  return a_i

x_0, y_0, and r_0 for Pt_0 and x_1, y_1, and r_1 for Pt_1. The value of a_i ranges from 0 to 1. Each cluster pair combination is visited twice. We consider the (x, y) correlation as vertical obstacles, which are basically projections from the ground plane. This cluster-by-cluster calculation is completed quickly because it only needs to consider a small quantity of clusters deemed significant by the initial clustering process. The final step of the Overlap approach batches adjacent clusters. If the overlap value between clusters in the cluster dictionary is greater than the threshold (0.01 for this study), those clusters are recursively associated. A low threshold value coupled with the 0.5 m radius addition in the cluster characterization step allows for association of irregular returns and obstacles that are not truly vertical. A point is grouped whenever its overlap value is larger than the threshold value for any other point in that group.

2.2.5 Neighborhood Clustering

DBSCAN Neighborhood clustering runs a second algorithm with a larger point spacing in place of the teal process blocks in Figure 2.4. Our goal was to consolidate the points that are associated with the same structure, e.g., the front and back of the right tower's open frame, so the depth correlation would not erroneously denote the sharp angle as belonging to different objects. The maximum distance between points d is:

$$d = R \cdot \tan \frac{\theta}{n} \cdot \%_{\text{buffer}} \tag{2.3}$$

where R is LIDAR range in meters, θ is elevation of the LIDAR window, n is the number of beams across the elevation window, and $\%_{\text{buffer}}$ is an allowance for uncertain point scatter. $\%_{\text{buffer}}$ was set to increase d by 25%. Since DBSCAN run time increases with eps spacing, this subsequent association after Initial Clustering takes advantage of the much smaller point cloud size.

2.2.6 Connected Components

Ref. [76] presents a connected components algorithm that discerns objects from an MLS. This algorithm determines angle of incidence from the sensor to a line connecting neighboring points in a range image. We modified Ref. [76] to discern the line of a narrow vertical obstacle since vertical obstacles tend to have a normal angle of incidence to aircraft. Whereas previous MLS primarily used a horizontal orientation, we had to consider 3D orientation of the β angle. Our sparse point cloud challenged us to infer depth in any direction due to the irregular vertical obstacle geometry. In place of the teal blocks in Figure 2.4, our adaptation of range-image-based connected components algorithm calculates complementary β angles for each initial cluster permutation.

Algorithm 5 Connected Components algorithm

procedure LABELCOMPONENTS(L_c)
  $Dict \leftarrow zeros(L_c \times L_c)$
  for $Pt_0 \in L_c$ **do**
    $Dict(Pt_0) \leftarrow L_c(Pt_0)$
    for $Pt1 \in L_c$ **do**
      $Dict(Pt_0, Pt_1) \leftarrow FindB(Pt_0, Pt_1)$
function FINDB(Pt_0, Pt_1)
  $\theta_0 \leftarrow Pt_0.\theta_{median}$
  $d_0 \leftarrow Pt_0.d_{median}$
  $\theta_1 \leftarrow Pt_1.\theta_{median}$
  $d_1 \leftarrow Pt_1.d_{median}$
  $\psi \leftarrow |\theta_0 - \theta_1|$
  $\beta \leftarrow \arctan \frac{d_1 \cdot \sin(\psi)}{d_0 - d_1 \cdot \cos(\psi)}$
  return β_i

2.3 Setting

An aerial vehicle must contend with a variety of environments and potential obstacles in each mission. These structures include electrical transmission towers, communication towers, light poles, utility poles, and cranes set among urban, commercial, suburban and rural scenes. Therefore, the evaluation environment should offer a range of expansive representative scenes and specifically labeled vertical objects. Due to the sparsity of points from vertical structures, most recent data sets that contain vertical structures do not differentiate them from the ground or road surface [81] [82] or combine them with other structures (”remaining hard scapes”[83], ”artifacts”[84], ”urban furniture” or ”vertical surface”[85], or ”street furniture”[86]).

2.3.1 Real world data set

The Dayton Annotated LiDAR Earth Scan (DALES) [67] is an Aerial LiDAR System data set that has the distinction of classifying points from various vertical structures in a range of scenes. Over 500 million points are distributed across 40 discontiguous tiles. Each tile has over 10 million points representing real world examples of urban, surburban, commercial and rural scenes. Most scenes contain utility poles and street lights, but not all of these vertical structures protrude from the surrounding terrain. For example, the tops of most utility lines are lower than treetops within a few meters. Twenty tiles with the most prominent protrusions were used for our analysis. Characteristics of those tiles are summarized in Table 2.2.

Table 2.2: Types of protrusions present in each scene

Scene Type	Number of Tiles	Transmission Tower	Communication Tower	Light Pole	Utility Pole	Crane
Urban	1	-	-	-	-	✓
Commercial	5	✓	✓	✓	✓	✓
Suburban	11	✓	-	✓	✓	-
Rural	3	✓	-	-	✓	-

A commercial tile example is shown in Figure 2.5. Less than 0.1% of the nearly 12 million points belong to the pole class. The logarithmic axis displays the dominance of ground, vegetation and building classes.

The data set presents a wide variety of actual vegetation and buildings. However, the DALES data set has some caveats for the aerial vehicle use case:

- Sensor perspective: The LiDAR was at the airplane’s nadir, increasing the number of returns from classes with perpendicular surfaces (typically every class except poles) due to incidence angle and low occlusion

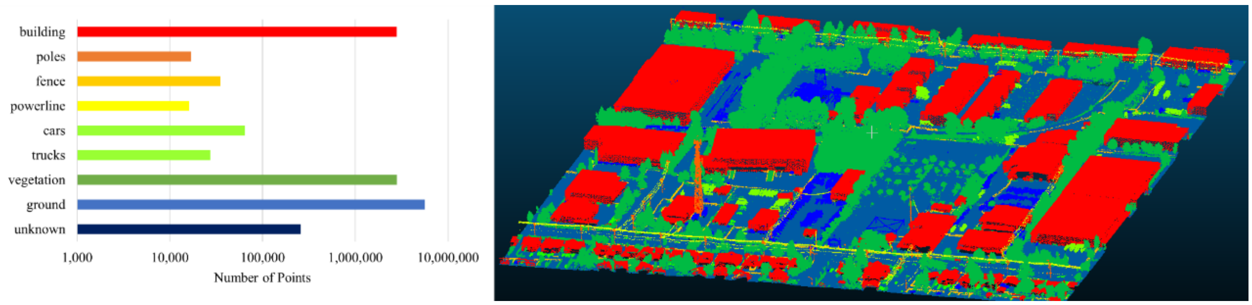

Figure 2.5: Distribution of points in a commercial scene.

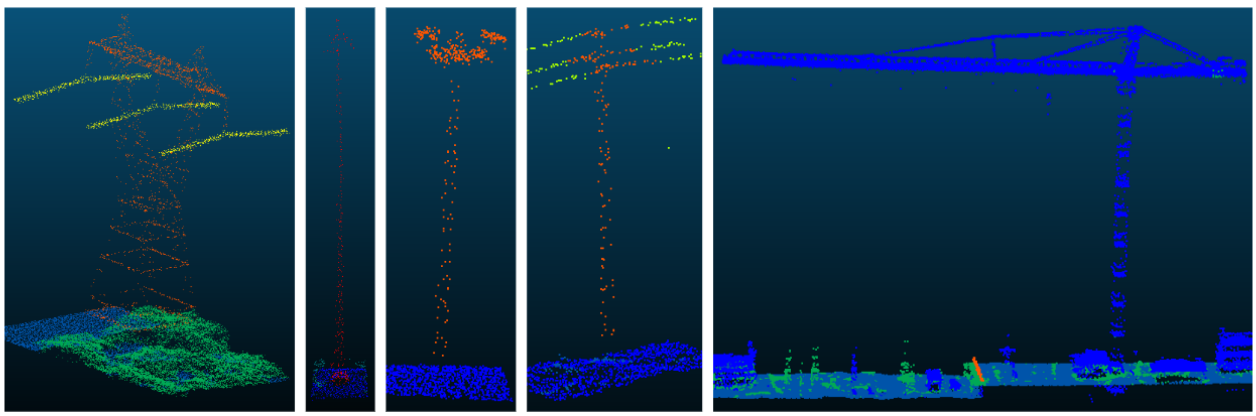

Figure 2.6: Examples (and corresponding height above ground) of transmission tower (28 m), communication tower (34 m), light pole (15m), utility pole (15 m) and crane (64 m).

- Greater than 400% sampling increased point density, especially for classes with significant horizontal surfaces
- Removal of outlying points prior to any analysis
- Inaccurate and incomplete classification: Cranes are labeled as ”unknown”, several prominent poles are erroneously labeled as ”ground” or ”vegetation”, power lines are labeled as poles and vice versus (some examples shown in Figure 2.32)
- Each 500 x 500 meter tile presents a large batch of points that a sensor on an aerial vehicle would not typically encounter simultaneously
- Except for a couple of communications towers and cranes, the most prominent obstacle on each tile would only be a hazard for vehicles flying within 10 meters of the treetop, rooftop, or other dominant terrain
- Most of the tiles with a protrusion are in a suburban setting

2.3.2 Simulated sparse and cluttered data set

Sparse and cluttered data to evaluate the overlap algorithm came from four tower combinations built in Unreal Engine simulation. We make several assumptions. First, we assume our point cloud comes from a LIDAR with minimal radiometric considerations or realities. For example, all surfaces in the model are perfectly reflective, i.e., each point has 100% value regardless of the surface type, angle, distance, or other considerations. We assume is no background radiation or other atmospheric interference, that all transmissions are received, and that calculated positions are accurate with no random or systematic error or uncertainty. Accuracy depends on the assumption that simulated vehicle position reporting (as the origination point of the laser beams) is accurate and precise. This work processes approximately 4 seconds of LIDAR data at a time, so we assume suitable performance on these samples will also be sufficient for online processing. We also assume the simulated tower structures will adequately imitate real world irregular returns.

Between eight and twelve 0.5 meter diameter floating spheres were added across the 3D volume to simulate range wrap and other sources of false positive LIDAR returns (Figure 2.7). Figures 2.8 and 2.9 show the fourth tower arrangement. Microsoft's AirSim quadrotor simulation provided a

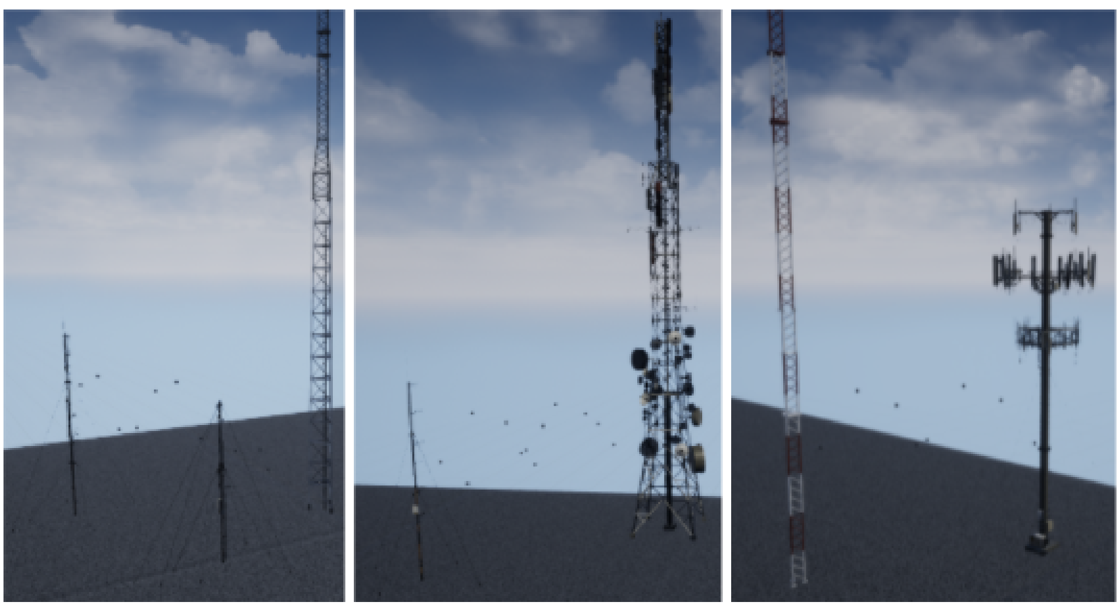

Figure 2.7: Case 1: Two short, thin towers in front of a tall tower (left); Case 2: One large, complex tower in front of a single thin tower (middle), and Case 3: Tall thin tower ahead of a cell tower (right)

basic plugin for the Unreal environment. We modified the LIDAR simulation to imitate a VLD-16 sensor and calculated distance and angle of incidence for each point received.

The AirSim quadrotor started just over 100 meters from the farthest vertical obstacle then proceeded forward at 2 m/s for 5 meters at an altitude of 10 meters. This acceleration and deceleration provided some amount of platform instability for the simulated LIDAR and resulted in more realistic scan lines. The starting distance close to VLD-16 maximum range ensured point cloud sparsity for the far towers. Figure 2.8 shows one starting point prior to takeoff. The bright green LIDAR collisions with the rear tower are just visible as the nearest portion of the back tower is just inside 100 m. The bottom panel of Figure 2.8 shows a top down view of the orientation of the towers and quadrotor. Tower case 4 and the corresponding point cloud is shown in Figure 2.9. The 68,677

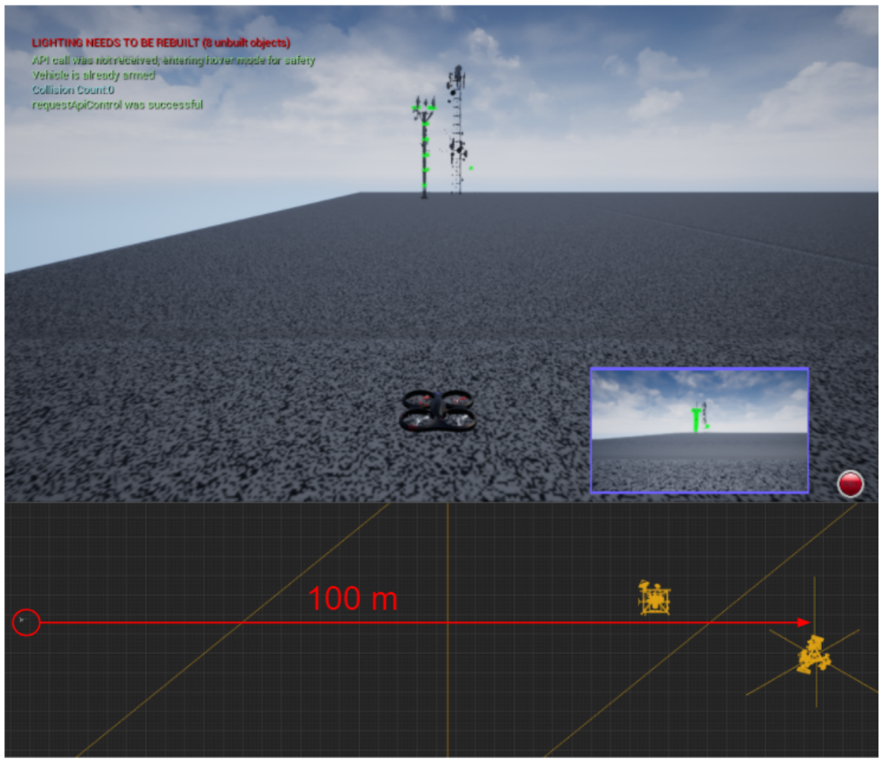

Figure 2.8: AirSim quadrotor start point facing towers (top); top view of start point (bottom).

points from the first 4,789 milliseconds of simulated flight are used to analyze our algorithmic approaches. We used Fowlkes-Mallows score [87] to compare effectiveness of DBSCAN and our Overlap clustering. Fowlkes-Mallows Index (FMI) is defined by:

$$FMI = \frac{TP}{\sqrt{(TP + FP)(TP + FN)}} \tag{2.4}$$

where TP is number of True Positives, i.e., number of point clusters that belong in both true and predicted labels, FP is the number of False Positives, i.e., the number of point clusters predicted to be a part of the tower structure but are not. FN is number of False Negatives, i.e., the number of point clusters that should have been associated with a tower structure but are classified as detached points. The FMI score range is from 0 to 1. A score of 1 indicates all points are correctly associated with a tower structure.

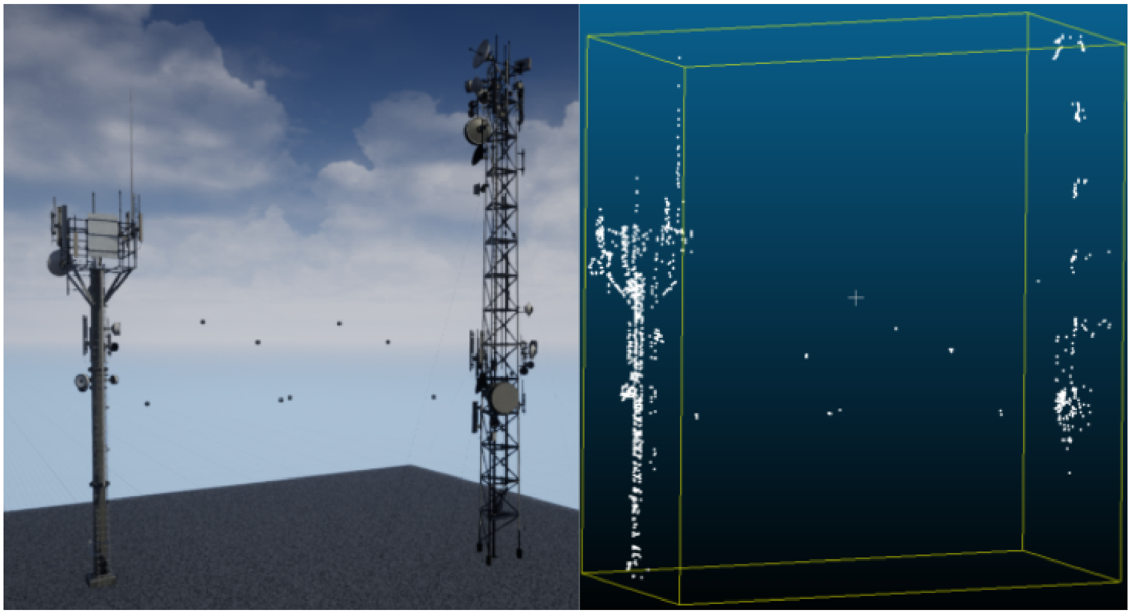

Figure 2.9: Tower arrangement 4 (left); point cloud (right)

2.4 Experimental Results

The DALES data set offers a variety of voluminous LiDAR scenes dominated by ground, vegetation, and building points. We seek to distill each scene into a more palatable subset. First, we find the best mesh particle spacing for each scene by maximizing the density and quantity of pole points using a new prevalence metric. However, many pole points in the DALES data set are associated with short utility poles and light poles. These less hazardous structures are frequently shorter than nearby buildings, trees and other large objects. We introduce the prominence metric to evaluate the effectiveness of the coarse mesh in identifying hazardous obstacles in each scene.

Next, we analyze the remnants that extend above the mesh. We attempt to remove treetops and other large objects by spherical association and connected component clustering. At the end of this stage, we hope to have isolated all the significant vertical obstacles.

Finally, we apply the overlap algorithm to some examples of vertical obstacles that were not identified with the proportional height filter and evaluate performance against the simulated sparse data set.

2.4.1 Maximizing prevalence with the mesh

The coarse mesh is designed to drape over low objects while allowing tall objects to pass through. Our first task is to investigate which particle spacing maximizes pole points while minimizing the intrusion of points from irrelevant objects. We assume that a higher number of pole points increases the height difference between points on a given obstacle due to the descending nature of the mesh. The most effective mesh spacing results in a high concentration of pertinent points while also considering the quantity of relevant points remaining. Therefore, we judge the goodness of mesh particle spacing based on the prevalence, P, of obstacle points in the remaining non-ground points:

$$P = \frac{|l|_{\text{obstacle}}}{|l|_{\text{total}}} \cdot 100 \cdot |l|_{\text{obstacle}} \tag{2.5}$$

where $|l|_{\text{obstacle}}$ is the quantity of obstacle points (tower, pole, or crane points) and $|l|_{\text{total}}$ is the quantity of all remaining non-ground points.

The particle spacing was varied on the 20 tiles with the tallest protrusions. Grid spacing investigations started at 1 m, followed by 2 to 12 m in 2 m increments. The best grid spacing corresponded to the largest P. For tiles where 1 or 12 m spacing resulted in the largest P, 0.5 and 14 m grid spacing, respectively, was also evaluated to ensure that smaller or larger spacing would not increase P.

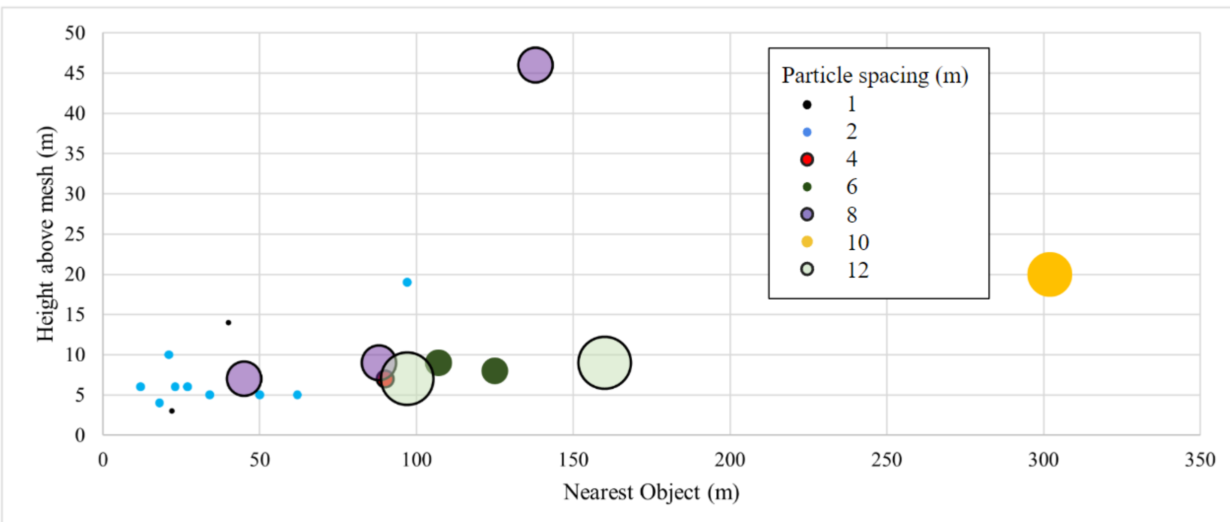

Figure 2.10: Best grid spacing for each combination of obstacle distance and height above mesh. The largest bubbles correspond to 12 meter grid spacing while the smallest bubbles correspond to 1 meter spacing.

Figure 2.10 shows the height and horizontal spacing for the tallest object in each of the twenty selected scenes. Larger circles indicate a larger particle spacing to maximize prevalence. One or two meter particle spacing results in the highest prevalence for most tiles. However, these tiles with smaller grid spacing also have fewer significant protrusions. Grid spacing between 6 and 12 m seems to maximize P in tiles with larger vertical structures. We delve into describing prominence in the next section.

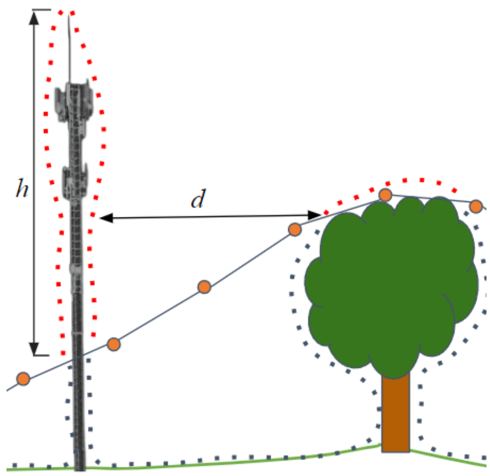

Figure 2.11: Prominence variables. The height, h, is the difference in elevation between the lowest and highest obstacle protruding points. The distance from the nearest object, d, is the horizontal distance from the nearest protruding object.

2.4.2 Considering mesh effectiveness with prominence

Each tile section contains a large variety of vertical objects, ranging from short fence posts to tall communications towers. The significance of a vertical object depends on it's altitude above nearby objects along with the object's proximity to those objects. We propose a metric of prominence, ψ, that succinctly describes the significance of an object by referencing the surface formed by the course mesh:

$$\psi = h \cdot \sqrt{d} \tag{2.6}$$

where h is the difference between the mesh penetration altitude and highest point on the vertical object and d is the distance from the closest object at the altitude where an object penetrates the mesh. In cases where the mesh settles onto the ground, h is the distance to the nearest non-ground object. These variables in a notional scene where the coarse mesh has settled onto a tree and tower are shown in Figure 2.11.

Communication towers are, unsurprisingly, the most prominent vertical structure type due to their height and typical distance from other obstacles. This data set only contains four examples, however, so we extend our investigation to the other types of vertical structures. Some electrical transmission towers are also prominent as they tend to extend high above the ground to support high tension power lines. However, the power lines in this data set are densely represented and usually collide with the descending mesh particles, as will be discussed later. Cranes are another tall object that tends to have a lower prominence score. This is largely due to the larger cross section from the airborne LiDAR's nadir which increases the density of points in the horizontal plane. Light poles are slightly more prominent than utility poles, even though they have similar heights. However, light

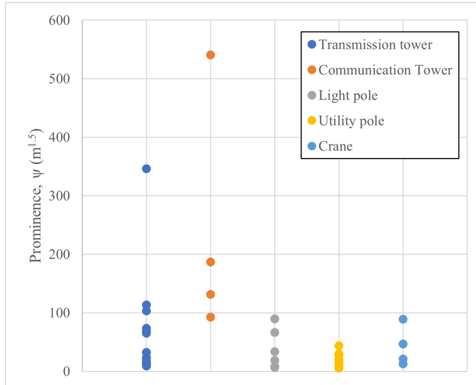

Figure 2.12: Prominence for each type of protrusion.

poles, unlike utility poles, usually do not support wires. Similar to transmission towers, the density of these points is increased by greater than 400% coverage from the airborne nadir perspective.

Figure 2.12 shows the prominence of each protrusion type in the various scenes. After finding the best grid spacing for each scene, the resulting protrusions were manually identified. The protrusion height above mesh and horizontal distance from the nearest object (typically vegetation or buildings) was measured. When there were multiple similar protrusions with few variations in height or geometry, such as a line of utility poles or grouping of light poles, the average protrusion height and nearest object distance was recorded. Communication tower, crane and transmission tower protrusions are characterized individually.

We present examples of each vertical structure in the following paragraphs. Two communication towers from the commercial scene in Figure 2.5 are highlighted in Figure 2.13. The taller 65 m tower on the left has the highest prominence, ψ, among all tiles in the DALES data set. In addition to its height, the tower is over 130 m away from the nearest tree top. The shorter communication tower in Figure 2.13 is approximately half as tall and is twice as close to surrounding tree tops that protrude through the mesh. Although the shorter tower has smaller cross section, its protrusion height and distance from trees results in the fourth highest prominence, ψ, across the entire data set. Only two 12 m tall utility poles running along the top of the scene penetrate the mesh due to mesh collisions with taller treetops.

Only one other scene contains communication towers. A 20 m tall communication tower ($\psi = 187m^{1.5}$) is much more dominant than a group of utility poles in Table 2.3. Close 2 m mesh particle spacing increases density and prevalence by 1,488% and 506%, respectively. As shown in Figures 2.5 and 2.13, the tallest 65 m communication tower ($\psi = 540m^{1.5}$) contains a large number of points. These plentiful returns result in the highest pre-mesh pole point density (0.15%) and prevalence ($P = 2,459$) (with the exception of scenes containing cranes, presented in Table 2.5). Wider mesh spacing of 8 m maximizes pole points in this second scene. The rich returns from the

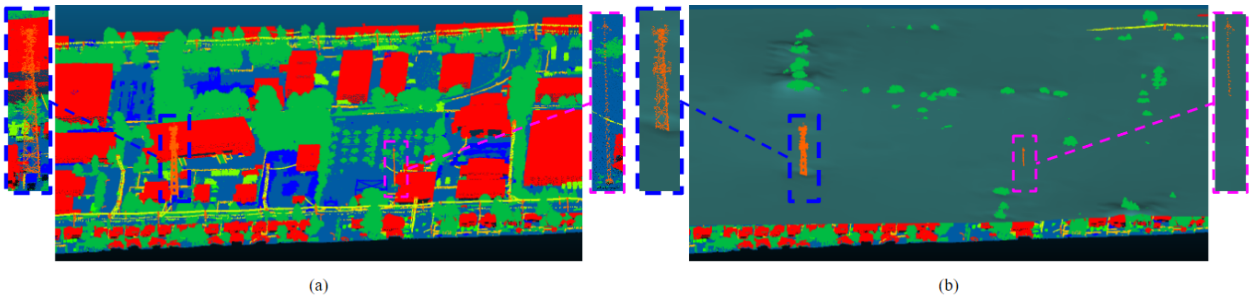

Figure 2.13: Two communication tower protrusions. The 65 m communication tower (blue dashed box) and 34 m communication tower (magenta dashed box) are shown in the unfiltered scene in (a). After applying a 12 m mesh, 46 m of the taller communication tower remain while to top 16 m of the shorter communication tower protrudes in (b)

taller and wider communication tower in the second line reduce the need for the mesh to settle among shorter protrusions.

Table 2.3: Mesh results for prominent communication towers (UP = utility pole, TT = transmission tower, LP = light pole, CR = crane, CT = communication tower). The 65 m communication tower ($\psi = 540$) and other aspects of Figure 2.13 are shown in blue.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Original density, %	Original prevalence, P	Best grid spacing, m	Post-mesh density, % (change, %)	Post-mesh prevalence, P (change, %)
20	-	-	-	187	0.08	770	2	1.27 (1,488)	4,667 (506)
16	-	-	-	540	0.15	2,459	8	4.52 (2,913)	21,074 (757)
				Average	0.12	1,615		2.90 (2,200)	12,871 (632)

A more commercial scene in Figure 2.14 displays examples of electrical transmission towers and a construction crane. Mesh particle spacing of 12 m results in the largest P (prevalence) of pole and crane points. The 16 m utility poles and even shorter light poles that run along the streets at the center of the scene do not reach the mesh due to the taller crane and transmission towers. The top 10 m of the 39 m transmission tower breach the mesh. Only the top 2 m of the crane show through. The taller protrusion of the transmission tower, combined with the greater distance from tree tops, makes the tower ($\psi = 114m^{1.5}$) nearly ten times as prominent as the crane ($\psi = 13m^{1.5}$, both in Table 2.4).

Grid spacing for the seven scenes with transmission towers varies widely in Table 2.4. Larger transmission tower prominence generally benefits from wider mesh spacing, although each scene offers a unique challenge. The top three rows describe scenes with less prominent utility poles ($\psi = 7m^{1.5}$ and $\psi = 17m^{1.5}$) and transmission towers ($\psi = 24m^{1.5}$). These three scenes maximize

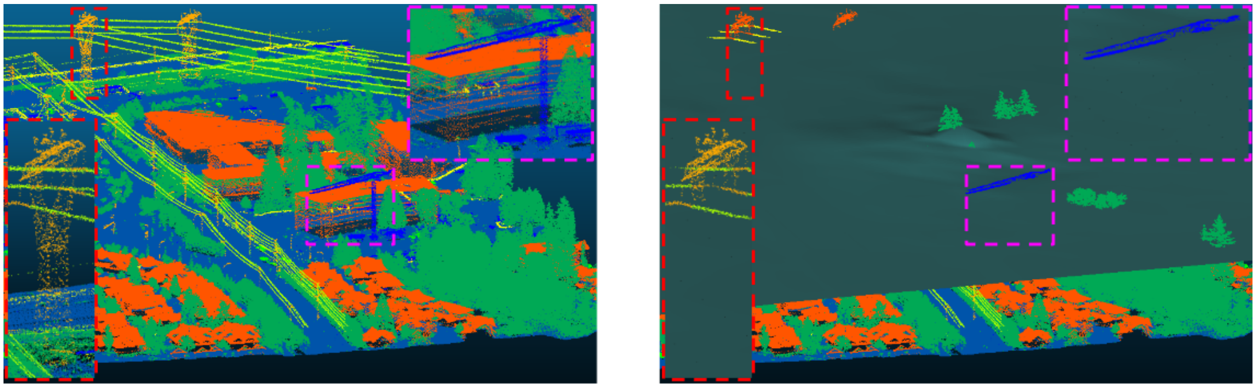

(a) (b)

Figure 2.14: Examples of transmission tower and crane protrusions. A 39 m tall transmission tower (dashed red box) is shown next to a 29 m crane (dashed magenta box) without any filtering in (a). After applying a 12 m coarse mesh, (b) shows the resulting protrusions.

pole point prevalence with close 2-6 m grid spacing. The mesh spacing in the last row was heavily influenced by a crane ($\psi = 13m^{1.5}$) in the same scene (Figure 2.14) which drastically increased the original density and prevalence of vertical structure points. The high density of power lines points emanating from transmission towers (as shown in Figure 2.14) resulted in collisions with the mesh particles, halting further mesh descent. Despite the interference of these large horizontal arrays, applying the mesh increased average density and prevalence by 4,292 and 1,242%, respectively.

Table 2.4: Mesh results for prominent transmission towers. The scene containing the 39 m tall transmission tower ($\psi = 114$) and crane from Figure 2.14 are shown in orange.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Original density, %	Original prevalence, P	Best grid spacing, m	Post-mesh density, % (change, %)	Post-mesh prevalence, P (change, %)
-	24	-	-	-	0.08	889	2	1.59 (1,888)	3,843 (332)
7	33	-	-	-	0.08	852	2	1.4 (1,650)	1,973 (132)
17	103	-	-	-	0.12	2,015	6	5.39 (4,392)	15,329 (661)
-	74	-	-	-	0.06	579	8	5.06 (8,333)	17,071 (2,848)
-	346	-	-	-	0.04	160	10	4.11 (10,175)	7,581 (4,638)
-	69	-	-	-	0.11	1,620	12	2.95 (2,582)	3,778 (133)
-	114	-	13	-	1.29	199,815	12	14.52 (1,026)	99,585 (-50)
				Average	0.25	29,419		5.00 (4,292)	21,309 (1,242)

Cranes only dominate two of the four scenes in which they are present. Nearly 80% of returns from the 37 m crane in Figure 2.15 come from the top 7 m which contains the horizontal boom. The abundance of points on the cranes' horizontal booms in the three other scenes also prevents protrusion of the thin vertical support structure. This top heavy density results in a modest change and, in the case of the most prominent crane ($\psi = 90$), a decrease in post-mesh prevalence after the mesh adheres to the boom. Table 2.5 shows mixed results when the mesh filter is applied to crane scenes. Optimal mesh spacing increases density by 1,506% but decreases prevalence by 29%.

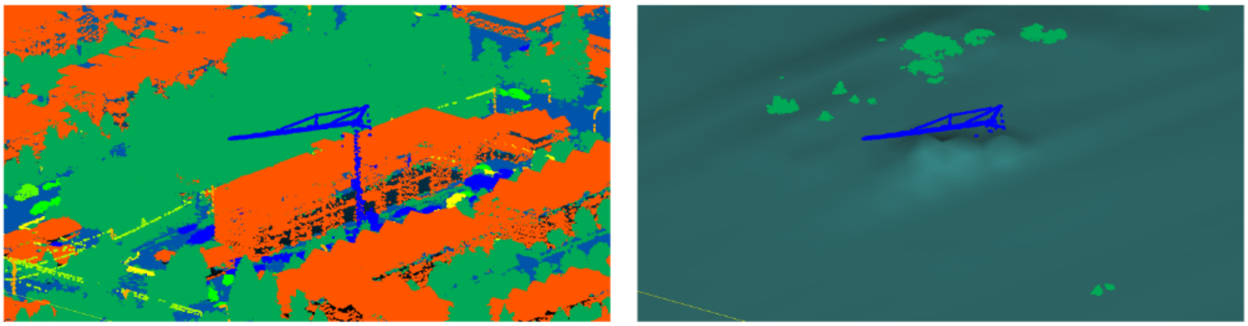

(a) (b)

Figure 2.15: Prior to applying the mesh to a 37 m tall crane (a), blue returns from the vertical structure are visible. Mesh particles collide with the horizontal boom in (b), largely removing evidence of the vertical structure.

Table 2.5: Mesh results for prominent cranes. Results from the scene containing the 37 m crane in Figure 2.15 are shown in blue.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Original density, %	Original prevalence, P	Best grid spacing, m	Post-mesh density, % (change, %)	Post-mesh prevalence, P (change, %)
30	-	-	90	-	3.17	1,177,721	6	28.86 (810)	331,176 (-72)
-	-	-	47	-	0.47	27,782	8	10.82 (2,202)	31,829 (15)
				Average	1.82	602,752		19.84 (1,506)	181,503 (-29)

Figure 2.16 shows examples of light pole and utility pole protrusions in a suburban scene. Due to the sparsity of high pole points, this scene maximized prevalence with a mesh particle spacing of 1 m. Most points from the 17 m tall light pole in the parking lot protrude through the mesh. The mesh also allows shorter 11 m utility poles to remain in the right side of the scene. The 17 m light pole is three times as prominent as the utility poles which protrude between 4 and 6 m. The light pole prominence is also greater because it is approximately 45 m away from the unclassified dark blue objects while the utility poles are within 12 m of the orange rooftops.

Results for the scene in Figure 2.16 are in the first row of Table 2.6. Scenes with the most prominent light poles have the lowest original prevalence. This sparsity combines with the generally low light pole prominence to require close mesh particle spacing that closely ”clings” to the terrain. Average density and prevalence above the mesh increases by 2,057 and 902%, respectively.

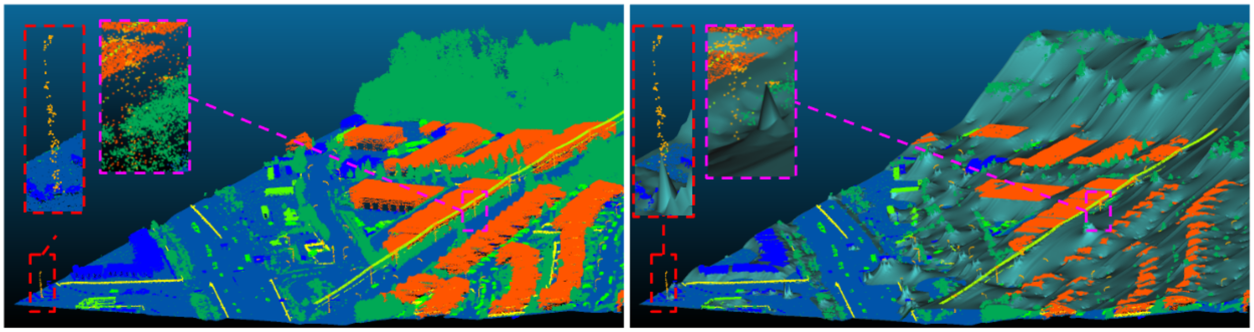

(a) (b)

Figure 2.16: Examples of light pole and utility pole protrusions. A 17 m tall light pole (dashed red box) and line of 11 m tall utility poles (dashed magenta box) without any filtering in (a). After applying a 1 m coarse mesh, (b) shows the resulting light pole and utility pole protrusions.

Table 2.6: Mesh results for prominent light poles. A quantitative description of the scene in Figure 2.16 is in the orange row.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Original density, %	Original prevalence, P	Best grid spacing, m	Post-mesh density, % (change, %)	Post-mesh prevalence, P (change, %)
30	-	90	-	-	0.03	133	1	0.28 (833)	683 (414)
6	-	19	-	-	0.05	280	2	0.78 (1,460)	846 (202)
7	-	34	21	-	0.06	396	2	2.52 (4,100)	11,898 (2,905)
-	-	66	-	-	0.09	1,045	4	1.74 (1,833)	1,942 (86)
				Average	0.06	464		1.33 (2,057)	3,842 (902)

Utility poles are usually the shortest vertical structure type in the DALES data set. The diminutive utility poles are aptly described by the corresponding small prominence values (Table 2.7). Utility poles also present the challenge of associated power line points which, in this data set, collide with the mesh. Figure 2.17 shows a case where a line of utility poles with dense power line returns typical for this data set protrude much less than utility poles in the same scene with sparse power line points. As with light poles, scenes with these subtle structures require close mesh spacing to maximize average density and prevalence.

Table 2.7: Mesh results for prominent utility poles. The scene with the most prominent utility poles ($\psi = 44m^{1.5}$) is described in the green row.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Original density, %	Original prevalence, P	Best grid spacing, m	Post-mesh density, % (change, %)	Post-mesh prevalence, P (change, %)
10	-	-	-	-	0.05	262	1	0.25 (400)	296 (13)
27	-	-	-	-	0.06	473	2	0.33 (50)	482 (2)
44	-	-	-	-	0.08	682	2	2.02 (2,425)	8,267 (1,112)
10	-	-	-	-	0.08	822	2	0.97 (1,113)	1,529 (86)
29	-	8	-	-	0.12	1,718	2	1.59 (1,125)	3,122 (82)
				Average	0.08	791		1.03 (1,123)	2,739 (259)

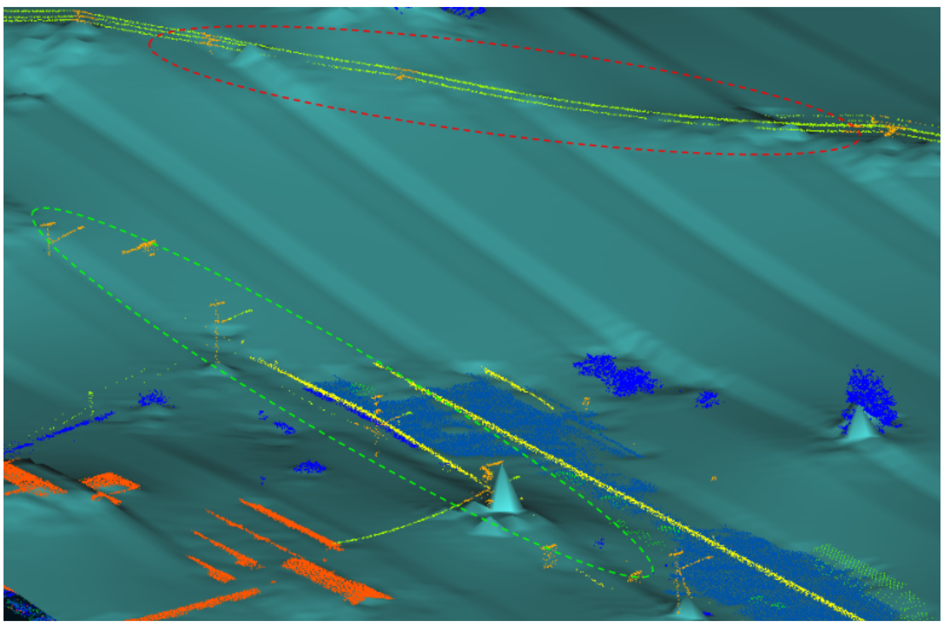

Figure 2.17: A line of utility poles with prominent power line returns (red dashed oval) protrude more than utility poles with less pronounced power line points (green dashed oval).

Most of the scenes with prominent structures maximize prevalence with particle spacing $\geq 6m$. Figure 2.18 shows the general trend that scenes with taller objects benefit from a more space between mesh particles.

The height above the mesh for the most prominent obstacle on each tile is greater than the space between grid particles in all but three instances (Figure 2.18). The two orange points that are less than the 12 m grid spacing are two transmission towers with particularly dense power line points. Despite the large particle spacing, particles collide with the nearly continuous power line returns before reaching the bottom of their catenary curve. The protrusion that is less than its 8 m grid

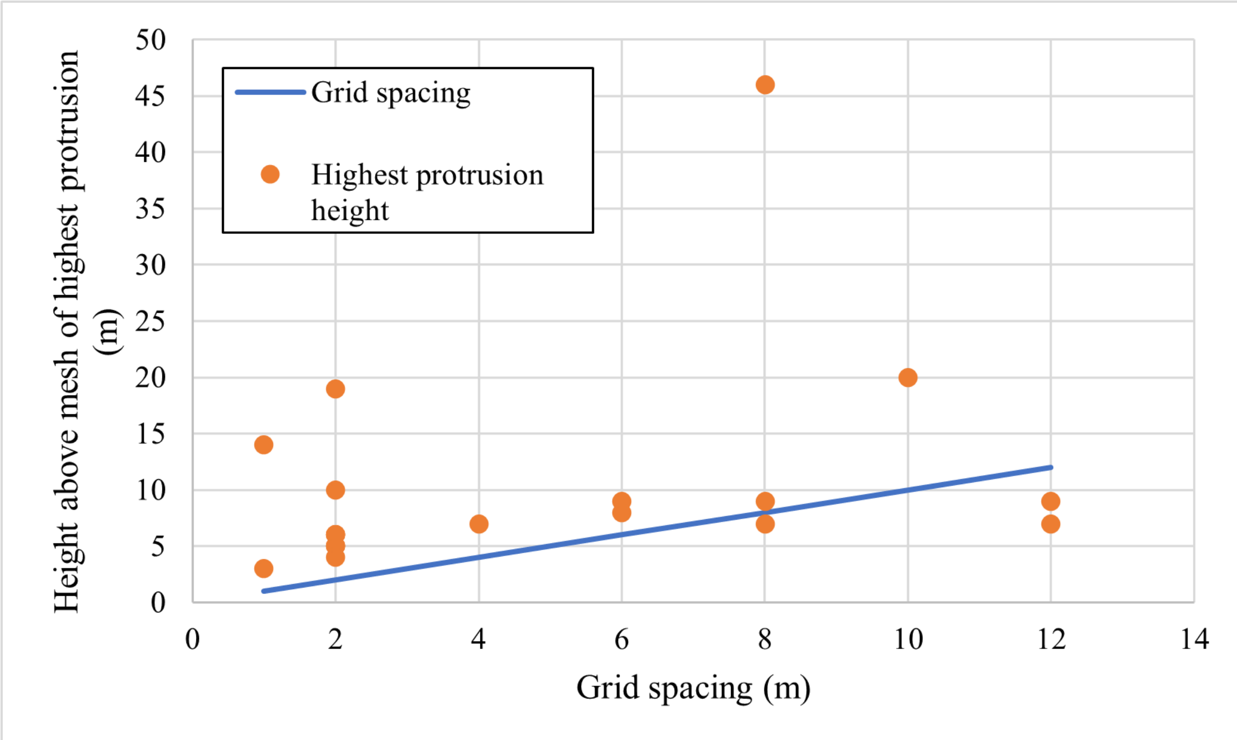

Figure 2.18: Height above mesh as a function of grid spacing for each tile's most prominent protrusion.

spacing is a crane with a high density of points on the horizontal surface of its boom (Figure 2.15). Cranes that protrude in other scenes also have reduced prominence because of the dominance of returns from their horizontal structure.

The time required to apply the mesh filter is highly dependent on the mesh particle spacing. Most tiles were 500 by 500 m squares. For $s = 1m$ mesh particle spacing, this meant that 250,000 particles were rasterized and iterated, taking an average of 7.7 seconds. For $s = 2m$, it only took an average of 1 second to conduct the same process with approximately 62,500 particles. The time required for mesh spacing between 4 and 12 m decreased to an average 0.4 seconds as the number of particles plummeted.

Overall, the mesh filter is effective at increasing density and prevalence in most scenes. Scenes with less prominent structures (such as those with only utility and light poles) require closer mesh particle spacing to maximize prevalence. The finer mesh allows more particles to descend until colliding with the terrain, as opposed to being constrained by immovable neighbors. This snug final position results in more non-pole points penetrating through the mesh. The average vertical structure prevalence, P, after applying the mesh filter for scenes dominated by utility (2,739) and light poles (3,842) is significantly less than more prominent cranes (181,503), transmission towers (21,309), and communication towers (12,871). The mesh filter increases average density (by 2,679%) and prevalence (by 778%) for the 20 tiles even though only a quarter of them contain high prominence structures ($\psi > 100m^{1.5}$).

2.4.3 Clustering

Even though the mesh filter drastically increases pole point prevalence, points from objects other than vertical structures still compose more than 95% of LiDAR points in most remaining point clouds. Vegetation is, by far, the most numerous source of returns. We start by removing spherically oriented points, then proceed with the proportional height filter to remove remnants of wide man made objects, such as rooftops, and remaining tree tops.

A filter for minimum height above mesh is not sufficient on it's own since, in most cases, there are vegetation protrusions. Many of these remaining tree point arrangements tend to have intrinsic sphericity. Figure 2.19 shows the top of a 33 m tall tree set among orange buildings. After applying mesh with 2 m particle spacing, most tree points are removed from the scene. However, a significant number of points (2,491) remain above the mesh. The remaining points for this tree have a vertical range of 14 m, which could be considered a significant obstacle if not for the associated 5-10 m width. The RANSAC approach detects a 10 m sphere. Doubling the radius encompasses all of these vegetation points.

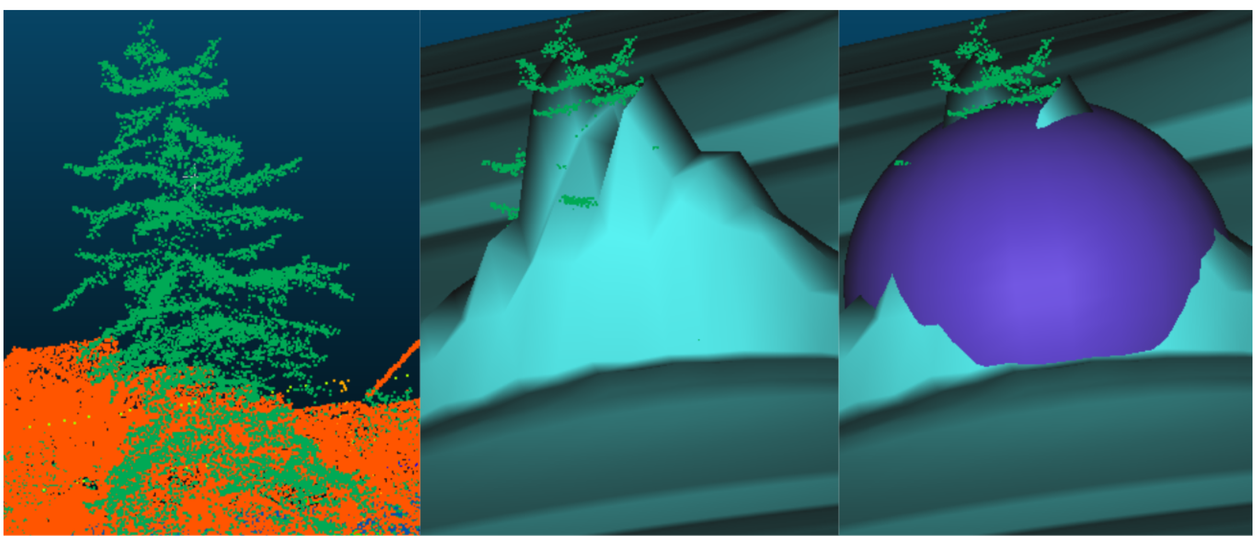

(a) (b) (c)

Figure 2.19: A 33 m tall tree protrudes above surrounding orange buildings in (a). In (b), a significant number of points protrude through the mesh filter. The RANSAC sphere shape detection process finds a purple 10.07 m sphere embedded among the tree points (c).

The subsequent step applies the proportional height filter. This filter checks point clusters associated according to the dimensions of the bounding boxes for the 3D connected components.

Table 2.8 shows that proportional height filtering greatly increased the prevalence of communications towers. The two prominent towers shown in Figure 2.13 successfully pass the criteria while low rooftops and trees are removed. The results for a scene with two less prominent towers is shown in Figure 2.20. In addition to the two towers encompassed by yellow bounding boxes in the foreground, there are also two barely perceivable annotations for a shorter light pole and building antenna in front of the yellow power line. This section of the scene also shows 6 tree tops that did not have sufficient points density for the sphere detector that were also taller than the 2 m mesh particle spacing. Nevertheless, the communication tower point prevalence and density showed a significant average increase (1,867 and 1,272%, respectively).

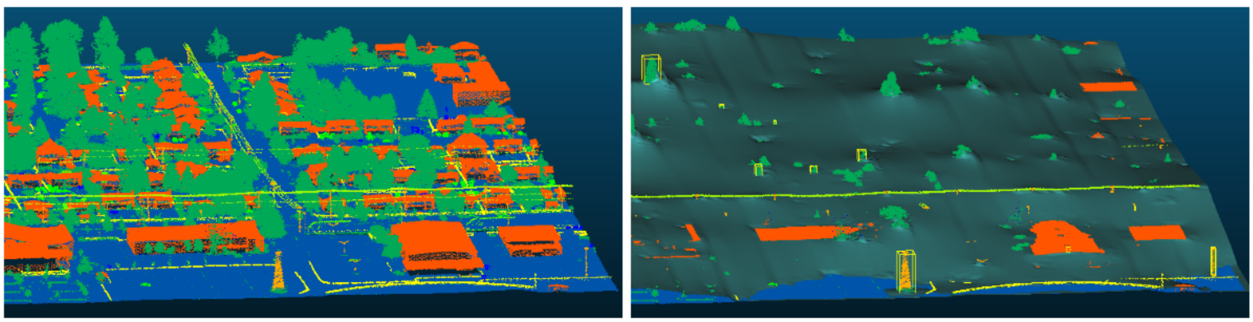

(a) (b)

Figure 2.20: Two communication towers (19 m and 14 m tall) are represented by orange points in the foreground of the unfiltered point cloud in (a). In (b), after applying the mesh and proportional height filters, the two towers are not the only objects highlighted by bounding boxes.

Table 2.8: Communication tower clustering results. Quantitative results from the scene in Figure 2.20 are shown in orange.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Post-mesh density, %	Post-mesh prevalence, P	Number of clusters (vegetation)	Post-clustering density, % (change, %)	Post-clustering prevalence, P (change, %)
20	-	-	-	187	1.27	4,667	19 (12)	21.86 (1,621)	35,432 (659)
16	-	-	-	540	4.52	21,074	2 (0)	100.00 (2,112)	418,200 (1,884)
				Average	2.90	12,871		60.93 (1,867)	226,816 (1,272)

Cranes were not accentuated by the proportional height filtering step. In both scenes where cranes were the most prominent vertical structure, no points passed the proportional filter check. It is worth noting, however, that in these same scenes there were no false positive bounding boxes around vegetation or other non-vertical objects. The spherical detection (shown in Figure 2.21) was more effective at finding the taller protrusions that tended to penetrate a mesh filter that did not settle as close to the ground as it did with closer particle spacing.

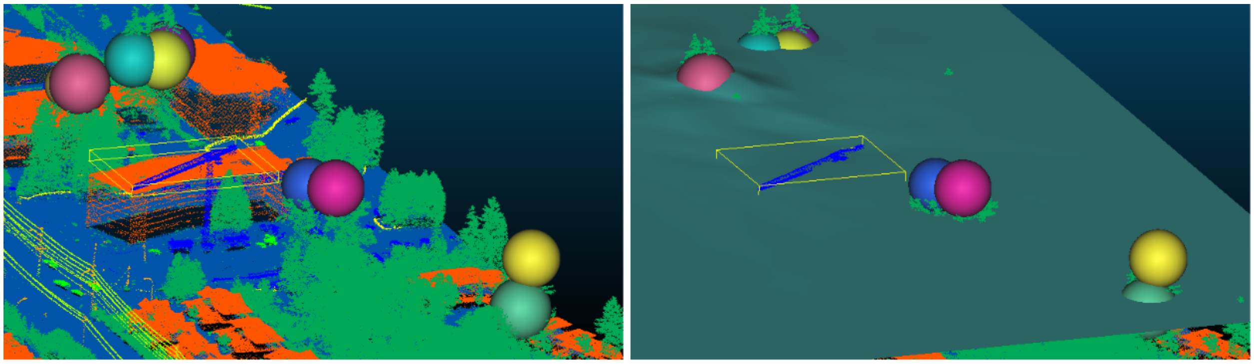

(a) (b)

Figure 2.21: Blue crane points in (a) fall within a wide and deep yellow bounding box. Nine colored spheres adequately find treetops above the mesh in (b).

Transmission tower representation was also stymied by the proportional height filter's requirement for tall bounding boxes. The most prominent tower among the analyzed tiles ($\psi = 346m^1.5$ in Table 2.4) is shown in Figure 2.22. The dense power line returns result in a connected component bounding box that is too wide, even for this tall tower.

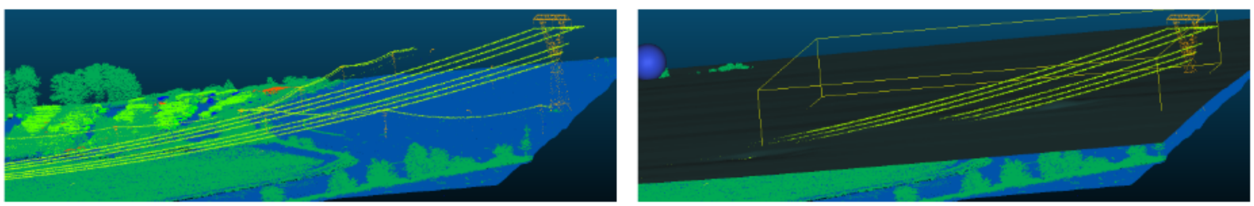

(a) (b)

Figure 2.22: A 38 m electrical transmission tower is shown before filtering in (a). After applying the mesh, (b), 23 m protrudes. The yellow bounding box also encompasses dense power lines, causing the bounding box width to be larger than the height.

Scenes with the most prominent light poles benefited significantly from the additional clustering. Light poles do not support power lines that inflated the bounding box width on other support structures. The prevalence and density of light poles also offered much room for improvement over taller cranes and communication transmission towers. The preceding mesh filter often required a close mesh particle spacing in order to settle on these shorter protrusions. This settling tended to include more vegetation and other points that are less important. Figure 2.23 shows an example where the clustering step greatly increased the prevalence of many subtle returns.

By zooming in on the point clusters in Figure 2.23b, the clustering effectiveness is more apparent. Figure 2.24a reveals that three light poles (including one wrongly classified as a vehicle) are highlighted along a road. The performance in Figure 2.24b is more nuanced in that it wrongly clusters some tall trees among 9 m light poles in a parking lot. Some protruding orange light pole

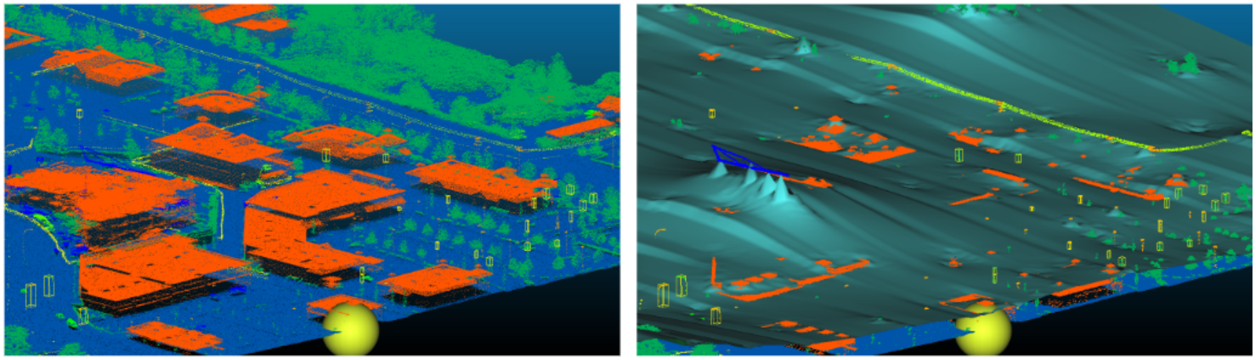

(a) (b)

Figure 2.23: A section of a suburban scene without (a) and with (b) 2 m mesh filter. The light poles along the road on the lower left are bounded along with light poles in a parking lot on the lower right.

points are not clustered due to their proximity to vegetation that causes their connected component bounding box to be too wide. Despite these difficulties, the proportional height clustering step increases average density and prevalence as shown in Table 2.9.

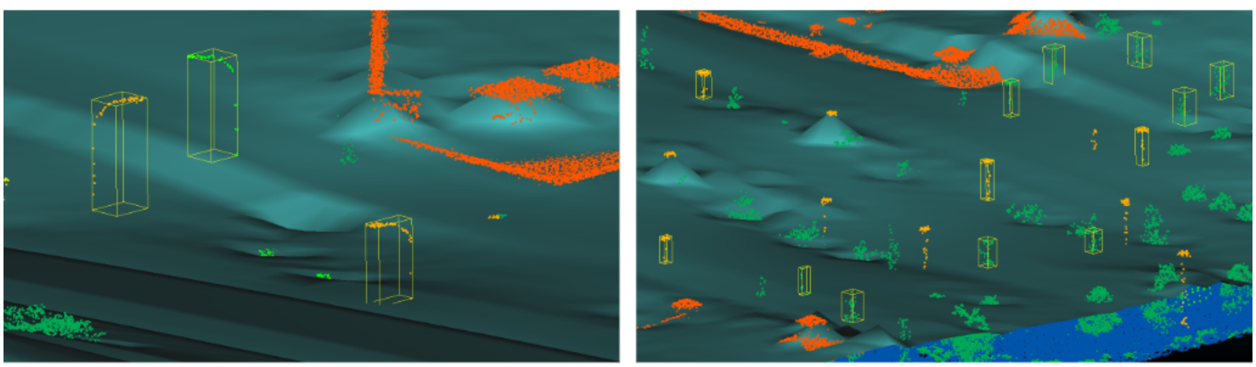

(a) (b)

Figure 2.24: Magnified portions from the lower left (a) and lower right (b) of Figure 2.24.

Table 2.9: Light pole clustering results. Quantitative results from the scene in Figure 2.23 are shown in orange.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Post-mesh density, %	Post-mesh prevalence, P	Number of clusters (vegetation)	Post-clustering density, % (change, %)	Post-clustering prevalence, P (change, %)
30	-	90	-	-	0.28	683	57 (49)	1.87 (566)	634 (-7)
6	-	19	-	-	0.78	846	13 (10)	3.24 (315)	343 (-59)
7	-	34	21	-	2.52	11,898	29 (16)	34.35 (1,263)	20,162 (69)
-	-	66	-	-	1.74	1,942	6 (4)	15.77 (806)	8,009 (312)
				Average	1.33	3,842		13.80 (738)	7,287 (79)

Utility poles were the last type of vertical obstacle considered. These shorter structures also tended to require close mesh spacing that resulted in the inclusion of distracting points. Similar to transmission towers, utility poles added the challenge of dense power line returns that collided with the mesh and resulted in low heights differences for the remaining pole points. Unlike the transmission towers, scenes with prominent utility poles had closer particle spacing, which decreased the minimum bounding box height requirement. Some short, but still narrow, bounding boxes passed through this filtering step.

Table 2.10: Utility pole clustering results.

Vertical structure prominence, ψ ($m^{1.5}$)									
UP	TT	LP	CR	CT	Post-mesh density, %	Post-mesh prevalence, P	Number of clusters (vegetation)	Post-clustering density, % (change, %)	Post-clustering prevalence, P (change, %)
10	-	-	-	-	0.97	1,529	6 (5)	0.34 (-65)	4 (-100)
27	-	-	-	-	0.33	482	8 (7)	1.56 (371)	89 (-81)
44	-	-	-	-	2.02	8,267	6 (2)	36.47 (1,706)	9,884 (20)
10	-	-	-	-	0.97	1,529	6 (5)	0.34 (-65)	4 (-100)
29	-	8	-	-	1.59	3,122	0 (0)	0 (-100)	0 (-100)
				Average	1.18	2,986		7.74 (370)	1,996 (-72)

The effectiveness of proportional height filtering with the DALES data set is mixed. It excels at increasing the prevalence of thin structures, such as communication towers and light poles, which do not support power lines. The proportional height filter also revealed several erroneously classified vertical structures. Beyond quantitative prevalence results, the proportional height filter showed potential to decrease the number of objects under consideration from thousands of raw LiDAR points to just a few bounding boxes.

2.4.4 Coping with sparse returns

Checking the proportional height of connected components only work if the vertical structures are represented by a single cluster of points. The filtering step should also cope with clutter and false returns that could originate from an object of interest. All the prominent vertical structures in this high density data set are represented by a nearly continuous line of points. The DALES data set also had LiDAR point outliers removed that could have been correlated to objects on the ground. Nevertheless, close examination of the data set does provide examples where the overlap algorithm could associate sparse points.

Three examples of distributed structural points is shown in Figure 2.25. The tops of three light poles return between 11 and 20 LiDAR points from their top light structure, probably due to their relatively wide horizontal surface. Careful examination reveals the light clusters are aligned and the height of each one is 9 m. Close inspection also shows that there are 2-3 dark blue points directly below the light pole tops. Were these structures to protrude above the mesh, the overlap algorithm could have readily identified these mis-classified points.

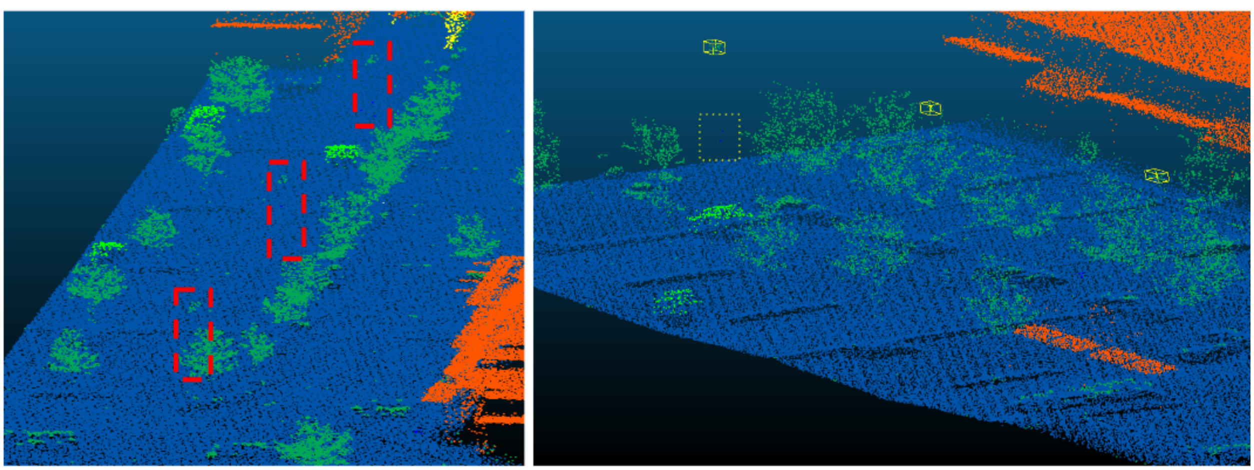

(a) (b)

Figure 2.25: A series of three 9 m tall light poles are outlines with red dashed rectangles in (a). The same three light poles from another perspective (b) shows the cluster of points at the top in yellow bounding boxes along with three stray dark blue points in the dotted yellow box.

The 3D connected component approach was not infallible in this data set. Figure 2.26 shows one example of erroneous clustering. Returns from the top of the light pole on the right side are batched with vegetation points that are 3 m away. The resulting wide connected component bounding box does not satisfy the proportional height filter. However, DBSCAN clustering correctly associates the red points from the light pole separately from the dark blue and green vegetation points.

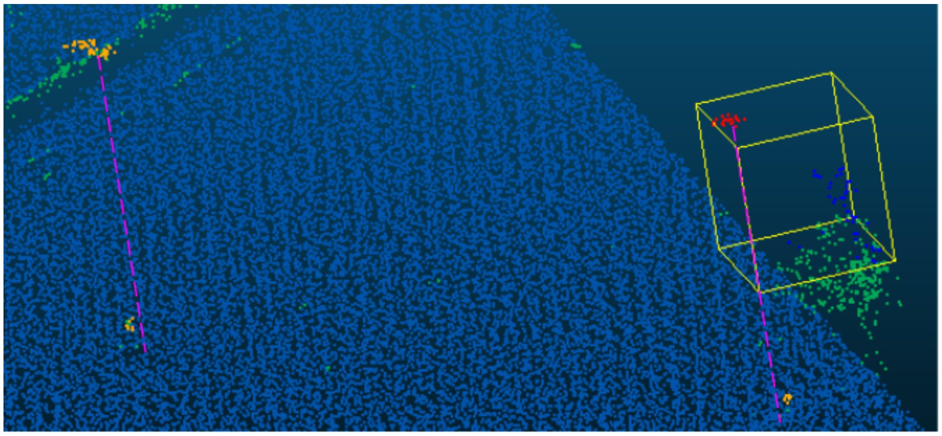

Figure 2.26: Two 9 m light poles inferred position is shown by the dashed magenta lines. Orange pole points are above the blue ground points. The yellow bounding box from 3D connected component analysis contains the top of a pole and vegetation points (turned red and dark blue, respectively.

These very low prominence structures contain very few points and, therefore, do not measurably improve vertical structure prevalence in their scenes. Most of the points associated with these objects were also incorrectly classified as points other than poles. Our simulated LiDAR point clouds provided sets of realistically sparse and cluttered point clouds. Next, we compared the effectiveness of three clustering approaches (DBSCAN, Connected Components, and the Overlap Algorithm) in finding vertical structure points from these point clouds.

The time required to complete initial DBSCAN clustering with larger point spacing necessary to bridge the gap between scan lines (4 meters per Eq. 2.3) takes over 8 times as long as 0.5 meter spacing used for initial clustering in the overlap algorithm. Increasing the eps distance also tends to incorrectly associate spurious points with actual objects. With two towers (Figure 2.27), increasing the eps distance steadily decreases the number of clusters, but still is unable to correctly associate all points from the large rear tower until eps distance is greater than the 4m raster line spacing in Eq. 2.3.

The 3D Connected Components approach struggled to distinguish tower components from false returns. The rear tower resulted in several returns from the same object that were almost directly behind one another in x. With initial eps spacing 0.5 m and beta range between 30 and 60 deg, the number of clusters was reduced from 90 to 38. However, clusters were not directly associated with false positives or actual tower objects (Figure 2.28).

Overlap algorithm performance versus two DBSCAN eps settings is shown in Figure 2.29. Initial DBSCAN clustering with an eps of 0.5 meters results in 90 clusters. Although a large

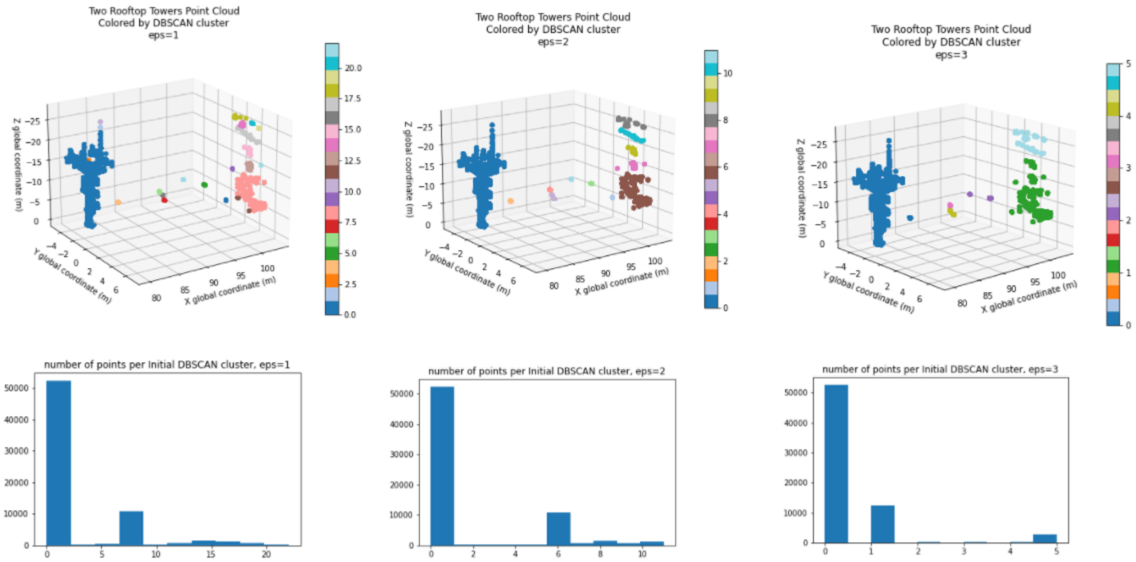

Figure 2.27: DBSCAN with increasing space between points (eps of 1, 2, 3 m, top row) and corresponding distribution of points per cluster (bottom row)

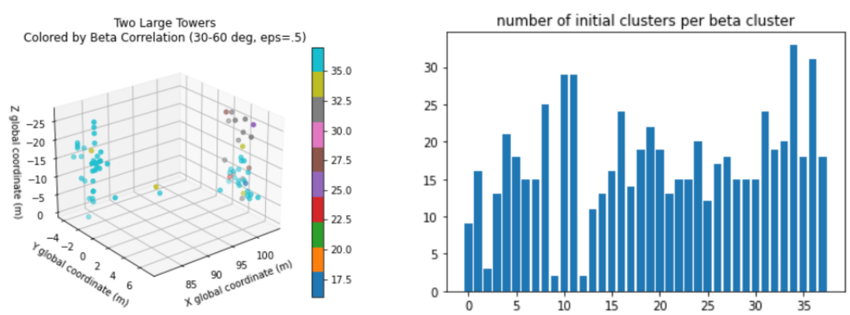

Figure 2.28: Connected components with $30° \leq \beta \leq 60°$.

portion of the near tower is clustered into the same object, the second tower is not clearly defined. The irregular shape of the top of the far tower is particularly challenging as even points in the same raster line are not correctly associated. The vertically oriented raster is farther apart than 0.5 meters and also not clustered. Increasing the eps distance to 4 meters per Eq. 2.3 to bridge the height difference between raster scans (thus associating points clusters in the same neighborhood) optimizes DBSCAN performance. All points within the near and far tower are correctly clumped. However, the larger distance setting also identifies false returns as tower structure. A string of yellow returns stretches close to 10 meters in front of the rear tower. Interestingly, DBSCAN with larger eps distance also fails to identify a cluster directly below the front tower. The overlap algorithm correctly identifies both towers and separates 6 of the 8 point clusters that originated from the small spheres. The two spheres not successfully broken out were within the overlap disc area. Structurally, the front false return was almost directly below the cell phone tower's wide top platform. The false

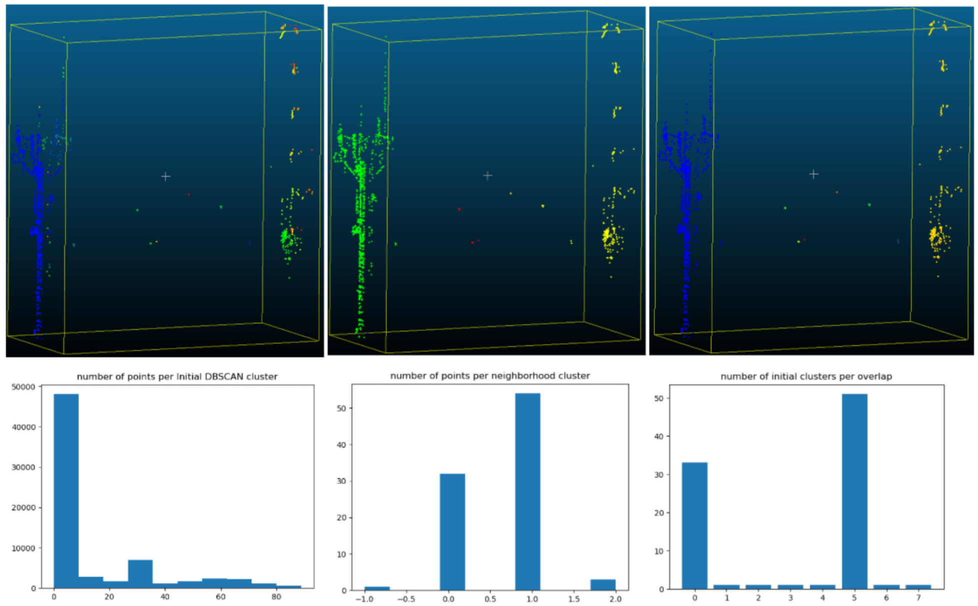

Figure 2.29: Initial Clustering (eps=0.5, left), Neighborhood Clustering (eps=4, center), and Overlap Algorithm (right)

return for the rear tower was close to a guy wire from the same tower.

Table 2.11 shows FMI scores for DBSCAN with Neighborhood point spacing versus Overlap Cluster performance for the four tower arrangements. The Overlap approach recognizes tower structures more effectively than DBSCAN in all four arrangements. The FMI penalizes an algorithm for incorrectly assigning detached points or not assigning attached points but does not consider correct identification of false returns. The larger number of Overlap clusters indicates successful identification of floating spheres not portrayed in the higher FMI score.

Table 2.11: DBSCAN neighborhood versus Overlap cluster performance for four tower arrangements.

Tower Arrangement	Initial Points	Neighborhood Clusters (FMI score)	Overlap Clusters (FMI score)
1	54992	7 (.79)	12 (.96)
2	45448	4 (.95)	12 (.99)
3	17568	7 (.81)	21 (.88)
4	68677	4 (.95)	8 (.99)

2.5 Discussion

The tallest protrusions (communication towers and transmission towers with less pronounced power lines) amplify their distinction with our coarse mesh. Points from horizontal surfaces on cranes and suspended power lines tend to collide prematurely with the descending mesh and make their associated vertical structures less obvious. Shorter structures, which composed most of the pole points in the DALES data set, can persist through subsequent filtering and automatic identification. Associating points based on their connectivity and proportional height decreases the number of objects under consideration by two more orders of magnitude. Relatively scant returns from smaller light poles require a separate overlap process to deduce thin structures.

This analysis operated within the limits of the DALES data set which has some key differences from what sensors on an aerial platform would encounter. The mesh descent step is particularly affected by the airborne LiDAR's nadir perspective. The top-down perspective greatly increases the density of points on horizontal surfaces such as power lines and crane booms. With a more flight representative sideways angle of incidence, most returns would come from surfaces below and ahead of an aerial vehicle. Poles are at least an order of magnitude wiser than their associated power lines. The total number of points would likely decrease while the number of pole points in a given scene would increase.

The consolidation of multiple nadir perspectives is another artificiality. An online sensor cannot expect this amalgamation of multiple perspectives. One example of this is simultaneous point returns from the near and far sides of a building in addition to the rooftop. This omnipresent perspective resulted in premature mesh collisions. The combined perspective also prevented transmission towers from passing the proportional height filter. Figure 2.30 shows relatively narrow side view where the 7 m portion that protrudes from the mesh is just over 1 m wide. The perspective on the bottom row shows a less pernicious 23 m width. Regardless of the perspective, the proportional height filter eliminates a 17 x 28 x 9 m bounding box. An online sensor data stream would provide a single sideways perspective. Figure 2.30 also shows how connected power lines increase the bounding box width.

Power lines associated with transmission towers and utility poles often diminished their protrusion height. Recognition of transmission towers is hampered by the inordinate quantity of power line points which prematurely collide with the mesh. In a more realistic point cloud, there would be far less power line points. This would result in more transmission tower points protruding above the mesh. An example of sparse power line returns is shown in Figure 2.17. Less than 2 m of 14 m tall poles protrude above the mess despite being 100 m from the treeline. A line of utility poles near orange buildings are much more prominent due to sparse power line returns.

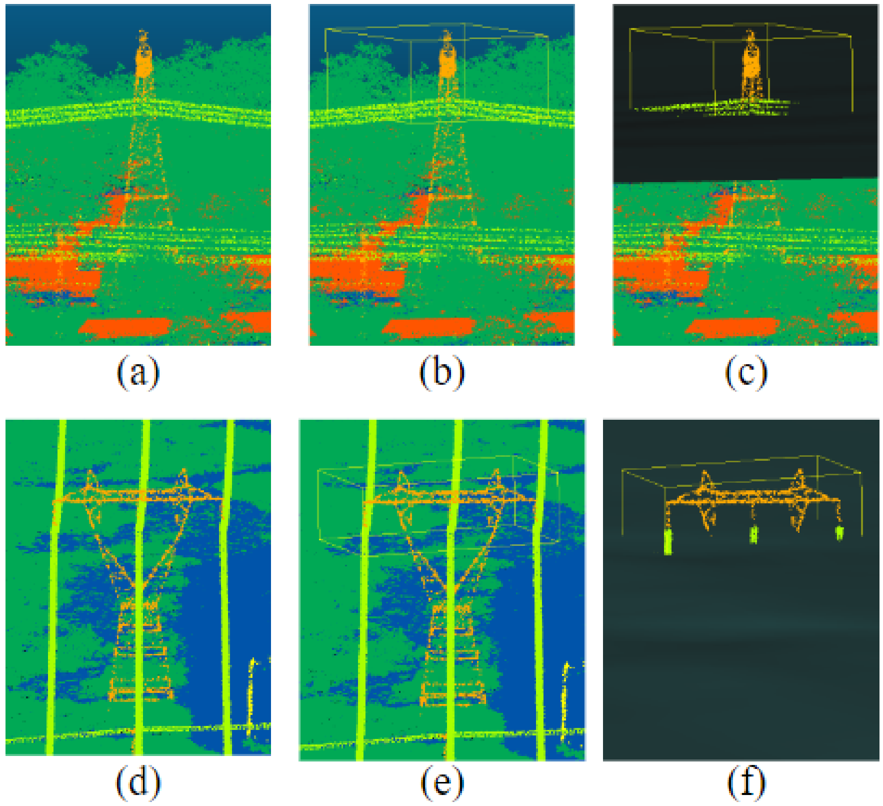

Figure 2.30: Two perspectives on a 35 m transmission tower. The top row shows a narrower side perspective (a), with proportional height bounding box (b), and with mesh (c). The bottom row shows the wider side perspective (d), with the same proportional bounding box (e) and with mesh overlay (f).

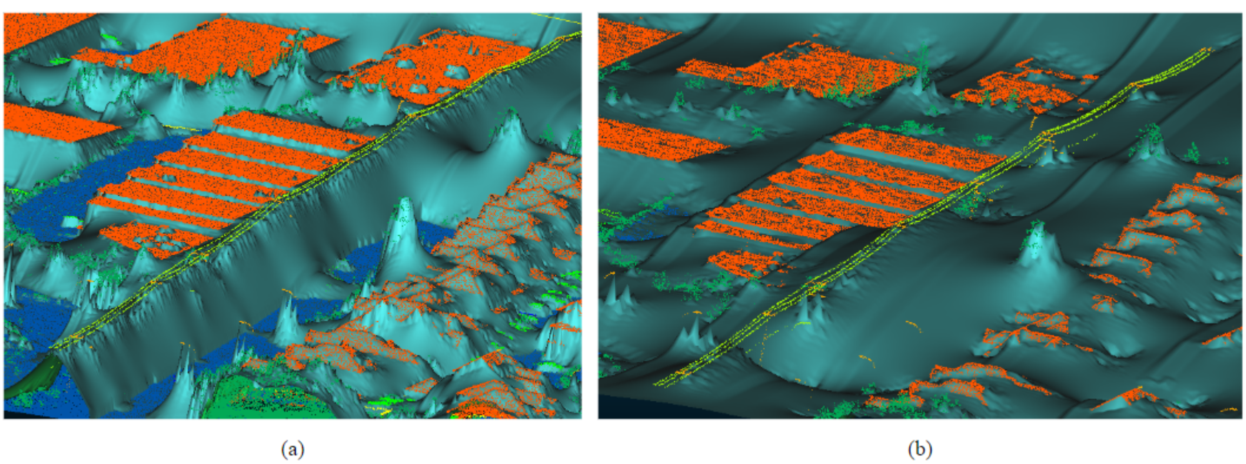

Figure 2.31: The same scene with 0.5 m (a) and 1 m (b) grid spacing.

The DALES data set was selected since it offered labelled pole points in a variety of scenes. However, Figure 2.32 shows examples of widespread incorrectly classified pole points. These inaccurate classifications tended to have an out sized impact on the proportionally small number of pole points in the overall data set. These inaccuracies show that the original human classification process that had the benefit of multiple perspectives and time are vulnerable to mistakes. This fallibility extends to previously mentioned machine learning approaches which rely on accurate training data. Geometric rule-based algorithms are not subject to these oversights.

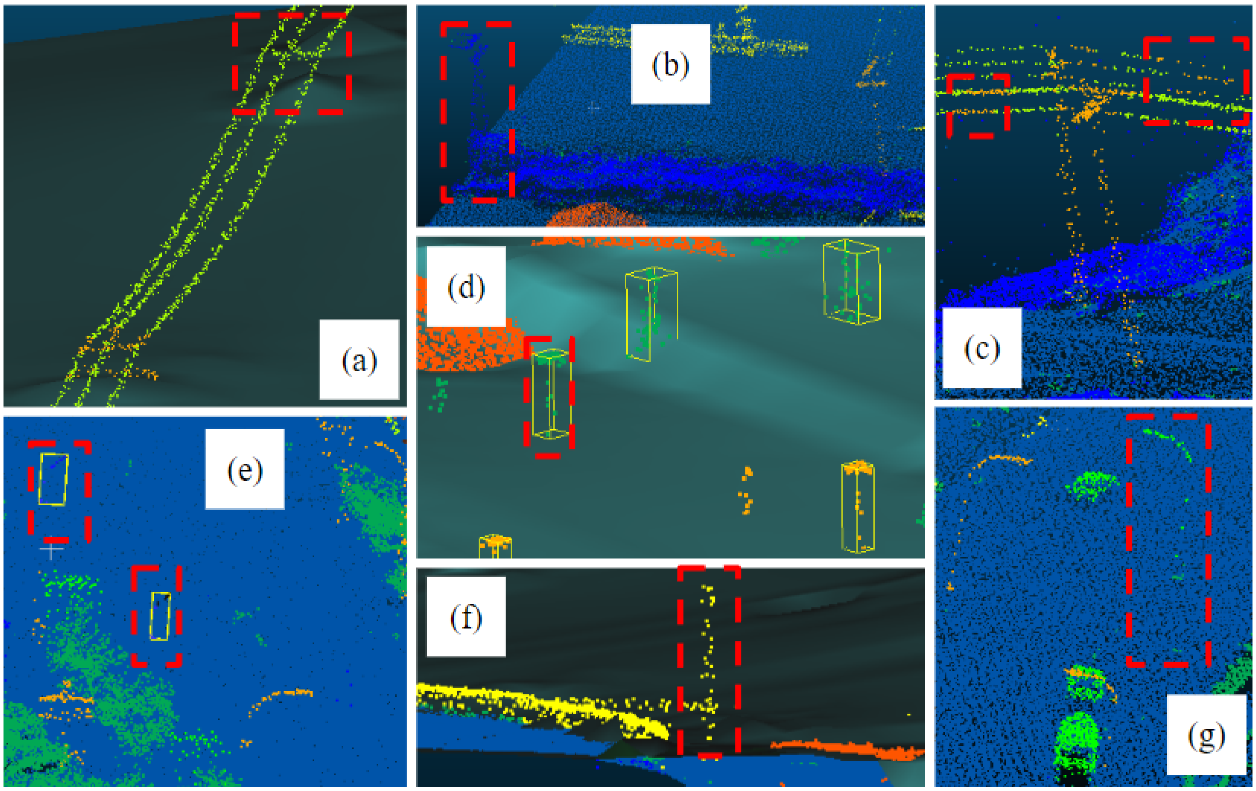

Figure 2.32: Examples of inaccurate labels. Correctly labelled pole points are orange, incorrectly labelled points are bounded by dashed red rectangles. Examples include: a utility pole classified as a power line (a) or unknown (b), power lines classified as poles (c), light pole classified as green vegetation (d), light poles labelled unknown (e), a 7 m pole labelled as fence (f), and a light pole labelled as a vehicle (g).

In general, the mesh filter allows key points to protrude. Figure 2.31 shows how slightly increasing the grid spacing allows the mesh to descend past most power line returns while also draping above blue ground surfaces. This ability to bypass suspended points portends that the coarse mesh would effectively descend past suspended outliers and other clutter. The experimental results showed how closer particle spacing was suited for scenes with less prominent vertical structures. However, selecting this close spacing also tended to increase the number of distracting points (from rooftops, the ground, and vegetation) that penetrated the mesh filter. Use cases where an airborne vehicle are concerned with more prominent ($\psi \geq 100$) vertical structures should start with larger particle spacing ($s \geq 8m$). In addition to mesh particle spacing, the minimum number of LiDAR points, $|l|_{min}$, should be proportionally adjusted according to the expected point density.

The time required for searching for embedded spheres depended on the number and arrangement of the remaining points. Finding spheres in the dense treetops generally took on the order of a second. The time required for an online use case will depend greatly on the sensor resolution and number of points above the mesh. Regardless, the searching process shows potential for online utilization. The proportional height and overlap filters similarly depend directly on the number of points remaining for consideration.

Other automated segmentation approaches use the harmonic mean of precision and recall to get an F_1 score. Our approach shows value in terms of precision by rejecting false positives. However, a high recall depends on finding more of the labelled points. The F_1 score does not prioritize points based on their significance. Our approach purposely rejects lower altitude points (including ones that belong to vertical structures) and prioritizes higher altitude points that protrude above the mesh. Quickly and reliably identifying these high obstacles is much more important to an air vehicle than classifying 100% of millions of points in a scene.

This study sought to identify returns from tall and thin vertical structure types that are typically associated with hard to see obstacles. The proportional height and overlap filters sorted points based on their geometric arrangement. In many cases, treetops satisfied the proportional height filter in addition to man made vertical structures. Although the treetops were not the original target, identifying thin and tall treetops (especially for leafless deciduous trees) could be a useful result.

An autonomous background process could conduct ongoing analysis of the intermediate products to provide valuable insight. For example, alerting an operator when vehicle flight altitude descends near or below the Coarse Mesh altitude could rapidly convey a cause for caution. This process could forecast future vehicle positions based on a flight plan or other planned states, providing a closed loop system to incorporate uncertainty. This uncertainty could then feedback into a coupled autopilot system to reduce the closure speed or steepness of a turn.

2.6 Summary

Current automated methods that use passive sensors require high contrast data which is difficult to achieve from an aerial platform. Existing algorithms that use active sensors struggle to efficiently and reliably segment vertical obstacles from large point clouds where the vertical structures are only sparsely represented. The large prominence of communication towers, ψ, conveys the efficiency of coarse mesh in finding these tall towers. However, the DALES data set only contained four communication tower examples among the shorter and more numerous pole types. Expanding the analysis to scenes with less prominent vertical structures was less fruitful due to a high density of returns on horizontal surfaces due to the nadir sensor perspective. Close inspection revealed that the

DALES data set had widespread labelling inaccuracies.

Our new Overlap correlation approach builds on the fast DBSCAN clustering algorithm and generalizes previous rule-based algorithms designed for ground applications. The Overlap algorithm outperforms tuned implementations of more basic clustering algorithms by successfully associating most simulated returns with the corresponding vertical obstacle while also correctly identifying simulated noise. The performance improvement ranges from 4 to 22%, indicating that the Overlap algorithm shows potential to identify vertical obstacles within a variety of cluttered and sparse point clouds. Our dual-filter approach shows potential to quickly distill both continuous and sparse obstacle data.

CHAPTER 3

Graphic Augmentation of Vertical Obstacles: Focus Groups and A Simulator Study

The previous chapter described a novel fast and modular process for identifying communication towers in real world scenes. It employs an algorithm that efficiently removes extraneous points and automatically distills remaining data to find significant vertical structures represented by sparse and cluttered point clouds. The approach is designed for online implementation and has minimal tuning requirements. Being able to identify the location of vertical structures quickly and reliably is necessary but not sufficient for avoiding collisions and catastrophic accidents when operating at low altitudes. The presence of those structures needs to be highlighted for pilots who are required to divide their limited attentional resources between numerous tasks, including scanning the external environment for obstacles but also navigating, monitoring flight instruments (altimeter, attitude indicator, etc.) and managing critical mission equipment (transponder, radios, etc.). Earlier research [88, 89, 90, 91] has shown that helicopter pilots spend between 57-67% of the time looking out the window during actual and simulated flights at low altitudes. Even when their visual attention is focused on the outside, [34] found that it took pilots between 2.4 and 6.1 seconds to detect unexpected obstacles in limited visibility. This four-second range can be the difference between a close call and a collision as supported by accident statistics indicating that failures to notice and avoid obstacles are only behind loss of situation awareness as the documented known cause of fatal civilian helicopter accidents [5].

In order to better understand and be able to address the challenges involved in obstacle avoidance, we first conducted 4 online focus groups with experienced pilots and engineers [4]. We then designed and evaluated, in a simulator study, the effectiveness of candidate graphic augmentation techniques to support the detection of vertical obstacles.

3.1 Focus Groups on Obstacle Detection

Reliable detection and avoidance of obstacles are fundamental to safe flight operations. The timely perception and understanding of the relationship between an obstacle position in relation to one's ownship position is critical. Detecting obstacles in time to take evasive action is an especially pressing task at low altitudes. The response time, the time-to-criticality (TTC), is dependent on the pilot's attentional focus. This response characteristic can range from less than 1/2 second for simple reaction time to visual stimuli,[92] increasing to between 1 and 2 seconds[93] for simple cognitive operations to a far greater period of at least 8 to 10,[93] 12.5,[94] or 13[34] seconds to take action when faced with an unexpected challenge. The simple awareness of the threat is necessary but not sufficient; additional factors, such as proximity, must also be quickly understood. Situation awareness is the ability to identify, process, and comprehend critical, perceived elements of information in one's environment in order to make decisions about future states of actions, and this situational understanding or representation is constantly being updated.[95][96] The nature and likely success of mitigations to obstacle collisions will vary greatly depending on whether the pilot is expected to manually manage and integrate that information or whether a semi-autonomous supervisory control approach is taken. The application of autonomous functions would potentially offload the demands on the pilot, but the demands on the system to reliably perceive, process and present critical information may also create challenges especially when the automation fails or does not behave as expected.[97] The understanding of the plethora of discrete challenges for detecting airborne vertical obstacles, and the range of potential mitigations was further explored in this study.

3.1.1 Method

Twelve subject matter experts (five helicopter pilots and seven engineers) were recruited for this study. Volunteers participated in in four two-hour focus groups. Each focus group consisted of a diverse set of 2-4 participants and two study investigators who guided the group discussions. All five helicopter pilots had experience with image intensification sensor systems (night vision goggles) and most (80%) also had experience with forward looking infrared and millimeter wave radar designed to detect obstacles and terrain. Two participants were AH-64 standardization instructor pilots. The majority of the pilots were very experienced, with over 1500 flight hours. Forty percent of the pilots were civilians. The seven engineer participants included two Federal Aviation Administration (FAA) rotor craft research engineers, a UAS certification manager, a future airspace technologist, an aviation vision systems engineer, an aviation geospatial practitioner and a survey drone software engineer. Although the engineer participants were not helicopter pilots, their roles often required them to consider flight operations in the current and developing low altitude flight environment. Panelists convened in online conference rooms. Three of the four groups had a mix of engineer and pilot participants to diversify the group's perspective. The investigators guided the discussion to address similar aspects of the research questions between the four focus groups. The research questions were focused on the challenges faced in obstacle detection, followed by the pursuit of information regarding available mitigations, and the possible limitations or shortcomings of those mitigations. Specifically, the research questions were:

1. How challenging is it to detect vertical obstacles?
2. What mitigates the risk of impacting an obstacle?
3. Are there any shortcomings to these mitigations?

The first question sought to discover what flight missions and environmental conditions were most closely associated with vertical obstacle detection and also how frequently each user operated in these conditions. This question also tried to elicit when obstacle avoidance arose as a concern and how a pilot establishes the location of an obstacle in flight. The second research question discussed ways to mitigate obstacle collisions with experience, preparation, and technology. Experience included training and any personal avoidance techniques. Preparation included best practices for planning and reconnaissance. Technology included sensors and electronic interfaces to maps and obstacle databases. The final research question asked whether there was room for improvement for these mitigations.

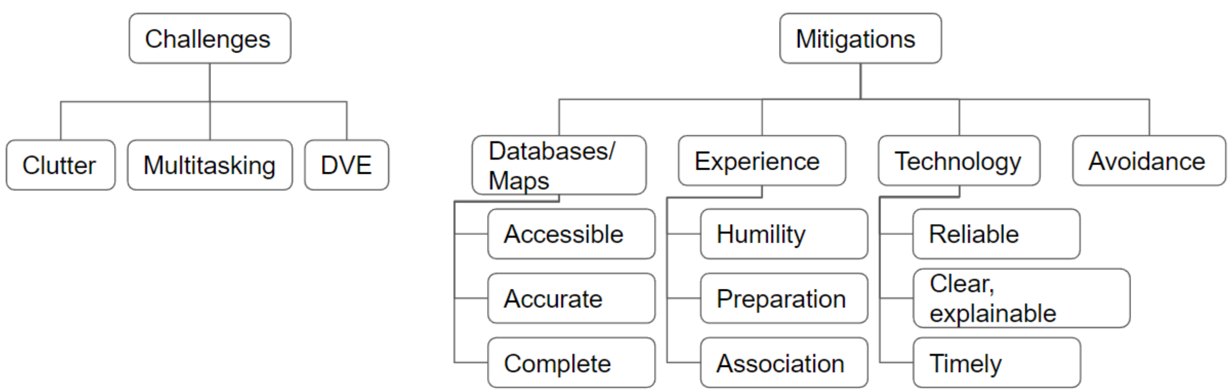

Figure 3.1: Initial categorization of comments arranged according to the first two research questions.

Discussions were recorded and automatically transcribed. The raw transcriptions were uploaded into NVivo, a qualitative data analysis software tool[98]. The 283 substantive comments were manually sorted into categories using the constant comparative analytic framework[99]. Separately, another investigator did a key word search of the transcripts based on his discussion notes to look for hidden categories or themes. Resulting categories corresponding to each research question and other emerging subjects are shown in Figure 3.1. Next, themes among the codes were deduced and correlated with the research questions in Figure 3.2 and Figure 3.3. Lastly, both investigators independently reviewed the raw transcripts again to verify that the captured themes adequately reflect the context of the surrounding discussion. The human research protocol was approved by the

University of Michigan's Institutional Review Board (study number HUM00202729). The study was subsequently reviewed by the United States Army's Human Research Protection Program office (administrative review number 21-025) due to the participation of a US Government employee as co-investigator and active-duty military panelists.

3.1.2 Results

The transcripts revealed several recurring themes across the participants and within participant groups (military pilots, civilian pilots, and engineers). These themes are delineated and described below. Unless cited to a reference, all quotations in this section are from focus group transcripts. Note that, according to many of the participants, when an obstacle appears unexpectedly "the mission stops", and the aircrew shifts their focus exclusively to identifying and avoiding the obstacle. This theme is critical to understanding how task interruptions affect flight safety.

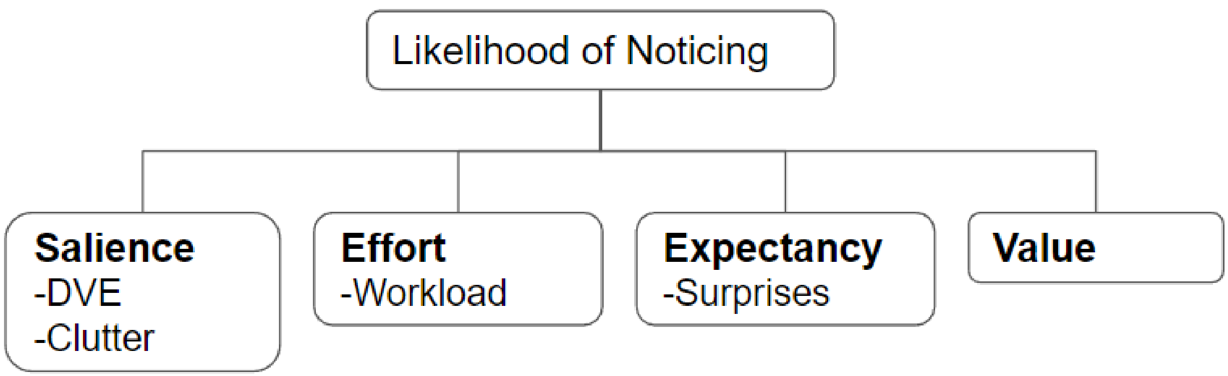

Figure 3.2: SEEV factors that decreased the likelihood of noticing vertical obstacles.

3.1.2.1 Salience

Salience is defined as "a signal-to-noise measure of the feature contrast between the target and the surrounding stimuli" [100]. It is one contributor to the fact that aircrews have "failed to detect a tower they were looking for and ran into it", as pointed out by pilots on the panel. Even under clear visual meteorological conditions detecting or distinguishing vertical obstructions in congested areas is challenging. Clutter exists when there is "close spacing between a target and surrounding distractors" and is exacerbated when these background features are similar the target [101]. Other objects and background with a similar textures, orientations, and colors are examples of clutter that decrease this signal-to-noise ratio and make it more difficult to detect an obstacle. Cluttered visual scenes are further compounded by any reduction in resolution due to reduced lighting or obscurants.

Thus, the salience of specific features is important to be able to confirm expected obstacles, as well as to identify new, novel, or other unanticipated impediments.

Clutter Panelists observed that the low cross sectional area of vertical obstacles make them difficult to discern when, as is often the case, they have a cluttered and non-uniform background. Detection despite clutter was the most frequently discussed detection challenge: "the biggest challenge [for detecting vertical obstacles] would be ground clutter". Clutter increases at lower altitude as more of the aircraft's field of regard becomes occupied with ground objects. When a tower was identified, there was suspicion that there are more towers not visible but close by as the same elevated terrain is commonly home to multiple separate towers. Nighttime operations in urban areas were "the most challenging environment" as the widespread illumination hinders detection of lit towers. At night, most pilots said that even lit towers among an illuminated urban background were especially difficult to discern: "in the city all the lights will drown out any lights that might be on the tower". The vertical orientation of towers also challenges detection as they typically run parallel to roads, buildings and other urban features that align in the same visual perspective. Outside of an urban environment, trees and other naturally linear features also can disguise vertical obstacles. Even aircraft structural design can hinder obstacle detection as no helicopter offers 4 pi steradian visibility where a pilot's field of regard is unlimited: "I can't tell you how many times aircraft were coming at us that I can't see [from the backseat of an Apache helicopter]". Pilots flying around transmission lines don't hit the big wire they are staring at. One panelist recalled an accident where "they hit a support cable for a tower that's out the right side of the aircraft and they're looking [out the] left [side]." Neutral tower coloring combines with the low cross section to decrease salience. Unless a tower is close to an airport or other air traffic area, they are typically grey or some other neutral coloring that blends in with the ground or sky. Away from urban clutter, pilots were concerned with the occasional cell phone tower. The varying size of a tower, a lack of size constancy, can also make accurate distance estimation difficult. There are few visual cues to distinguish a close short tower from a taller tower in the distance. Tower contrast also varies based on the location of the sun and the structure material.

Degraded Visual Environments Degraded Visual Environments reduce visibility to such a degree that obstacle awareness cannot be maintained as comprehensively as in Good Visual Environments[102]. Examples of DVE include night, fog, rain and dust. While most difficulties with obstacle detection occur even in day visual meteorological conditions, the impact of degraded conditions created additional difficulties. The topic of DVE arose in most of the groups even though it was not a directed line of questioning; however, it is a natural extension of challenges in the low altitude environment. Night flight is the most frequent DVE condition, and detecting obstacles at night with little or no natural illumination was cited as being particularly difficult. Furthermore, a panelist explained that unforecast reduced visibility conditions could make route planning that relied on visual detection of obstacles no longer tenable. Lastly, reduced visibility due to fog and clouds also arose as a challenge to vertical obstacle detection since the range of visual detection decreases even more. Multiple pilots agreed that even Visual Flight Rule prevailing visibility of three miles has haze that hinders visual detection.

3.1.2.2 Effort

Pilots flying at these low altitudes have many tasks competing for their attention (including maintaining aircraft control, navigating, monitoring radio traffic, and managing other mission systems) and are typically not solely focusing outside the aircraft. One pilot stated that "if you're focused on looking for obstacles you're not really focused on flying". Another panelist stated that "if your sole focus cannot be on scanning for obstacles (which it never is), it doesn't matter if you're an Apache or any other helicopter or aircraft flying low; it is very challenging to detect and avoid obstacles". Most civilian operations have only a single pilot, which further increases the workload concentration. Urban areas and operations in congested environments also require more precise navigation, flying, and coordination over the radios which consumes more attention that is needed to overcome the previously described visual clutter. Accurately identifying and correlating a tower in a crowded area also takes extra effort as similar unknown towers can cause confusion. One pilot recalled a leading course designed to reduce the risk of flying in the wire and obstruction environment. The class purports that detecting these obstacles should heavily rely on mental models of the infrastructure since wires and vertical obstacles are so difficult to see [103]. This course of instruction explains that it is more effective to predict the presence of obstacles and wires by becoming familiar with the design and arrangement of powerlines.

Other studies have also found that noticing obvious changes is more difficult when there are competing tasks requiring time sharing. The ability to notice changes dropped to low levels when the task of noticing was combined with complex flight control[104]. In another study, almost half of participants engaged in a competing task did not notice a change when looking directly at it[105]. Awareness of changes decreases further with more tasks[106]. The maneuverability of aircraft at these low altitudes often requires the pilot to consider a wide horizontal and vertical field of regard. This necessitates significant head and body movements to properly clear a future airspace. However, one panelist stated for aircraft with slewed sensors, there is a similar effort to point and adjust them to an area of interest. If attentional focus distracts a pilot or if the pilot dwells too long in one area, then the likelihood of missing other otherwise salient obstacles is reduced.

Reaction time Panelists estimated that this multitasking drove required reaction time to obstacles up to between 16 and 20 seconds depending on pilot experience and comfort. This timeline is consistent with findings from related research that documented the process of searching the visual field for hazards, detecting/perceiving those hazards, understanding the relationship between the hazard and the intended flight path, creating a plan for modifying the flight path, and then implementing the necessary control actions to avoid the hazard[34]. This time-to-criticality must be understood within the context of the pilot's attentional focus, and available response timeline.

3.1.2.3 Expectancy

In SEEV studies the detection rate for rare events was low. Consistent with that finding, all of the pilots had first hand experience encountering unexpected vertical obstacles in flight. An obstacle is less apparent to pilots that were not looking for an obstacle and did not notice it until it was uncomfortably close. Almost all pilots recalled close calls with towers that were seen only as they went right beside or under their aircraft. Wariness about vertical obstacle detection increased with more experience.

3.1.2.4 Value

Our panelists never questioned the importance of the awareness of the environment. However, depending on the phase of flight, there were other more pressing tasks than scanning for obstacles that consumed their attention. This discrepancy is ripe for further exploration.

3.1.3 Mitigations

To address the challenges identified in the panel discussions, experience-, training-, and technology-based mitigations were discussed frequently. For the sake of this paper, experience and training are integrated, and the mitigations are organized into measures meant to increase expectancy and value. In the words of one panelist, the key challenge is injecting the pilot with "with sufficient timing, knowledge, and awareness for them to be successful". Figure 3.3 shows how the effort, expectancy and value mitigation measures rely on prior knowledge from databases and maps. Expectancy and value mitigations also rely on automated processes that can result in automation complacency.

3.1.3.1 Avoidance

Many panelists brought up the current strategy of avoiding vertical obstacles preemptively by adhering to controlled airspace or avoiding known obstacles. When an obstacle was known or suspected, most pilots stated that they increase lateral spacing (when possible) and reduce

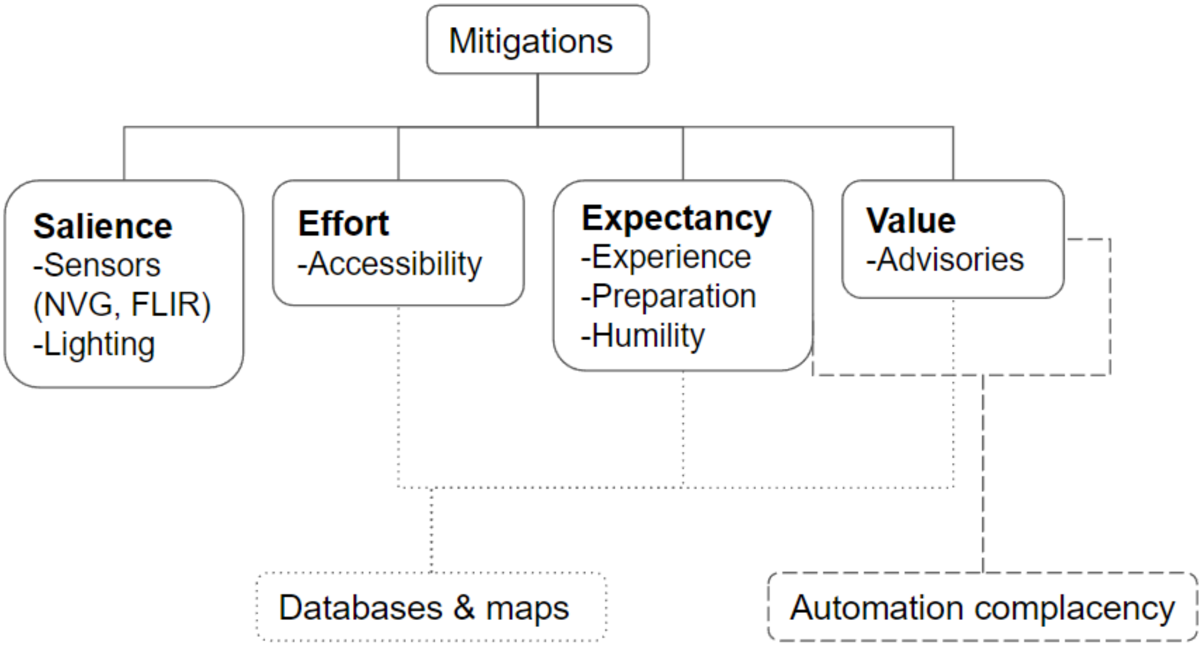

Figure 3.3: Possible mitigations for the difficulty of obstacle detection.

their airspeed to decrease the closure rate. However, regardless of the level of preparation, the 'surprise' effect of novel, unexpected, or unobserved obstacles was still a relatable experience. In addition, avoidance strategies may come in conflict with other factors such as unplanned rerouting into unfamiliar areas, or higher level direction (ATC in civil operations; tactical threat in military operations).

3.1.3.2 Salience

Beyond the out-the-window visual scene, infrared imaging systems (commonly called Forward Looking InfraRed, FLIR) was mentioned as one way to detect vertical obstacles. Although FLIR systems could be used to detect vertical obstacles despite low illumination and through some obscurants, the display is vulnerable to clutter due to a tendency for horizontal or vertical artifacts. FLIR imagery can provide a 2D camera ready scene that can be interpreted by the human pilot, but without a reliable means of size constancy in scene elements (i.e., unity magnification) and without an independent means of establishing range and bearing to objects in the scene, depth perception and optical flow will be adversely affected. Although Night Vision Goggles (NVGs - image intensification devices) were cited as the primary mitigation for night DVE, they are not suited for tower detection. Widespread LED lighting on towers themselves is not sensed by the NVGs and was cited as a shortcoming by both military and civilian aviators.

3.1.3.3 Effort

Database availability and quality arose as topics in one focus group. Even when a tower is depicted correctly on a map, "there's a high probability that I can find a few towers along the route that the student [pilot]never perceived" as the tower symbols "get completely lost in the clutter". Daily obstacle updates are published more frequently as Notices to Air Missions (NOTAMs). However, reading and correlating these text messages is "completely useless at communicating that [info to] the pilots" according to one senior pilot. Panelists also cited software applications that don't offer a "detailed obstacle planning" interface even though they suspected they referred to the same obstacle database. Thus, the effort required to verify and integrate information and plan ahead could be discriminating factors.

3.1.3.4 Expectancy

"In urban areas, it's easier to see [vertical obstacles] if you know they're there." This statement from one of the panelists highlights the importance of top-down information processing and the availability of good data before the flight. This challenge is confounded by the near certainty of changes to a planned flight route. Flight boundaries (airspeed, altitude, routes) can be established *a priori*, but that does not mean that additional measures or conditions will not also be required to maintain flight safety in a dynamic environment.

Experience and Training Training can allow a pilot to associate obstacles with other, more apparent, objects. Most helicopter pilots are familiar with the saying that "all roads have wires". One panelist offered that "all hills have antennas". Another shared that while crossing a bridge at low altitudes, he looks "outside a little bit to look for shadows." Unfortunately, even the best training is not fail safe. One pilot panelist had undergone an extensive training program designed to increase perception of wires and associated hardware and structures. Sadly, even this extensive training did not prevent two other similarly trained pilots from recently perishing after impacting tower wires. A fellow pilot who is more familiar with hazards along a flight route provides confidence, but their experience must be recent. A frequently cited mitigation to obstacle collisions was using Cockpit Resource Management where at least one crew member can "really focus outside of the cockpit" while the other performs other duties; however, this cannot be accomplished in single pilot operations.

Preparation Pilots agreed that the mere possibility of obstacles on an air mission increases the time and effort required for planning. However, given that many missions are urgent, investing additional time and effort is not always an option. One can increase knowledge about specific obstacles with route planning and detailed aircrew briefings. In cases where the anticipated workload is high, “you run out of radios and you can’t talk to everybody. That’s where you’ve got to rely on the planning for your route...and crew coordination of scanning and avoiding obstacles.” When there is sufficient time and resources, a researcher pointed out that “chair flying” or simulating a route can be helpful. However, simulated obstacles should be visually, temporally and spatially accurate so that they correspond to what will be encountered during the flight. The aircrew briefing provides a forum to understand “how you’re going to react to things before you take off.” Obstacle anticipation before takeoff can also include studying “satellite imagery, Google earth, or other nonstandard aviation tools” or IFR departure and arrival procedures (if instrument approaches exist) in order to develop control measures for locating or avoiding an obstacle in flight. High aerial reconnaissance flights above potential obstacle locations can also provide last minute preparation “based on the assumption that the pilot is able to perceive the hazard” and has time to conduct the reconnaissance.

Humility One pilot declared that an air crew’s “overconfidence in their ability to see and scan for obstacles is an absolute recipe for disaster.” There is always the possibility that a crane or other new obstacle has arisen since one last flew a route. Maintaining a sense of humility keeps the expectancy that an obstacle collision could happen.

3.1.3.5 Value

Technology can accentuate the hazard of obstacles and support their detection. Panelists cited displays that had pop-up visual or audio advisories that are designed to alert them of their proximity to an obstacle based on their aircraft’s database.

Advisories False positives have drastic effects on the effectiveness of any supplemental advisories. False alarms quickly decreased confidence in the advisory system. In one case, the potential for these alarms resulted in an FAA recommendation to deactivate the Terrain Awareness and Warning System[107]. In another case, a panelist recalled how erroneous alerts can block the entire screen and forced him to focus inside the cockpit to clear it. The bottom line is that when pilots do not trust the system, they simply will not use it.

3.1.4 Databases and Maps

Most mitigations designed to decrease effort while increasing expectancy and value rely on accurate and complete obstacle data, as shown in Figure 3.3. Thus, discussion about database accuracy was more nuanced. Panelists who were most familiar with surveying and mapping obstacles were most critical of their accuracy. For a lot of towers “it wasn’t uncommon for them to be inaccurate outside of their accuracy tolerances.” New survey data is often unverified or correlated with existing data. It is especially difficult to find the owner of private facilities to verify obstacle data. Interestingly, almost all panelists without firsthand surveying experience overestimated the accuracy of the obstacle databases. Assumed obstacle accuracy ranged from a "couple of feet" to "within a few hundred meters". Flights in proximity to large airports (such as Dallas-Fort Worth) seemed to be more accurate. Maps do not depict uncertainty associated with an obstacles position and this can cause confusion. Panelists explained that a user assumes the obstacle should be in the mapped location. If it is in a "slightly different location...that becomes a distraction from other hazards that are actually present." There were other instances where discrepancy misguided attentional focus. One pilot recalled that "the road will be on the wrong side of the airfield" when switching between data sources. Panelists recalled situations when this inaccuracy caused confusion when sensors declare an obstacle in a different location from the database, making one question whether the sensed tower is "the same thing...did you actually see it."; or whether there may be more than one obstacle in close proximity.

Completeness of the database was the subject of the most vociferous discussion. There was skepticism about how official databases are updated: "the most effective NOTAM is me telling my buddy 'watch out, there's a new tower out there'". Database completeness is challenged with limited sensor resolution. Most aerial and satellite imagery offer only "six inch pixel resolution". LIDAR data offers more dense sensor source data but is often not an accepted by the FAA. The focus on large airports results in "a ton of the obstacles are at 19 of the top 30 busiest airports." Helicopters usually operation from one of "9000 helicopter facilities which are private use" and therefore have limited access. This sheer number forces the FAA to focus on the basic task of updating facility coordinates and contact information, leaving low capacity for deploying survey teams. Databases also can consolidate multiple obstacles that are within a specific horizontal distance, only reflecting the altitude of the highest object[108]. Additional confusion comes from crowded areas like Las Vegas where "it was a disaster, trying to even figure out what they were referencing". Some panelists were aware that if an antenna is less than 200 ft tall it was not required to be reported as an obstacle[109]. One recalled over a decade of effort trying to get obstacles "charted below 200 feet." The long lead time for updating the database, typically over a month, was another source of skepticism. For example: "rehearsals certainly won't work for new obstacles that have popped up since the last database update". NOTAMs can also be incomplete: one pilot encountered "a crane that wasn't in the NOTAMs, wasn't in anything". Completeness also depends on prompt database updates for aircraft systems which "is not always true". Other users have "reported obstacles and two months later it's still not on the hazard maps." Although towers don't grow, they can change as antennas and other structure are added and taken away.

3.1.5 Automation Complacency

Users of automation are complacent when they tend to miss more malfunctions under automated control than they would under manual control[110]. Figure 3.3 shows how mitigations designed to increase expectancy, such as computer-based flight planning, or those designed to increase value by automatically presenting advisories rely on automated control. Many panelists recalled using flight planning software to quickly determine if there were any obstacles along their route of flight. However, the planned route“is only clear if you’re staying within that narrow corridor” according to a panelist familiar with the advanced settings of the flight planning software. Obstacle clearance along the planned route also requires accurate and precise navigation. Judging navigation accuracy depends on the skill and experience of the pilot. Experienced panelists expressed skepticism that junior pilots likely did not take these uncertainties into account. The details of how route accuracy affects automated route clearance was likely hidden in levels of menus that a typical user does not access

3.1.6 Discussion

This research gained current perspectives on the difficulty of detecting vertical obstacles from four focus groups. The panelists offered examples of how they overcame the detection challenge along with commentary on the suitability of these mitigations. All focus groups indicated that divided attention of aircraft operators at low altitudes is insufficient for visual detection of inherently low salience vertical obstacles. This challenge looms despite multiple examples of mitigations designed to increase salience, reduce effort, increase expectancy, and increase value. These mitigations range from sensors with better visualization than the naked eye (such as NVGs) to more autonomous systems which automatically determine proximity to an obstacle. In this discussion, we organize the mitigation options into levels of automation to compare potential challenges and benefits for informing future pilots.

Human-automation interactions range from low levels where there is no automatic assistance up to high levels where the computer processes raw data on its own, only presenting highly filtered information to the human. Parasuraman et al.[111] proposed a framework for selecting appropriate levels of automation for each of four stages of information processing – information acquisition, information integration/analysis, action selection/decision making, and action implementation. Information acquisition and information integration/analysis encompass the detection task whereas action selection/decision making refer to obstacle avoidance. The appropriate level of automation support at each of the four stages depends on factors including mental workload, situation awareness, complacency, timeliness, and reliability.

Based on the results of the panel discussions and leveraging the first stage of human information

Table 3.1: Levels of automation for information acquisition.

Level	Perspective		Sensor Data Mode	Filtering/ Highlighting	Acquisition, memory	Example
0	First person	-	-	-	Working memory	Naked eye
1			Passive	Manual	Video recording	NVGS, FLIR
2				Smart switching		-
3			Passive & active	3D proximity to an object	Data recording	Integrated sensor image
4		World (a priori)	Passive	Proximity to DEM	Route tracking	Integrated moving map
5			Passive & active	3D proximity to known obstacle	3D model	-

processing,[111][97] we propose the framework in Table 3.1 for comparing the levels of automation support that were mentioned during our discussions. We delineate six levels of automation where each level has distinct perspectives, data modalities and corresponding amount of automated filtering capabilities. This is not to be confused with previously-defined levels of human-system interaction, such as Parasuraman et al.[111], Sheridan and Verplank[112], and Copeland[97]. Lower levels of information acquisition only have a first-person perspective that relies on the fixed sensor position on a platform, whether that is the viewpoint out the window or mounting location of the sensor(s). Information about the environment is only available within the range of the sensor (or eyeball). Noticing information for a naked eye with level 0 automation relies directly on the pilot's perceptual abilities and working memory. Level 1 acquisition automation introduces passive sensors that send information to displays for presentation without filtering or highlighting. The pilot is responsible for choosing between which sensor(s) display or their naked eye. FLIR is one example of displayed information that can be recorded for level 1 acquisition automation. The next level of acquisition automation uses rules to automatically switch sensing modality based on the environmental or flight conditions. Active sensors that emit energy (such as LIDAR) are incorporated in level 3. The three-dimensional attributes of this data acquire the proximity to an object. There is no *a priori* data at this level, so all returns are valid at the moment they are received from a first-person perspective and there is no identification or correlation of these three dimensional returns. The sensor stream can record these additional 3D details for future analysis. Level 4 of automation acquisition introduces *a priori* data (such as a Digital Elevation Model) from a world perspective beyond the sensor's capabilities. Data from beyond the first-person view could also include lower latency reports of obstacle location. In level 4, passive sensors track aircraft position (with systems such as Inertial Navigation Systems) to depict aircraft position on map images, for example. These "moving maps" can have georectified symbology in addition to graphics on the original map. In the highest level of acquisition automation, active sensor(s) add another layer of three-dimensional detail in the context of *a priori* data, offering the capability of correlating real time returns with known obstacle and terrain databases. Three-dimensional data from the world perspective also increases the detail at this highest automation level.

Table 3.2 shows a framework that compares levels of automation for information analysis automation. We present six levels of information analysis automation. Higher levels of automation correlate and integrate information sources while propagating data quality. Advanced information analysis automation levels also automatically extract features from the sensor data that increase situation understanding. Level 0 of information analysis requires manual correlation and feature extraction from an image, whether from the out the window view or from a raw sensor image. At level 1, the visual sensing systems (such as the FLIR display) benefited from image processing meant to enhance contrast or remove noise without interpreting the image content. In both of these first two levels of automation, the goodness of the analysis is proportional to the image quality. Level 2 considers range data from a single sensor (such as a radar altimeter or LIDAR) to determine proximity to a solid object. Advisories informed by federated (radar altimeters) and combined sensors (Terrain Awareness and Warning System, TAWS) are other examples of information analysis automation that arose in the focus groups. This level of automation starts to provide situational understanding of the range and bearing to an object. Level 3 starts to incorporate *a priori* data and provides the ability to portray this known information alongside more recently sensed data. However, newer sensed data typically overwrites the older *a priori* information. The pilot maintains the burden of correlating known and sensed obstacle information along with resolving or ignoring data discrepancies. One example would be automated terrain banding where a low-resolution Digital Elevation Model does not depict a new tower visible in NVGs. Level 4 uses data from a single sensor to pursue automatic feature extraction. One notable shortcoming of this approach is that true validation requires previous ground truth information. Knowledge of the obstacle environment is limited to what has been within range of the sensor. Newly perceived obstacle information from series of images provides various options for automatic classification. The last level of information analysis automation fuses multiple data sources and dynamically compares sensed objects to select the best data to provide the most thorough feature perception.

Parasuraman et al.[111] offered that the number of levels of automation will differ between stages. For example, Sheridan and Verplank[112] proposed ten levels of automation between humans and computers for the decision and action stages while controlling underwater robots. Although our proposed frameworks both have six levels, the level of automation does not necessarily remain constant for the same system. For example, an airborne system with a passive FLIR sensor that acquires information with level 1 automation could use automation level 0 of the information analysis stage to present a raw two-dimensional image or it could use level 4 of automated information analysis to deduce structure from motion.

Table 3.2: Levels of automation for information integration/analysis.

<table>
<tr><th>Level</th><th>Data correlation</th><th>Data integration</th><th>Data quality propagation</th><th>Feature perception/ extraction</th><th>Situation understanding</th><th>Example</th></tr>
<tr><td>0</td><td rowspan="4">-</td><td rowspan="3">-</td><td rowspan="2">Image quality</td><td rowspan="4">-</td><td rowspan="2">-</td><td>Raw 2D image</td></tr>
<tr><td>1</td><td>Noise removal, edge detection</td></tr>
<tr><td>2</td><td rowspan="2">GO / NOGO relies on sensor's built in test</td><td>Discrete alerts</td><td>Radar altimeter</td></tr>
<tr><td>3</td><td>Reads from or overwrites existing database</td><td>Rule-based multidimensional awareness</td><td>Terrain banding</td></tr>
<tr><td>4</td><td>Only with self</td><td>Creates new data</td><td>Can estimate convergence/ divergence</td><td>Various automatic classification options</td><td>Limited by what sensor has seen</td><td>Structure from motion</td></tr>
<tr><td>5</td><td>Database lookup</td><td>Dynamic comparison</td><td>Dynamically quantify and select "best" data</td><td>Reinforced with correlated data</td><td>Dynamic world model</td><td>-</td></tr>
</table>

3.1.7 Summary

Longstanding challenges to detecting vertical obstacles are likely to become more pressing with emerging aviation concepts such as Advanced Air Mobility that will operate in increasingly congested airspace. Visually detecting these obstacles in cluttered displays and environments is a leading challenge for current operators. Pilots at these low altitudes also have especially high workloads with many competing tasks that increase the lead time required to effectively avoid obstacles. Current mitigations designed to decrease obstacle detection effort, while increasing salience, expectancy and value have significant shortcomings. Unexpected vertical obstacles continue to be a threat despite detailed planning, training, and other techniques that rely on prior knowledge. The accuracy and completeness of current databases is overestimated. Panelists observed that automation designed to reduce workload required detailed understanding that is not widespread. Improving these mitigations and enabling future advanced autonomy rely heavily on accurate, complete and up-to-date obstacle databases while carefully considering the impact on busy operators.

The findings from the focus groups underscore the need for some form of support to make pilots aware of vertical obstacles when they operate at low altitudes, in cluttered and/or degraded environments and are engaged in multiple competing visual tasks. This support is needed both when pilots are engaged in actively scanning their surroundings for hazards – visual search - and when their attention is focused on other tasks and needs to be reoriented. Two attention filter models [113] that provide insight on visual search and noticing while multitasking will be reviewed briefly as they can inform the design of such visual guidance: Guided Search [114] and Noticing Time-Salience Effort Expectancy Value (NT-SEEV) [115].

3.2 Guided Search and NT-SEEV

The Guided Search model describes variables that affect the time required to find an object of interest (such as a tower) when looking for it in a crowded environment (as faced by pilots during low altitude operations). According to the model, search time depends on other objects in the field of view, also known as distractors, and the degree of clutter in the scene. Clutter [101] refers to irrelevant scene information that decreases search efficiency. Less time is required to find a target when it shares fewer features (size, shape, or color) with distractors or when the target background offers a higher contrast [13]. The model posits that we first process basic preattentive and salient features such as color, shape, and motion simultaneously across a large visual field [116, 117]. The information gathered at this early stage then guides our visual attention to a particular object or location for more in-depth processing. Prior knowledge of specific target features, a top-down influence on attention, decreases search time. The latest version of Guided Search [118] discusses three factors that affect visual search: Priming, Scene Guidance, and Value. Priming affects which parts of a scene are examined based on previous attentional deployments. This concept is related to Inhibition of Return [118] which refers to people's tendency to avoid returning too soon to locations that have recently been sampled. Scene Guidance leverages information in the scene other than the actual target. Scene Guidance suggests that targets with logical syntactic placement (e.g. keyboards don't float) and semantic relationships (e.g. keyboards are typically close to a computer) are found faster. Finally, the Value of a target affects visual search efficiency (e.g. prioritizing red targets). The Guided Search model does not apply to dynamic scenes where targets move and fade in and out of view.

This aspect is included in the NT-SEEV (Noticing Time -Salience, Effort, Expectancy, and Value) model which predicts the allocation of attention and the likelihood and time to notice a discrete event (such as the appearance of a vertical object in a scene) in the context of routine task-driven scanning across large scale visual environments [14]. In an aviation context, the model successfully predicted when an infrequent event was noticed by pilots engaged in divided attention [119]. Similar to Guided Search, this model includes two top-down factors: expectancy and value. Expectancy increases the likelihood of noticing an event or object by relying on prior knowledge (e.g., knowing that there will be a vertical object at a given location). Value refers to the relevance of attended data or events to a given task. Bottom-up factors in the NT-SEEV model are salience (how much an object of event stands out against the background) and effort (how much of the visual field a person has to scan to find the object or event of interest).

Both Guided Search and NT-SEEV highlight the role of salience and expectancy in visual search and noticing. We therefore choose to adjust these factors in the design of sensor visualizations and graphic augmentation candidates to be tested in a subsequent simulator study. Earlier research has shown that graphic augmentation of known objects is a promising means of supporting the timely detection of unknown and/or unexpected hazards [35]. For example, pilots initiated avoidance maneuvers sooner for obstacles placed along a simulated helicopter flight route at 1,000 or 200 ft Above Ground Level (AGL) when these objects were flashing or brightened on a moving map display [36]. In another study, helicopter pilots operating at low altitude (100 ft AGL) noticed obstacles farther away when they were augmented with graphics on a heads up display [34]. Separate ground-based studies [37, 38] showed that, in the absence of visual cueing, objects were noticed later or missed completely while precise cueing decreased detection time in another visual search study [15].

In the low altitude flight environment, the default method for detecting obstacles is looking out the window. The salience of potential obstacles in this expansive and dynamic scene can be degraded by low ambient light or obscurants (such as visible moisture or smoke). Technologies designed to overcome these degraded visual environments, thus increasing overall scene salience, include Night Vision Goggles (NVG) and Infrared (IR) thermal sensors. In addition to changing the overall scene salience, graphic augmentations can increase expectancy by directing visual attention to an area of interest.

The next section reports on a simulator study that compared the effectiveness of sensor images and graphic augmentations in facilitating hazard detection with divided attention in the low altitude flight environment.

3.3 Simulator Study on the Effectiveness of Sensor Visualizations and Graphic Augmentations for Supporting the Detection of Vertical Objects

This simulator study examines the effectiveness of two interventions – sensor visualizations and graphic augmentations – for supporting the detection of vertical structures (towers) that represent a potential hazard during low altitude operations. Specifically, this experiment aims to:

1. Compare the noticeability of expected and unexpected vertical structures between simulated unaided (naked eye) and aided (image intensification or thermal imaging) visualizations

2. Compare the effectiveness of two types of augmented conformal graphics (Obstacle Visual Augmentation, OVA) for improving the detection of vertical obstacles

Previous static visual search studies [13, 116, 120] indicated a faster detection time with color and other additional acuity (motion, orientation, size, and shape). Our experiment varied scene salience with two sensor visualizations for each of the three Ambient Visual Conditions (day, dusk and night). Based on the previous visual search studies, we expected the following outcomes:

1a. In day Ambient Visual Conditions (AVC), obstacles will be noticed sooner (and/or more reliably) with unaided video than IR video

1b. In dusk AVC, obstacles will be noticed sooner (and/or more reliably) with IR video than with unaided video

1c. In night AVC, obstacles will be noticed sooner (and/or more reliably) with NVG than IR video

Our study has three levels of graphical augmentation: none, *a priori* circles, and more precise bounding boxes that would rely on an onboard active sensor. Previous studies in dynamic airborne [34, 36] and static environments [35] indicated faster target detection time with graphical augmentation. Separate ground-based studies [37, 38] showed objects without cueing were noticed later or missed completely. More precise cueing decreased detection time in another visual search study [15]. Therefore, we expect faster and more reliable detection with more specific cueing.

2. In all AVC, obstacles augmented with bounding boxes will be noticed sooner (and/or more reliably) than those with a priori circles, followed by those without cueing

3. Obstacles augmented with a circle or box will be noticed sooner (and/or more reliably) with IR video than with the corresponding unaided/NVG video

The study (protocol number 2022-026) was approved by the United States Army Medical Research and Development Command's (MRDC) Institutional Review Board (IRB), Log Number M-11028. The University of Michigan's Health Sciences and Behavioral Sciences IRB ceded oversight (HUM00219941) to MRDC's IRB according to an Educational Partnership Agreement between the University of Michigan and the United States Army Aeromedical Research Laboratory (USAARL).

3.3.1 Methods

3.3.1.1 Participants

Twenty-five U.S. Army helicopter pilots with current medical flight clearances were recruited from the U.S. Army aviation community at Fort Novosel, AL. All participants had normal or corrected-to-normal visual acuity which allowed the use of eye tracking equipment. Unfortunately, eye tracking data was lost for three participants. Of the remaining twenty-two pilots, twenty-one were qualified in the UH-60 and one was an AH-64 pilot. Years of service as a military aviator ranged from 1 to 36 years ($M = 11.9, SD = 10.3$). Total flight hours ranged from 200 to 9,500 hours ($M = 951, SD = 2,551$). For pilots qualified in the UH-60, UH-60 flight hours ranged from 60 to 8,000 ($M = 1,459, SD = 2,175$). Experience with Aviation Night Vision Goggles (AN/AVS-6) ranged from 25 to 2,000 flight hours ($M = 429, SD = 541$).

3.3.1.2 Apparatus

All flights were conducted in the USAARL NUH-60FS Blackhawk simulator. The NUH-60FS is fully accredited by the Directorate of Simulations (DoS) and by the Program Executive Office Simulations, Training, and Instrumentation (PEO-STRI), as a 6-Degree of Freedom (DOF), full-motion, and full-visual (Level D equivalent) NUH-60FS Black Hawk helicopter flight simulator with interactive UH-60M control heads and Multi-Function Displays. To simplify data gathering, the simulator was stationary during our experiment. The flight simulator is equipped with an Rsi CV10R dome and eight Barco FS40 projectors. CATI Training System's X-IG 5.0 Image Generator System simulates natural helicopter environment surroundings for: day, dusk, night, dust, snow, rain, clouds, and mid-range (1.4 to 3 micron wavelength) Infrared (IR) characteristics. The right seat perspective is shown in Figure 3.4. The participant's right side Multi-Function Disply (MFD) showed flight instruments while the left MFD showed a moving map overlaid with the current flight route.

Figure 3.4: Right seat perspective with unaided dusk Ambient Visual Conditions.

Flight routes (Figure 3.5) were based in a 140 km^2 San Francisco terrain model created by PLW Modelworks LLC. Imagery with 10 cm/pixel resolution was overlaid on a rolling, predominantly urban terrain model supplemented with three dimensional objects that represented houses, skyscrapers and other structures. These objects were supplemented by 25 thinner vertical structure models reported in the Federal Aviation Administrations (FAA) Digital Obstacle File (DOF). In addition to these 25 vertical structures, four towers were placed in ecologically valid settings in the North, East, South, and West sections of the terrain model. The top 125 ft of each tower was the same object.

Tower height Above Ground Level (AGL) ranged from 130-220 ft depending on the altitude of the underlying surface.

Each flight route included a unique audio file that combined actual tower or approach control recordings from LiveATC.net with a separate, scripted conversation. The scripted conversation was between NORCAL approach and five aircraft entities. Three male and one female volunteer responded to a male NORCAL approach entity according to a randomized script. Aircraft call signs were United 0430, Army 2280, Pick 2874, Delta 3120, and Army 0474 (the simulator call sign). Information requests for Army 0474 were followed by a 10-20 second pause (except for the real world background recording) to allow time for the participant to respond. The scripted conversations were structured to avoid participant response requirements in the 30 seconds prior to a tower encounter while the other aircraft entities sustained conversation with NORCAL approach. Brownian noise and light distortion were added to the scripted recording to imitate degraded aircraft radio sound quality.

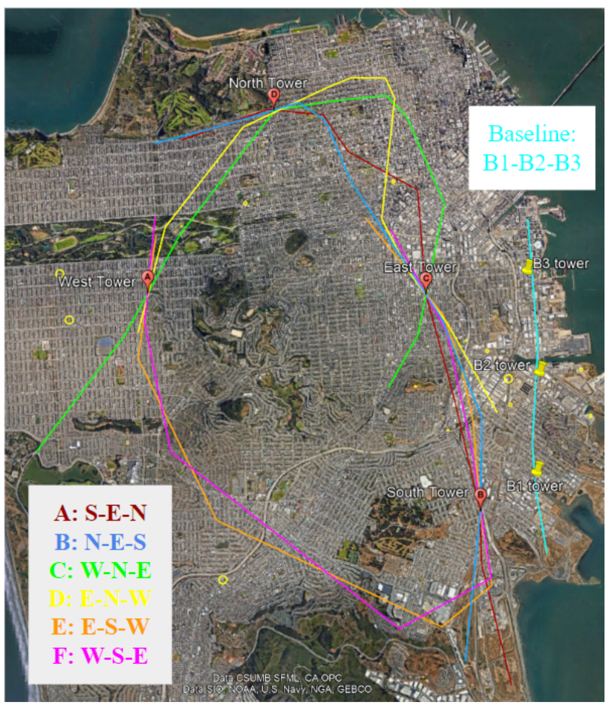

Figure 3.5: Route Overview. Each flight route (colored lines) encounters three towers. The light blue baseline route coincides with the three baseline towers (yellow pins). Red place markers show the location of the North, East, South and West towers.

3.3.1.3 Ambient Visual Conditions and Sensor Visualization

Routes were flown in three simulated Ambient Visual Conditions (AVC): day, dusk, and night. Settings for each AVC are shown in Table 3.3. The entire out the window view simulated the selected sensor visualization. The base of each opaque cloud layer was measured from the terrain height at the initiation point for each flight route. This terrain altitude ranged from 30 to 240 ft AGL. Day- and night-time weather simulated the lower bound of Marginal Visual Flight Rules (MVFR): 3 Statute Miles (SM) visibility with 1,000 ft ceiling [121]. This ceiling afforded cloud clearance as the coupled helicopter traversed the rolling terrain.

Table 3.3: Simulator settings for each sensor visualization.

AVC	Sensor visualization	Cloud layer (ft AGL)	Visibility (SM)	Sun elevation (deg)
Day	Unaided IR	1,000-2,500	3	168
Dusk	Unaided IR	1,000-6,000	1.2 1.7	5
Night	NVG IR	1,000-2,500	3	-11 168

Dusk AVC approached the permissible limit for VFR flight. Illumination was reduced by thickening the cloud layer and lowering the sun elevation to just prior to sunset. Visibility was reduced towards the 1 SM minimum for Special Visual Flight Rules (SVFR) operations [122]. With the same reduced visibility setting, tower objects could be seen farther with the unaided visualization than they could with the IR visualization option. However, since IR imaging has the ability to sense radiative surfaces (such as metal) despite obscurations [123], the dusk-IR visibility was increased so that towers and other terrestrial objects were visible at a slightly farther distance. All sensor visualizations were qualitatively validated by USAARL pilots with experience flying in platforms with IR and NVG sensors. The dusk sensor visualization is shown from the right seat in Figure 3.4.

The night AVC was paired with a simulated gray scale IR or green scale, image intensifying low light camera to simulate night vision goggles (NVGs). Solar twilight (-11 deg elevation) was supplemented by a moon at 53 deg elevation. The simulator's IR simulation model required the sun elevation to remain high to present a sensor image that had the same brightness as the NVG image.

3.3.1.4 Graphic Augmentations

Towers along each flight route had either no graphic augmentation, an *a priori* yellow circle (Figure 3.6 right), or a yellow bounding box (Figure 3.6 left). The size and shape of these graphic aids are based on a U.S. Government study on Vertical Obstacle Visualizations [46]. Circles and box graphic augmentation gradually faded into view starting at 1,500 m from the helicopter and are solid at a distance of 1,000 m. This gradual appearance reduced scene clutter and salience due to sudden appearance, while imitating reasonable detection range for an airborne LiDAR sensor.

The radius of the yellow circle varied according to the reported horizontal accuracy in FAA's DOF database. Towers that coincided with the flight route with circular graphic augmentation had a 200 ft radius circle. Additionally, the center of the circle was offset so that the tower is in the left (or right, depending on direction of flight) third of the circle. This offset imitated the relative uncertainty about the exact location of the tower. The structure could be anywhere within the circle's radius, not necessarily in the center. Yellow bounding boxes, on the other hand, conveyed obstacle height and location to within 10 ft. This precision would result from onboard sensors using automatic vertical structure detection (Chapter 2). There is a significant difference in cost and complexity of relying on the circle's *a priori* data and active sensing required for generating bounding boxes. These two levels of graphic augmentation were designed to compare the effectiveness of these different information sources.

Figure 3.6: A tower augmented with bounding box (left) and *a priori* circle (right) in day AVC.

3.3.2 Tasks and Procedure

Participants first read and signed the informed consent form, completed a demographic questionnaire, then read the aircrew mission briefing in Appendix A. In the simulator, participants flew a baseline route followed by six evaluation routes. Each flight route started at, and remained

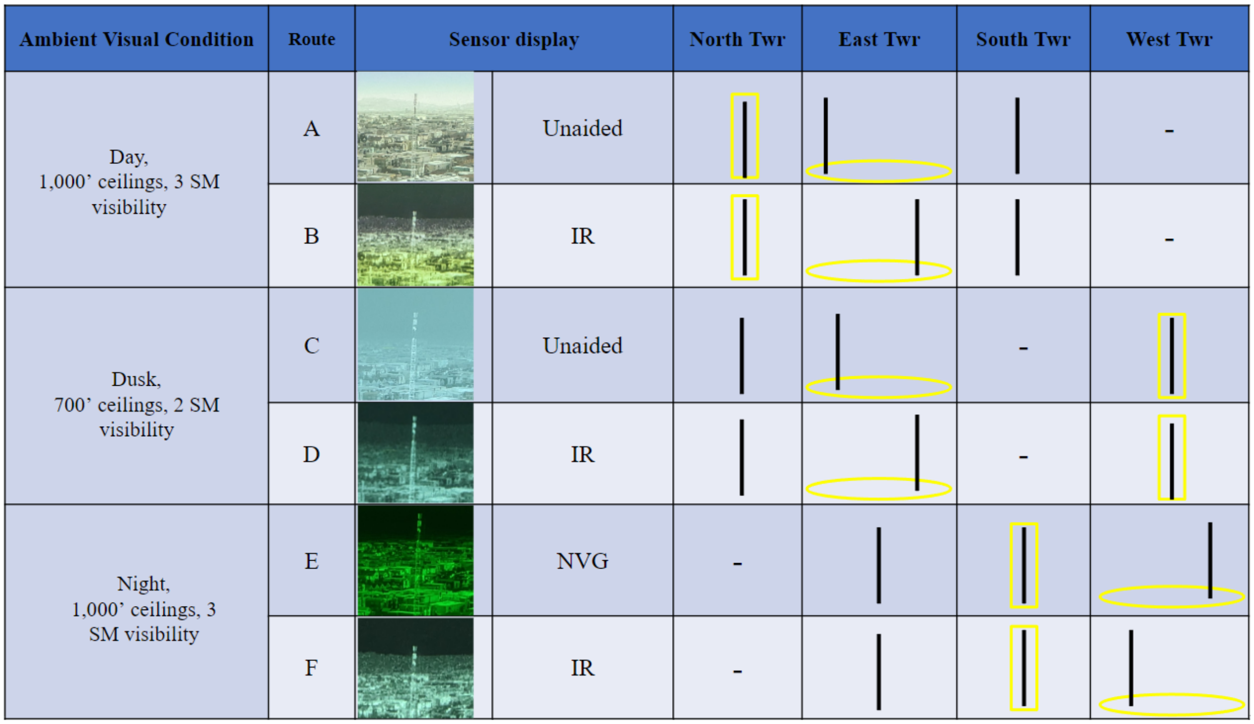

Figure 3.7: Test Conditions Matrix.

coupled to, 200 ft radar altitude and 80 knots indicated airspeed. Each route followed a course with identical way points to ensure a consistent visual experience. In the baseline route, pilots encounter three examples of the tower object in day AVC. The first tower was annotated with a bounding box, the second tower was within a yellow circle, and the final tower example has no graphic augmentation. The observer ensured the participant saw each tower example and that the Air Traffic Control recording was audible through the participant's headset. The observer also confirmed that the participant could locate aircraft state information (such as aircraft heading and fuel remaining) along with adjusting the transponder code and radio frequency in the simulator cockpit.

During each subsequent 5-7 minute long evaluation route, the participant was responsible for performing the duties of a pilot not on the controls according to an aircrew mission briefing in Appendix A. These duties included monitoring and responding to Air Traffic Control, tuning radios, and airspace surveillance. The participant was seated in the right seat. The observer in the left co-pilot seat performed pilot on the controls duties according to the aircrew mission briefing checklist. As foreshadowed in the participant instructions, the observer/co-pilot was usually distracted and never noticed any of the obstacles unless prompted by the participant. Along each of the six evaluation routes, the aircraft comes within approximately 50 ft of three towers with the sensor visualization and graphic augmentation combinations as shown in Figure 3.7. Each participant was instructed to announce explicitely when they (a) noticed any graphic augmentation (circle or box) and (b) when they were able to discern the tower itself. The order of the evaluation flight routes changed for each participant so that the three AVCs and associated visualization options were counterbalanced in a Latin square design.

3.3.2.1 Experiment Design

The experiment employed a 2 (sensor visualization: unaided/naked eye or night vision goggles (NVG) and infrared (IR) imaging) x 3 (obstacle visual augmentation (OVA): none, a priori circle, or bounding box) within-subjects design. Day and dusk AVC involved either unaided/naked eye or IR sensor visualization while night AVC employed NVG or IR.

3.3.2.2 Dependent Measures

Dependent measures, shown in Table 3.4, included obstacle detection time (seconds before impact) and additional metrics to monitor workload and task management throughout each flight. Participant eye movement and world view was recorded with a Pupil Labs Core headset [124]. The five-point computer screen calibration was considered adequate when it achieved gaze accuracy of less than 2 degrees. The world view camera recorded 1080p video at 30 Hz while the eye camera recorded at 200 Hz. The Pupil Player software showed gaze location on the world video. A fixation detector plug in found fixations for each data file that met angular dispersion and minimum and maximum duration. These fixations were used to determine whether, when, and for how long each participant fixated on each obstacle. Obstacle detection time was the difference between the time that the participant fixates on a tower in the world video and the moment that the tower comes within the rotor disc of the helicopter on the eye tracker's world view video. Based on previous studies [125, 126], we defined a fixation as having a duration between 200-400 ms and maximum angular dispersion of 2 degrees. In addition to using eye tracking data, obstacle identification was confirmed with coincident audio files. Our pilot study revealed that eye tracking data, on its own, did not reveal exactly when a participant saw a tower and did not delineate when a participant saw a graphic augmentation or the associated tower. We therefore recorded coincident audio files throughout the experiment to capture when participants announced verbally the detection of an augmentation feature or a tower. The fixation which preceded an annunciation of a graphic augmentation or tower determined the participant's detection time for that visual feature. The observer also annotated whether the participant correctly responded to simulated air traffic controller calls and avoided each of the 18 obstacles.

All routes were designed such that towers would be detectable only within 1,500 m (approximately 36 s, at 80 knots ground speed), due to their small cross section and placement within urban

Table 3.4: Dependent variables and outcomes.

Outcome Measured	**Dependent Variables**
Task management	Completion of workload inject tasks
Obstacle identification / detection time	Verbal announcement of augmentation / obstacle detection
	Fixation on object / tower
Workload	Post-flight Survey (Appendix B)
Visualization preference	

clutter. However, a significant number of participants unexpectedly detected some towers beyond 1,500 m (approximately 40 s prior to collision) during day AVC. Since the obstacle placement and the tower itself were identical between participants, these early detections may have been the result of inadvertent unequal opportunities to notice the object when it was approached from different directions. For example, towers that had detections especially far out, such as the East tower within the circle with an IR sensor ($M = 38.3s$), were centrally located in the participant's scan sector field of view over 2,700 m prior to the encounter. The same tower was detected much closer ($M = 19.6s$) during approaches from the South, likely due to a small hill which partially masked the East tower until the aircraft was within 1,500 m of the vertical structure. The two approach paths for this tower are shown in Figure 3.5. The additional (29 s) time in the central field of view provided participants the opportunity to detect the tower farther out when it was momentarily silhouetted or otherwise more salient against the dynamic background. Two tower encounters in day AVC and one tower encounter in night AVC benefited from this unequal exposure and were excluded from detection time analysis.

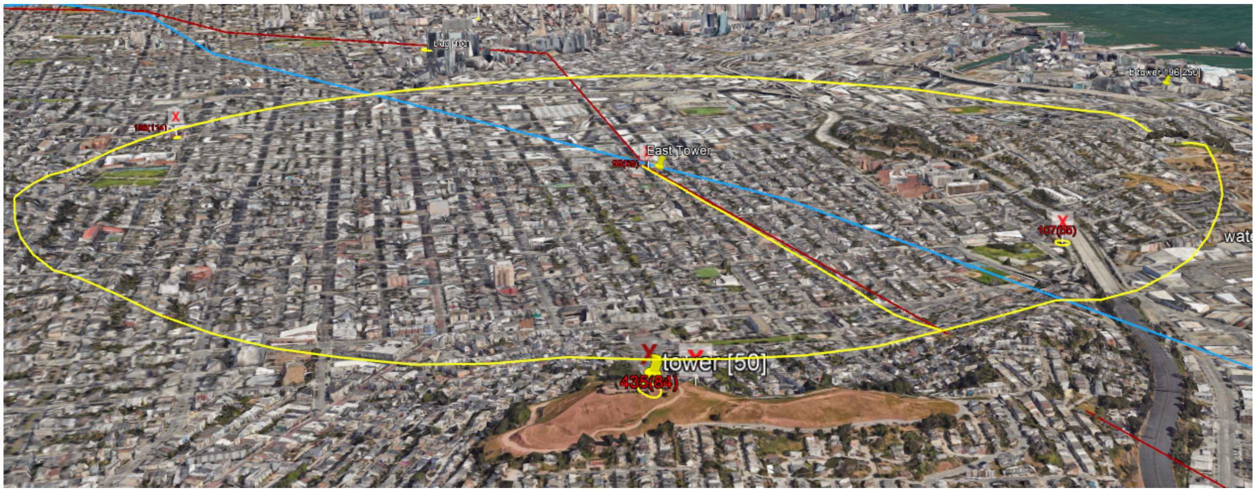

Figure 3.8: Comparison of flight routes A (red line) and B (blue line). Route A flew from the top left of the scene towards the bottom right. Route B flew from the bottom right to upper left. The yellow circle is centered on the East tower and has a radius of 1,500 m. The red route A passes over a small hill which partially obscures the East tower outside of the yellow ring.

Data from this study were analyzed using generalized linear mixed effects models in R studio [127]. The models were tested using the Linear Mixed-Effects Models [128] and lmerTest[129] packages. Detection time was the dependent variable. The fixed effects were sensor type (unaided, NVG or IR) and graphic augmentation type (none, circle or box). Participant ID was included as a random effect. Main effects were evaluated using Chi-squared tests between a null model (without the main effect of interest) and another model with the remaining main effects. Tukey post-hoc tests were used for pairwise comparisons.

Coincident eye tracking and audio data highlighted the delay in communicating the presence of a hazard. Depending on the cadence and workload, participants did not start to enunciate an obstacle until 4 to 5 s after their fixation. Drawn out obstacle descriptions could easily consume 15 s. Noticing a tower was typically associated with a sequence of vertical fixations at the top and base of the tower. Gaze also tended to shift away from the tower before the participant started describing the location. Pilots commonly overestimated the distance to obstacles, usually resorting to stating a distance of 2-3 miles even when the obstacle was within 1,500 m. There was also a tendency for pilots to state that an encounter was ”no factor” when the tower was, in fact, dangerously close.

3.3.2.3 Detection Rates

With the exception of one tower encounter with one pilot, all participants detected all towers. The single missed tower coincided with a simulator volume control malfunction which prevented the participant from adjusting the volume of the audio recording. The participant was focused on reducing the volume knob on the center control panel and did not look outside while a tower was in view. This instance was excluded from the subsequent analysis.

3.3.2.4 Detection Times

Day Ambient Visual Conditions The detection times for the East and North tower that benefited from the unbalanced route designs are accentuated with an orange arrow in Figure 3.9 among other day AVC detection times. These two sets of early detection times were removed from the following analysis of day AVC data.

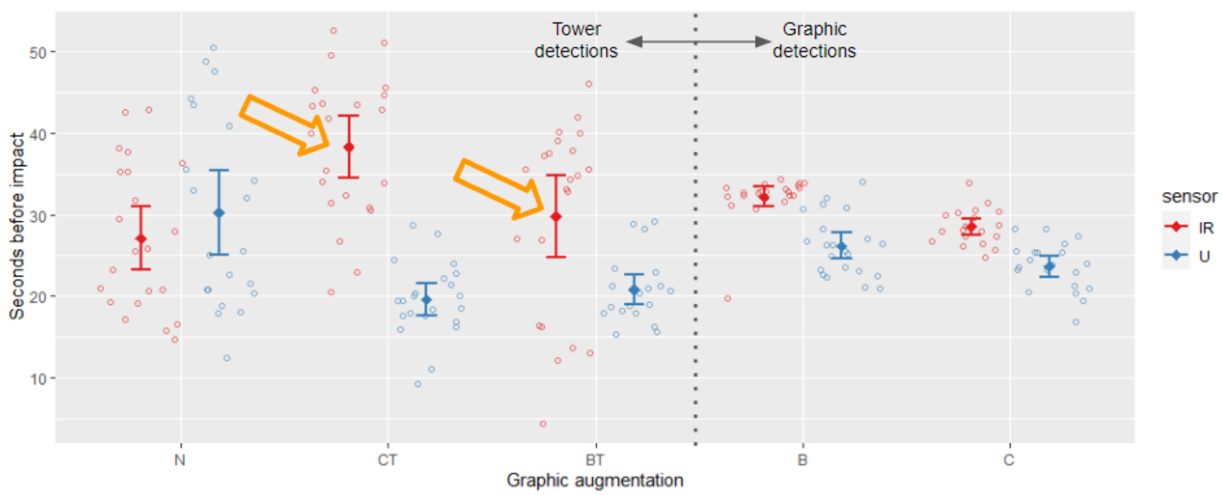

Figure 3.9: All detection times during day AVC for towers with no augmentation (N), tower within circle (CT), tower within box (BT), circle (C), and box (B). Tower detection times that were affected by route design are annotated with an orange arrow.

Sensor visualization did not affect tower detection time but there was a main effect of graphic augmentation on tower detection time, as shown in Figure 3.10 ($\chi^2(4) = 35.2, p < 0.0001$). Due to the unbalanced routes, detection times for graphically augmented towers (CT and BT) are presented only for the unaided visualization. The mean detection time for a tower without graphic augmentation ($M = 28.7s, \beta = 10.7, SE = 1.7, 95\% CI = [7.4, 14.0], t(141.0) = 6.3, p < 0.0001$) was faster than for towers within a circle ($M = 19.6s, \beta = -10.7, SE = 1.7, 95\% CI = [-14.0, -7.4], t(141.0) = -6.3, p < 0.0001$) or within a box ($M = 20.8s, \beta = -9.5, SE = 1.8, 95\% CI = [-12.8, -6.1], t(141.8) = -5.4, p < 0.0001$). Detection times did not differ significantly between towers within a circle and those within a box.

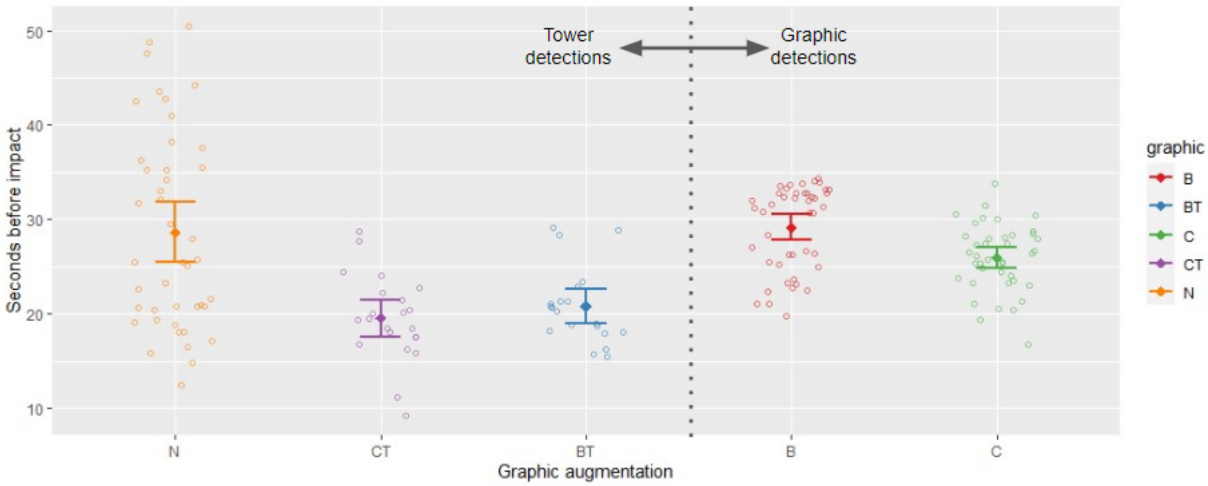

Figure 3.10: Tower and graphic detection times during day AVC for towers with no augmentation (N), tower within circle (CT), tower within box (BT), circle (C), and box (B). IR-CT and IR-BT detection times are excluded due to excessive exposure.

Regarding graphic augmentation objects, circle detection time ($M = 26.0s, \beta = 12.0, SE = 2.5, 95\%CI = [7.3, 16.7], t(141.4) = 4.9, p < 0.001$) was slower than unaugmented tower detections. The box was seen only slightly before ($M = 29.1s, \beta = 15.8, SE = 2.4, 95\%CI = [11.1, 20.4], t(140.8) = 6.5, p < 0.001$) a tower without graphic augmentation.

There was also an interaction between graphic augmentation and sensor visualization (Figure 3.9, $\chi^2(2) = 16.7, p = 0.0002$), such that the detection times for the circle ($M = 28.4s, \beta = 7.9, SE = 2.5, 95\%CI = [3.2, 12.7], t(141.2) = 3.2, p = 0.003$) and the box ($M = 32.1s, \beta = 7.9, SE = 2.4, 95\%CI = [4.5, 13.8], t(140.8) = 3.8, p = 0.0004$), but not for the unaugmented towers, were significantly faster in IR compared to the unaided condition. A possible interaction effect for towers within circles (CT) and towers within boxes (BT) could not be explored. Those towers had to be removed from the analysis because of prolonged exposure due to uneven route design.

Dusk Ambient Visual Conditions Dusk reduced ambient illumination, lowered cloud levels to just above the flight altitude and decreased visibility to 2 Statute Miles. Simulating these Special Visual Flight Rule conditions provided an opportunity to compare the effectiveness of sensor and graphic augmentations at the lower limit of permissible visual flight. Unlike with the day and night AVCs, dusk did not offer opportunities to perceive the designated towers beyond 1,500 m when they were particularly silhouetted or centrally located in the participant's field of view. Therefore, data for all towers were included in the analysis, presented in Figure 3.11.

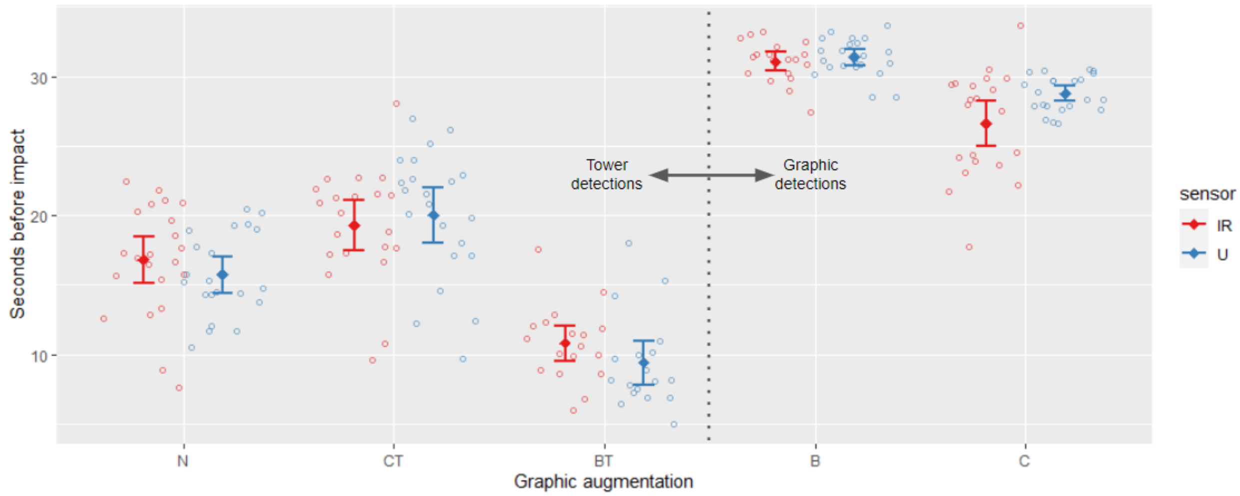

Figure 3.11: Dusk results with sensor and graphic type delineated. Mean detection time for each sensor and graphic combination is denoted by diamond, standard error is shown with brackets. Graphics and objects noticed are tower without any graphics (N), circle (C), tower within circle (CT), box (B), and tower within box (BT).

There was no main effect of sensor visualization on detection times in dusk AVC ($\chi^2(1) = 0.2, p = 0.68$). However, there was a main effect of graphic augmentations ($\chi^2(4) = 386.1, p < 0.001$) on detection time during dusk AVC, as shown in Figure 3.12. Post-hoc pairwise comparisons revealed that towers within a circle (CT) were seen significantly sooner than unaugmented towers ($M = 19.7s, \beta = 3.4, SE = 0.7, 95\%CI = [2.1, 4.7], t(176) = 5.1, p < 0.001$). However, the tower within the box (BT) was seen later than a tower without augmentation ($M = 10.1s, \beta = -6.2, SE = 0.7, 95\%CI = [9.0, 11.2], t(178) = -8.8, p < 0.001$).

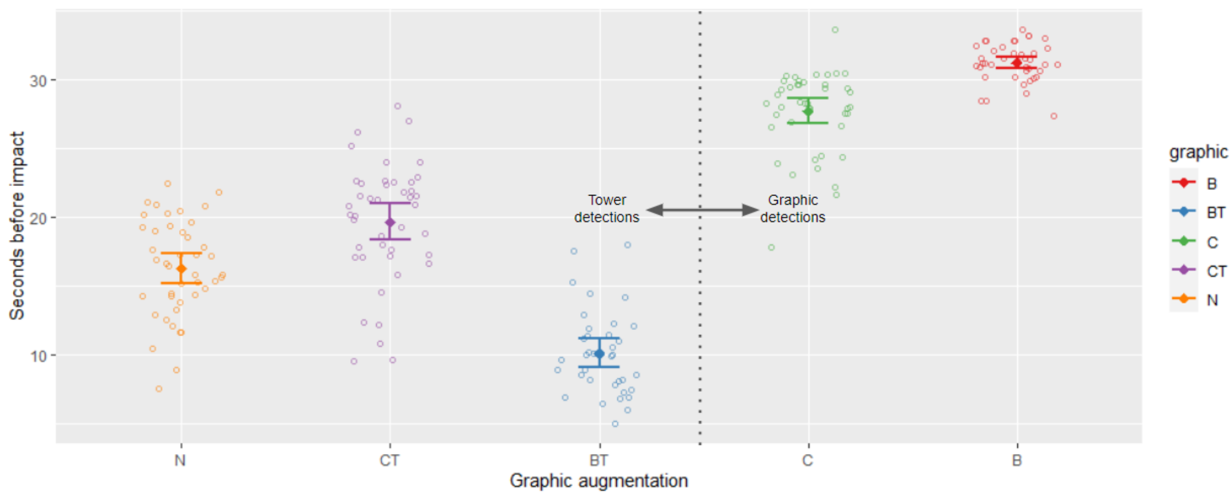

Figure 3.12: Dusk results consolidated by graphic type. Mean detection time for each graphic combination is denoted by diamond, standard error is shown with brackets. Graphics and objects noticed are tower without any graphics (N), circle (C), tower within circle (CT), box (B), and tower within box (BT).

For graphic augmentation objects, the detection time for the circle (C) was faster ($M = 27.7s, \beta = 11.4, SE = 0.7, 95\% CI = [10.1, 12.8], t(176) = 17.0, p < 0.001$) than for unaugmented towers. Boxes (B) were seen prior to circles and farthest away from each tower encounter (($M = 31.3s, \beta = 15.0, SE = 0.7, 95\% CI = [13.6, 16.3], t(176) = 22.1, p < 0.001$)). There was no significant interaction between graphic augmentation and sensor visualization ($\chi^2(4) = 9.2, p = 0.06$).

Night Ambient Visual Conditions Like day AVC, night AVC also offered 3 SM visibility and opportunities for detecting towers beyond the expected 1,500 m range due to unbalanced route design. In this case, unequal exposure was present only for the West tower (annotated with an arrow in Figure 3.13) which offered a particularly silhouetted perspective when approached from the north. Figure 3.13 shows the disparity between IR and NVG detection times for the tower within the circle (CT) where nearly half of participants detected the tower before the augmenting circle ever became visible. Due to this unequal exposure, notice times for the tower within the circle with the IR sensor were excluded from the following analysis.

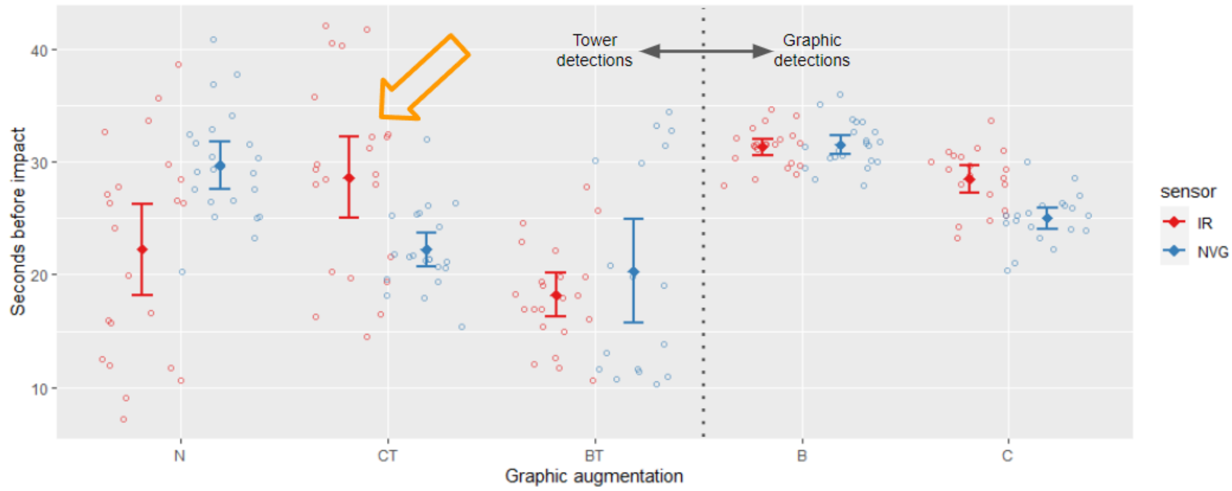

Figure 3.13: Night results with sensor and graphic type delineated. Unbalanced route annotated with orange arrow. Mean detection time for each sensor and graphic combination is denoted by diamond, standard error is shown with brackets. Graphics and objects noticed are tower without any graphics (N), circle (C), tower within circle (CT), box (B), and tower within box (BT).

There was a main effect of sensor visualization ($\chi^2(1) = 4.0, p = 0.05$). NVG detection times were slightly faster ($M = 26.5s$) than IR detection time ($M = 25.0s, \beta = 2.0, SE = 1.07, 95\%CI = [-0.05, 4.1], t(146) = 1.9, p = 0.05$). There was also a main effect of graphic augmentation ($\chi^2(3) = 87.3, p < 0.001$). The detection time for the unaugmented towers ($M = 25.9s$) was slower than for the box ($M = 31.4s, \beta = 5.4, SE = 1.0, 95\%CI = [3.1, 7.8], t(139) = 5.4, p > 0.001$) but faster than for the tower within the box ($M = 19.1s, \beta = -6.8, SE = 1.1, 95\%CI = [-9.4, -4.5], t(140) = 6.5, p > 0.001$). Circle detection time was not significantly different from unaugmented towers.

The was a significant interaction between graphic augmentation and sensor visualization in this subset of night AVC data ($\chi^2(3) = 30.2, p < 0.001$). The unaugmented towers, but not the box nor the tower within the box, were detected significantly sooner with NVG ($M = 29.7s$) than they were with IR sensor ($M = 22.2s, \beta = -7.5, SE = 1.4, 95\%CI = [-11.5, -3.0], t(139) = -5.2, p > 0.001$).

Table 3.5: Overall results for tower and graphic detection time with corresponding expectations. Since all towers were detected, no factor affected detection rate.
*denotes unbalanced routes where one approach direction was excluded from analysis

		Results (detection time, seconds before impact)		
Expectation		**Day**	**Dusk**	**Night**
1. Obstacles will be noticed sooner and/or more reliably with unaided/NVG video, compared to IR video		There was no main effect of sensor visualization on tower detection time		Towers were detected sooner with NVG visualization
2. Obstacles augmented with bounding boxes will be noticed sooner (and/or more reliably) than those with a priori circles, followed by those without cueing	1. B / C	B (29.1)	B (31.3)	B (31.4)
	2. BT	N (28.7)	C (27.7)	C (26.7)
	3. CT	C (26.0)	CT (19.7)	N (25.9)
	4. N	BT* (20.8)	N (16.3)	CT* (22.2)
		CT* (19.6)	BT (10.1)	BT (19.0)
3. Obstacles augmented with a circle or box will be noticed sooner (and/or more reliably) with IR video than with the corresponding unaided/NVG video		No effect on tower detection time		Box with NVG resulted in earlier tower detection
		Circles and boxes were noticed sooner with IR	Only circle seen sooner with IR	

Effect of age and experience on tower detection time Participants' age and experience level may have affected tower detection time. Figure 3.14 presents the average tower detection times for the unaided- and IR-unaugmented (N) towers, IR-tower within circle (CT), and IR-tower within box (BT) along with participants' self-reported age and flight hours. Note that other towers are not included since they were affected by unbalanced route design. Detection times did not differ significantly as a function of age, total flight hours and NVG flight hours .

3.3.2.5 Post-Flight Survey

Participants completed the post-flight questionnaire (Appendix B) after completing all routes. All participants overestimated the number of unique towers they encountered, with estimates ranging from 5 to 30 ($M = 13.5, SD = 6.9$). Participants reported having the least amount of mental

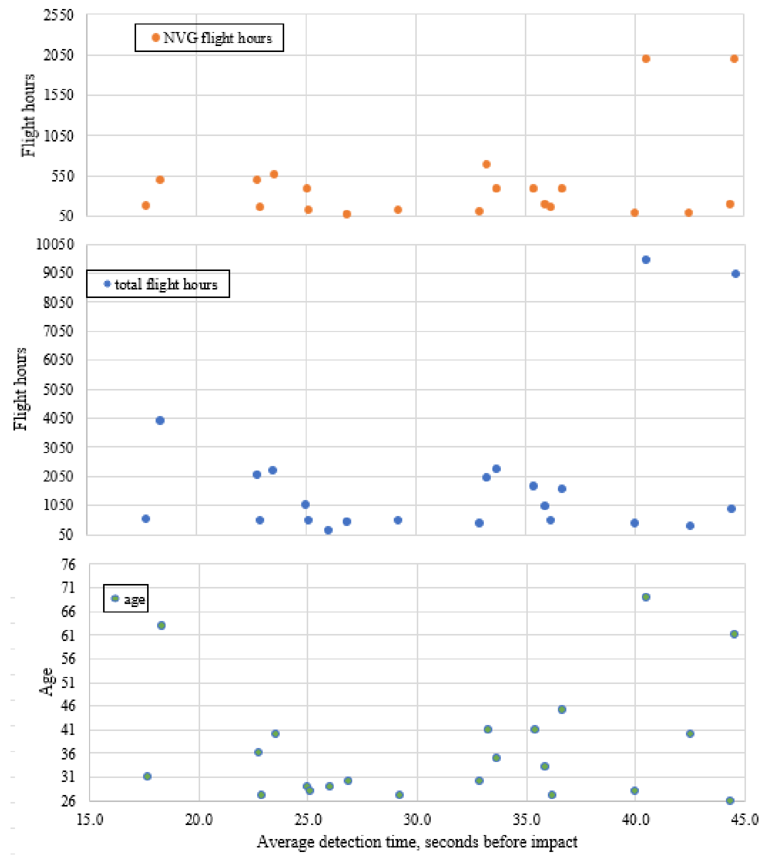

Figure 3.14: Average detection time for towers in day AVC plotted against flight hours (top), NVG flight hours (middle), and age (bottom).

capacity when responding to requests from Air Traffic Control (Figure 3.15).

	Average across all routes	When identifying an obstacle	When responding to ATC
1 (plenty of capacity)	18.2%	18.2%	0.0%
2	36.4%	40.9%	18.2%
3 (some capacity)	40.9%	36.4%	54.5%
4	4.5%	4.5%	27.3%

Figure 3.15: Reported mental capacity.

The highest external awareness (proximity to buildings and terrain) (Figure 3.16) was experienced during day AVC with the unaided visualization. The lowest awareness was reported for dusk-unaided.

In day AVC, participants stated that they had slightly more awareness of the external environment with unaided visualization over IR on a scale of 0 to 100 (Figure 3.17, where 0 was unaided and 100 was IR) ($M = 46.8$).

In dusk AVC, participants stated that they had more awareness of the external environment with the IR over the unaided visualization (Figure 3.18) ($M = 69.0$).

	Dusk-Unaided	Day-Unaided	Day-IR	Dusk-IR	Night-NVG	Night-IR
1 (low)	4.5%	0.0%	0.0%	0.0%	0.0%	0.0%
2	9.1%	0.0%	0.0%	9.1%	0.0%	0.0%
3 (moderate)	50.0%	9.1%	27.3%	36.4%	27.3%	40.9%
4	31.8%	59.1%	68.2%	45.5%	68.2%	50.0%
5 (perfect)	4.5%	31.8%	4.5%	9.1%	4.5%	9.1%

Figure 3.16: Reported external awareness during each AVC/sensor combination.

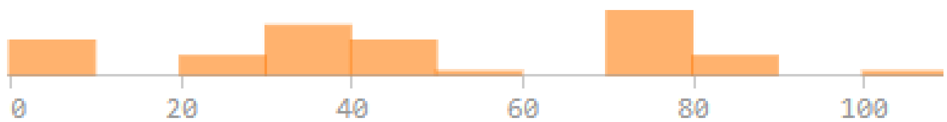

Figure 3.17: Preference between unaided and IR visualization for providing awareness of the external flight environment in day AVC (0=unaided, 100=IR).

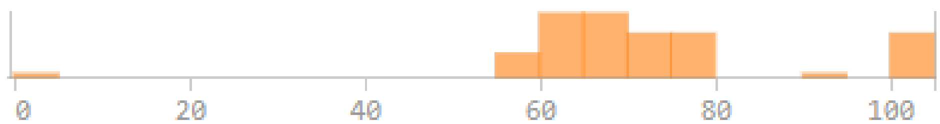

Figure 3.18: Preference between unaided and IR visualization for providing awareness of the external flight environment in dusk AVC (0=unaided, 100=IR).

In night AVC, participants stated that they had more awareness of the external environment with the NVG over the IR visualization (Figure 3.19) ($M = 71.0$).

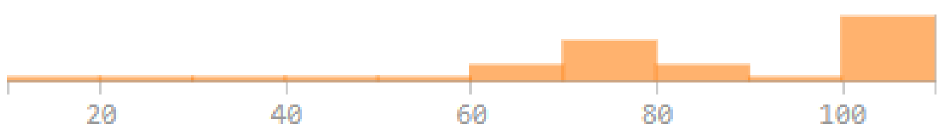

Figure 3.19: Preference between NVG and IR visualization for providing awareness of the external flight environment in night AVC (0=IR, 100=NVG).

Among all sensors and graphic augmentations, participants reported that the box provided the most useful information (Figure 3.20).

	Yellow circle	Unaided, no sensor	Gray IR	Green NVG	Yellow bounding box
2	13.6%	9.1%	9.1%	0.0%	0.0%
3	18.2%	45.5%	36.4%	36.4%	13.6%
4	31.8%	22.7%	40.9%	40.9%	13.6%
5 (very useful)	36.4%	22.7%	13.6%	22.7%	72.7%

Figure 3.20: Information quality between sensor and graphic augmentation options.

Participants reported that the boxes aligned with obstacles better than the circles (Figure 3.21).

1 (barely)	5.3%	4.5%
2	0.0%	18.2%
3 (moderate)	15.8%	31.8%
4	31.6%	36.4%
5 (perfect)	47.4%	9.1%

Figure 3.21: Alignment of graphic augmentations (boxes: left column, circles: right column) with obstacles.

Finally, as shown in Figure 3.22, on a scale of 0 to 100, where zero was the circle and 100 was the box, participants said that boxes were more useful for finding obstacles ($M = 59.5$).

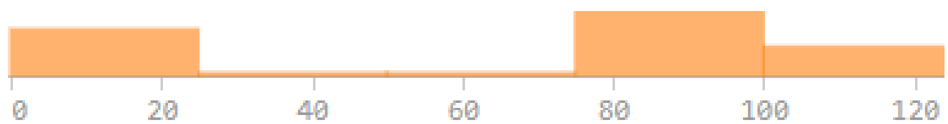

Figure 3.22: Preference between the circle and box graphic augmentation for finding obstacles (0=circle, 100=box).

3.3.2.6 ATC task compliance

The workload inject task of monitoring and responding to Air Traffic Control (ATC) was designed to be easily accomplished by the average aviator. However, given the competing attention demand and possibly higher prioritization of the tower detection task, only 2 out of 22 participants correctly responded to all radio calls. One participant spoke English as a second language and was excluded from the summary in Table 3.6.

Table 3.6: Radio call response summary for 21 participants.

AVC	Route	Air Traffic Control Requests	Correct response: Mean	Correct response: Standard deviation	Correct response: Percent
Day	A	3	2.6	0.9	86
	B	4	3.9	0.3	98
Dusk	C	4	3.6	0.6	89
	D	5	4.3	0.9	86
Night	E	7	6.3	0.7	90
	F	7	6.3	0.8	90

3.3.3 Discussion

The current experiment examined the effectiveness of visualizations (unaided, thermal/infrared or image intensification/night vision goggle) and graphic augmentations (none, *a priori* circles, or sensor-informed boxes) for supporting obstacle detection in three Ambient Visual Conditions (AVCs).

Except for one participant who missed a tower due to a simulator malfunction, all obstacles were noticed before coming inside the rotor disc of the aircraft. Obstacle detection times differed considerably across sensor visualizations and graphic augmentations (4.4 to 52.6 s before impact). The overestimation of the number of towers indicated that there was a low learning effect. Sensor visualization type affected detection time only during night AVC where participants detected towers sooner with the NVG sensor than they did with the IR sensor. Bounding boxes, one form of graphic augmentation, were detected faster than circles and unaided towers in most cases. However, subsequent detection time for the tower itself, within the box, always lagged behind detection time for an unaugmented tower. Towers within circles were seen prior to towers within boxes during dusk and night AVCs.

The significant amount of time (5 to 15 s) that elapsed from when a pilot first fixated on an tower to when they finished describing the obstacle and its relative location showed the need for early obstacle detection that compensates for delayed response or recognition. This delay would likely be exaggerated in an operational environment where pilots are less primed and focused outside the aircraft. The time required to communicate obstacle presence also reveals the limitation of current airspace surveillance procedures which rely on manual identification followed by spoken descriptions. A previous study [34] which also evaluated graphic augmentation effects in a helicopter flight simulator had much faster response times, but also severely restricted visibility which prevented pilots from clearing airspace more than 20 s past their current position.

Pilots often overestimated the distance to hazards, typically declaring an obstacle was two to three miles away when it was actually less than 1,500 m. This inaccurate perception could have been influenced by the simulation visualization. Nevertheless, this overestimation reveals that, even with a nearly constant ground speed and altitude and moving map display, pilots' sense of distance can be exaggerated in the flight environment. This discontinuity between perceived and actual distance to a hazard could cause confusion in future cockpits where automated alerts will state actual, not perceived, distances to objects of interest.

There was no significant correlation between age or flight experience and obstacle detection time. This insight reinforces the conclusion from a previous study [130] that found that more experienced pilots are not less likely to have wire strikes. In other words, more flight experience does not necessarily translate into better attention management skills.

3.3.3.1 Overall Impact of Graphic Augmentation

The significant delay in detection time for towers within boxes showed that the graphic augmentation meant to highlight the tower object tended to obscure the hazard itself. In addition to the proximity of the box lines, the relative brightness of the yellow color also delayed identification of the less salient object within the box. This finding emphasizes the need for careful trade off decisions in design. The right level of relative salience needs to be determined empirically to ensure reliable attention capture by the box without masking of the actual object of interest (the tower). Masking refers to the situation where the visibility of a target stimulus (in this case, the tower) is decreased by presenting it in close spatial and temporal proximity to a so-called 'mask' (the graphic augmentation). Post-flight surveys revealed that participants thought the boxes were most useful for the airspace surveillance task. This preference may result from two factors: first, it may indicate that the expectancy [115] of a potential obstacle conveyed by the highly salient boxes was highly valued and overshadowed the tendency of the box to obscure the obstacle it contained. In addition, and in contrast to the circle, the box provided a more precise indication of the tower. Another study [35] using a Helmet Mounted Display found that a precise cue pointing at the target decreased detection time.

Augmenting obstacles with a highly salient and precise cue runs two risks, however. First, pilots may get fixated on the salient box while trying to discern its message and contents. This fixation will reduce the amount of attention paid to the rest of the environment and could cause them to miss other unaugmented hazards due to inattentional blindness [15], similar to the fixation on a SVS display [33]. The second risk is that a high level of salience and precision may result in over reliance on the cue which can be problematic in the context of imperfectly reliable underlying automation [131]. Compelling graphics also have the potential to overwhelm the pilot, as was shown in previous studies [43, 44, 45].

Despite the less definite nature of the circle graphic augmentations, detection time for the tower within the circle was significantly faster than identification of towers within boxes. This suggests that the circle may be the preferable design, despite pilots' preferences [132]. In the end, the goal of the augmentation is to support the detection of towers, not of surrounding symbology. The circles did not mask the tower, and they may be less compelling and thus avoid automation over reliance.

3.3.3.2 Overall Impact of Sensor Visualization

With the exception of the NVG option, sensor visualization did not affect tower detection time. Unbalanced tower encounters which were excluded from the analysis tended to have early tower detections due to momentary silhouetting of the tower against a dark sky in the IR visualization. Flying low to have a better perspective was mentioned in the focus groups (Section 3.1) as a flawed, albeit effective, detection strategy.

Across all AVCs, the horizontally oriented circle was seen sooner with the IR visualization than it was in the corresponding visualization option. This faster detection time shows the increased noticeability when a colored symbol is placed against a gray scale background. Decreased search time for contrasting colors is in line with the Guided Search model [118].

3.3.3.3 Impact of Graphic Augmentation and Sensor Visualization During Different AVCs

Participants expressed an approximately equal preference for the unaided and IR visualization in day AVC. This coincided with similar obstacle detection times between the visualization options. Graphic augmentations had the least impact on tower detection time during day AVC. The IR visualization enabled earlier detection of the graphics themselves against the gray scale background. These results show that in a rich daytime visual environment, there is less utility for graphic augmentations. In dusk AVC, there was also no main effect of the sensor display despite the increased visibility range (1.2 SM for unaided, 1.7 SM for thermal). The similar detection times countered our hypothesis that the thermal sensor's ability to penetrate light obscurations will result in faster detection time. Detection time is not simply a function of range but is also affected by color and more apparent shapes with the unaided visualization. This tendency to detect towers with the unaided visualization aligns with the Guided Search model [118] which predicted faster search times when a target shares fewer features with its surroundings. The same model also offered that finding an item in clutter (such as the urban environment which surround the hazardous towers) is faster when the target (the tower) and distractors (the buildings and other clutter) share fewer features. The colorful unaided visualization provided an additional contrast than the size and shape which was also present in the IR visualization. During night AVC, tower detection times with the NVG sensor were slightly faster than with the IR visualization. This faster detection time shows that, given equal visibility (3 SM), the slightly higher resolution from the NVG's image intensifying visualization provided an opportunity for the participants to discern tower shapes sooner. Both visualizations were monochromatic, which indicates that tower colors and textures which were present in the unaided visualizations, have less of an impact when the target (the tower) has a small cross section.

3.3.3.4 Implications for Obstacle Visualization Systems

Obstacles augmented with graphics were not always seen before unaugmented obstacles. The distinction between less definite *a priori* circles and tighter bounding boxes was less pronounced than anticipated. Quantifying the potential benefit of a precise though possibly less salient bounding box will inform the cost-benefit analysis of the increased cost and complexity of adding an onboard sensor. There was a significant difference in detection time when visibility was reduced towards the minimum allowable visual flight conditions during dusk AVC. Graphic augmentation clearly contribute to faster detection time in these conditions. An automated obstacle alert or avoidance system could take this visual detection performance into consideration by only alerting pilots or taking action when an obstacle is within its sensing range (and thus should have already been noticed by the human operator).

Visual detection can also be hindered by the large amount of visual stimuli. Tactile and audio cues draw visual attention to an area of interest if a potential hazard exceeds some critical value and have been implemented to warn of stalls and other exceedances during flight [133]. Cross-modal cueing can guide visual attention via signals in other sensory channels [134]. Tactile cueing was shown to increase hover performance over a moving target in a simulator [135]. Tactile and audio cueing was also evaluated during high-workload approaches to hover in severely degraded visual conditions [136, 137]. Neither modality affected flight performance. Tactile cueing, however, was the least preferred cueing mode and was ”distracting or frustrating”. Occasional alerts for discrete obstacle encounters should avoid this overwhelming tendency. Tactile and audio cueing shows potential to guide visual attention via signals in other sensory channels.

3.3.4 Limitations

The study involved some limitations. Not all participants responded correctly to radio calls, meaning that workload and task interference may not have been as high as intended. For some participants, this could have been because they were listening but not comprehending the radio conversation. Other participants could have chosen to ignore the radio calls despite the imperative to respond to ATC. Despite our instructions, pilots may have prioritized tasks differently and devoted more attention to looking for towers than they otherwise would have. Another limitation is that the simulation was set to a moderate 80 knots ground speed and had mild turns due to limited coupled flight control authority in the simulator. This closure rate provided a generous amount of time for participants to see the obstacles. Finally, given the high-fidelity nature of the simulator and terrain database, and given the need to counterbalance the order in which towers were experienced across routes, their visibility could not be controlled perfectly.

3.3.5 Summary

This study compared the effectiveness of different means of supporting target detection in cluttered environments. Advanced Aerial Mobility and other near term, low altitude concepts increase the density of air traffic without relieving the fundamental requirement to see and avoid obstacles while monitoring numerous information channels. Our study examines ways to increase efficiency of this essential task and has the potential to apply more broadly to visual scanning in other time critical domains.

CHAPTER 4

Efficient Vertical Structure Correlation and Power Line Inference

The vertical structure localization methods in Chapter 2 enable us to develop and evaluate approaches for comparing found vertical structures with current data sets. Databases with improved overall position accuracy allow us to infer power line presence. Efficiently consolidating the existing albeit imperfect data is essential given the difficulty of finding vertical objects at range (which can lead to false negatives) coupled with the imperative to avoid nuisance alerts caused by false positives [138, 139].

Wire strikes cause even more accidents than collisions with vertical structures [140] due to their near invisibility to the naked eye and sensors. In addition to their small cross section, power lines are hard to sense due to their irregular catenary shape, occlusion from parallel wires, movement due to wind and sagging, irregular point density and highly varied background. The vast majority of power lines are less than 200 ft above ground level, owing to the fact that most transmission towers are less than 200 ft tall [141]. Finding and annotating these lines is not a priority for agencies focused on "obstructions to air navigation."

The goals of this chapter are to first efficiently consolidate vertical structure position information then find power lines based on this structure locations. This chapter offers a method for efficiently correlating and updating existing vertical obstacle databases with new observations. Next, towers within the updated database are compared to infer the presence of power lines. Specific contributions are:

- A new method to efficiently correlate vertical structures,
- A novel approach to reliably finding potentially hazardous wires based on their arrangement, proximity and similarity,
- Evaluation of these approaches against the current Delaware Digital Obstacle File

Section 1 summarizes related work. Section 2 provides an overview of of the database update and power line inference methods. Section 3 describes the FAA Digital Obstacle File used for our experimentation. Section 4 presents results of our methods. Section 5 discusses overall performance of our method, while Section 6 offers conclusions.

4.1 Background

4.1.1 Current uses of vertical structure data

There are several examples where large amounts of raw information is organized into useful maps. This integrated data can then inform motion planning or other forms of action selection and decision making. Towers (from wind turbines to transmission towers) are an essential part of the growing information [48] and electrical [49] infrastructure around the world. Monitoring the condition of these structures is a continuous process that requires tremendous efforts. Due to the ever changing nature and sheer quantity of infrastructure, there have been several efforts to automate the mapping process. A crowd-sourcing approach [142] has been proposed but is no longer online. Recent efforts use night-time lighting patterns in satellite imagery to predict infrastructure position to within 1,000 m 70% [50] to 75% [51] of the time.

Another ground based application is vehicle localization based on pole location. The large data scale requires that the area of consideration be down sampled due to onboard computer system storage and computation time constraints. Detecting numerous poles that maintain the same physical characteristics over time provides an opportunity to match detected pole patterns with an *a priori* map [52]. The technique in [52] used quantization to compare an average of 18 poles within 50 m of the moving vehicle with an *a priori* pole map with one meter accuracy to locate a vehicle's position. A more recent mapping approach minimizes residual error between extracted pole and road curb points in successive frames to create a local feature map [53]. Neither of these techniques update the high accuracy *a priori* pole map. Both approaches rely on high resolution at close range for efficient correlation. Refs [54, 55] project extracted points to the ground plane as squares. If the squares overlap, they are combined into a single pole entity. Poles that are not seen repeatedly are removed from the map using this sliding window. The extracted poles are compared with a 30 m square reference map for localization. Recent approaches that use semantic labelling for localization and mapping rely on continuous, dense surfaces that are associated with rich imagery [56].

Above the ground level, the Federal Aviation Administration's (FAA's) Digital Obstacle File (DOF) is the definitive, publicly available source for vertical structures that could be a hazard to flight operations. Currently, adding and revising man-made obstacles that are far away from airports with instrument approaches largely relies on voluntary reporting from infrastructure builders. These reports are occasionally supplemented by imprecise observations from aircrews or ground personnel. Reported vertical structures are manually entered into the database and are seldom confirmed by onsite inspections or correlation with other data sources. High resolution 3D data, such as LiDAR point clouds, is not allowed to be a primary source of obstacle information [143]. Obstructions greater than three miles from designated airports that are less than 499 ft above ground level are not considered "obstructions to air navigation," [47] making low-altitude obstacles especially prone to oversight. Figure 4.1 shows an example of the disparity in two leading vertical obstacle data bases: the Federal Aviation Administration's (FAA) Digital Obstacle File (DOF) and the National Geospatial-Intelligence Agency's (NGA) Digital Vertical Obstacle File (DVOF). Although some obstacles are mutually represented, a significant number of vertical structures are present in only one of the two databases.

Figure 4.1: Vertical obstacles in San Francisco. Towers and other vertical structures from DOF are shown with yellow pins. Obstacles from DVOF are denoted by red X's.

DVOF and DOF use cases focus on two domains: ground level and flight altitudes greater than 500 ft Above Ground Level. These focus areas omit the low altitude flight environment that is essential for traditional helicopters, emerging Advanced Aerial Mobility (AAM) aircraft, and small Uncrewed Aircraft Systems (UAS).

4.1.2 Current data structures

Efficiently representing the environment is a long-standing challenge. Several categories of metric maps are shown in Table 4.1. These data structures either compile (in that they add new observations) or match (in that they quickly aggregate new obstacle information along with uncertainty). The data structure for the ground level domain supports either short range forecasting or high level exhaustive comparison. Ref. [52] used quantization that divided the area of consideration based on radius (50 m) and position certainty (1 m). Infrastructure predictive mapping ([51]) consolidated structures along with other map features and tiles in the GeoPackage data structure [144].

The data structure for aerial obstacles must include a large geographical areas, but data latency depends on manual examination and verification. Obstacle accuracy and completeness wane at lower altitudes away from major airports. In DOF, accuracy better than 250 ft horizontally and 50 ft vertically is only motivated by the need for obstacle clearance when descending to or departing from airports [145]. As a result, over half of the DOF obstacles have a position uncertainty larger than 250 ft horizontally and 50 ft vertically.

Other work has found ways to efficiently map more general features in other safety critical applications. Recent approaches use feature matching between previously gathered point clouds and current perspectives [146]. Although this approach minimizes the memory required to match features in a variety of locations, it relies on a pre-processed, high resolution point cloud previously gathered from a viewpoint that is similar to the current 3D sensor perspective.

Table 4.1: Comparison of existing feature mapping processes. New entries are manually (M) or automatically (A) added. For mapping that absorbs found objects, new observations overwrite (O) or supplement (S) previous feature data.

	Object position accuracy	New entries treatment	Matches objects arrangement	Matches individual objects	Object correlation
DOF, GeoPackage [144]	✓	M	-	✓	S
2D Projection [55, 54]	-	A	-	✓	O
Quantization [52, 53]	-	-	✓	-	-
Predictive mapping [51, 50]	-	A	-	-	O
Point cloud matching [146]	-	-	✓	-	-

4.1.3 Power line mapping

Current wire finding methods depend on detecting wires directly. Airborne methods that automatically segment power lines rely on continuous contact [57] [58], known location [59] [60], and/or very close range [61] [62] [63]. We propose to leverage the fact that power lines and other wires are intrinsically associated with more apparent vertical structures.

4.1.4 Problem statement

The previously described approaches are vulnerable to false positives (when an already detected tower is identified as a new tower) or false negative (when a new tower is merged with an existing object). Quickly consolidating ever evolving information is another challenge. Given multiple incomplete and inaccurate databases and current observations, how can obstacles be reliably and automatically correlated and updated? Given additional position accuracy, can the tower arrangement and relationship be used to infer the presence of power lines?

4.2 Methods

Chapter 2 presented methods for efficiently locating prominent vertical structures in large point cloud data. We will compare these precise positions with the known *a priori* obstacle databases to create a more comprehensive and complete obstacle listing. Figure 4.2 presents a proposed method for correlating detected vertical structures to an existing database. If the vertical structure is not within the uncertainty bound of an existing structure, it will be added to the database as a new item. If the detected structure is within the uncertainty boundaries of a structure described in the database, they will be correlated and the database meta data will retain the more accurate position attributes.

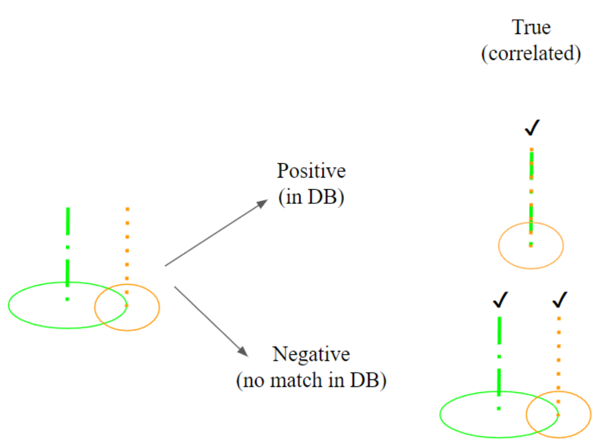

Figure 4.2: Tower correlation options. Comparing an existing vertical structure entry and associated horizontal uncertainty (green) with an observation (orange), the observation and entries are either correlated and consolidated (true positive) or the observation is found to not exist in the database (true negative).

4.2.1 Efficient database updates

We create a hash table to enable lookup and correlation with $O(1)$ complexity. Hashtables rely on a succinct and descriptive key. Our key is based on the current World Geodetic System (WGS) 84 projection. Selecting a resolution of 0.25 seconds equates to approximately 25 ft in latitude and between 19 and 24 ft in longitude for the continental United States. This resolution is within DOF's most accurate horizontal uncertainty, thus avoiding repeated entries for the same coordinates.

In degree, minute, second format, representing latitude or longitude requires 7 digits (DD MM SS.SS) when we remove the character specifying the hemisphere. We retain degrees to retain relevance throughout a hemisphere. Also, a single degree of longitude is only 48 miles at the northern section of the continental United States, which is well within the span of an aerial vehicle's operating range. Combining latitude and longitude provides a unique 16 character string.

Next, we build an index hash table (IHT). This table reads in the data fields from an existing DOF or DVOF database and adds a spatial hash key based on rounding the given coordinates up to the nearest 0.25 second. Simply searching for the nearest neighbor is not sufficient due to numerous vertical obstacles and their associated large, uneven, and overlapping horizontal uncertainties. An updated vertical structure coordinate may not be closest to its existing database location. Exhaustively searching within the overlapping uncertainties (even with the reduced spatial hash resolution) requires checking at least 9 (for 20 ft horizontal uncertainty) and rapidly increasing to over 54,000 cells with a horizontal uncertainty of one nautical mile. Given that most obstacles in DOF have an uncertainty of at least 250 ft (143 cells), the search process must be efficient.

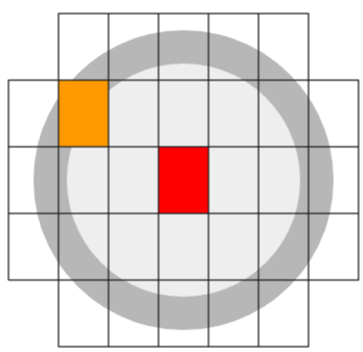

Figure 4.3: A 100 ft horizontal uncertainty is encompassed by 31 rectangular spatial hashes. The extent of the radius is imprecise due to the spatial hash's 0.25 second resolution. The 30 outlying cells encompass this uncertain circumference.

We use the IHT with spatial hash keys to create a second hash table. The purpose of this second hash table is to readily determine if a queried location is within the horizontal uncertainty of an existing vertical object. The Uncertainty Hash Table (UHT) builds a list of vertical structure hashes whose uncertainty encompasses a given spatial hash. For example, the spatial hash for the center red cell in Figure 4.3 would be the value for the 31 keys that are the spatial hash for each of the cells that encompass the provided horizontal uncertainty.

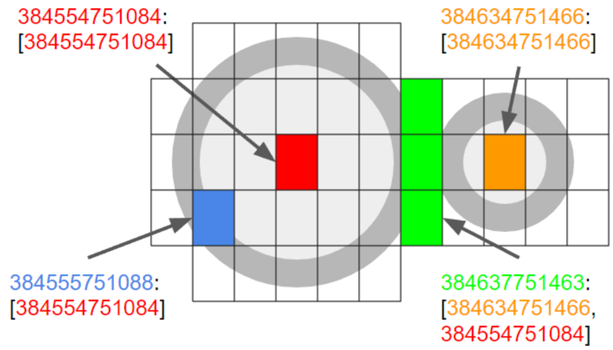

Figure 4.4: Three cells that coincide with 100 ft (left) and 50 ft (right) horizontal uncertainties are shown in green. In the Uncertainty Hash Table (UHT), the entries for the red and orange center cells of the 100 and 50 ft circles have the same spatial hash for the key and value. The off-center blue cell UHT entry has the blue spatial hash for its key, but the corresponding value is the spatial hash for the center coordinate of the 100 ft circle. Each green cell has its coordinate's spatial hash for the key, but because the cells are within the uncertainty of two vertical structure entries, each green key has two values.

When the array of uncertainty cells for an entry in the IHT overlap with another, as they do in Figure 4.4, the UHT key includes values for each relevant index. We instantiate a multimap to allow multiple entries for each key in the UHT.

After creating the initial IHT and corresponding UHT, we incorporate new observations. An observation consists of the latitude, longitude, height AGL, horizontal uncertainty, and observation date. Algorithm 6 provides an overview of the update process. If an observed vertical structure's coordinate exists within the UHT, the observation is compared to each IHT element in the UHT list. First, we determine that the observation and existing IHT entry are similar when the observed height is within the uncertainty range of the existing entry. If the heights coincide, we determine the horizontal distance between the observation and the prior entry. Instead of updating the entry to the midpoint, we place the new location based on the age of the observation (Figure 4.5). Since the majority of DOF entries are many years old, we make displacement proportional to the relative age of the observation and the prior database entry. The horizontal uncertainty is updated to the observed horizontal uncertainty as long as the observed uncertainty is less than the prior IHT entry.

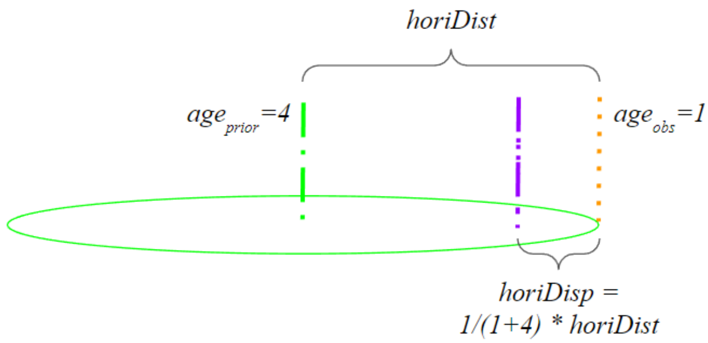

Figure 4.5: Horizontal displacement methodology. After determining that a prior four year old entry (green) and a one year old observation (orange) are similar, the revised entry's horizontal position (purple) is biased towards the more current observation. In this case, the horizontal displacement is 1/5 of the horizontal distance between the prior entry and the observation.

Algorithm 6 Update algorithm. Array of existing database entries represented with *prior* subscript. Array of aspects of newly observed structure denoted by *obs* subscript.

procedure UPDATEEXISTING(IHT, UHT, Obs)
  $LatLonHash \leftarrow SpatialHasher(lat_{obs}, lon_{obs})$
  if $UHT[LatLonHash]$ **then** ▷ if observed hash exists in prior
    for $center \in UHT[LatLonHash]$ **do**
      $prior \leftarrow IHT[center]$
      $Upper_{prior} \leftarrow Ht_{prior} + VertUncert_{prior}$
      $Lower_{prior} \leftarrow Ht_{prior} - VertUncert_{prior}$
      if $obsHt \geq priorLower\ \ \&\ \ obsHt \leq priorUpper$ **then** ▷ similar height
        $horiDist \leftarrow distBetween(lat_{prior}, lon_{prior}, lat_{obs}, lon_{obs})$
        $age_{prior} \leftarrow \Delta date_{prior}$
        $age_{obs} \leftarrow \Delta date_{obs}$
        $age_{weight} \leftarrow \frac{age_{obs}}{age_{obs}+age_{prior}}$
        if $horiUncert_{obs} \leq horiDist$ **then**
          $horiDist \leftarrow horiUncert_{obs}$
        $horiDisp \leftarrow age_{weight} \cdot horiDist$
        $latLon_{update} \leftarrow pointOffset(horiDisp, lat_{prior}, lon_{prior}, lat_{obs}, lon_{obs})$
        if $horiUncert_{obs} \leq horiUncert_{prior}$ **then**
          $horiUncert_{update} \leftarrow horiUncert_{obs}$
        else
          $horiUncert_{update} \leftarrow horiUncert_{prior}$

4.2.2 Power line inference

Using our more accurate and comprehensive vertical structure catalogs, we then implement and evaluate a power line finding algorithm.

We hypothesize that at least three towers are necessary to support power lines and other wires. We also hypothesize that a line of at least three power line pylons would share three characteristics:

1. The height above ground for each tower will be within a certain range
2. The angle between this set of towers will be less than 90 degrees
3. The spacing between successive towers will be within a certain range

We assume a tower detection range of 1,000 m. This conservative detection range coupled with a sensor azimuth of 40°could encompass three transmission towers supporting power lines with over 700 m uniform spacing ($\frac{2{,}128m}{3towers}$) as shown in Figure 4.6. This wide field of consideration is also suited for the low-altitude flight profile where heading and flight path is consistently varied. The approach equally considers obstacles across the entire field of view versus depending on narrow foveal vision. Although this approach would be designed for relatively simple online implementation, it could function just as well in an offline application where power line presence is inferred prior to takeoff.

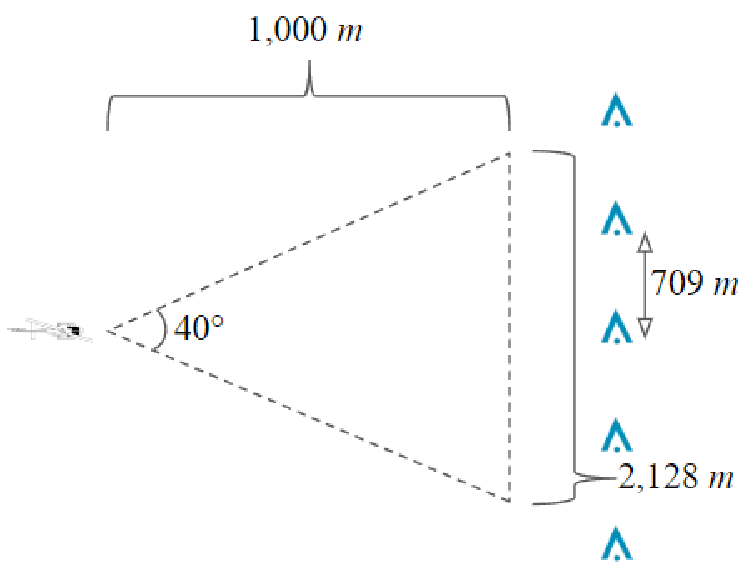

Figure 4.6: Top down view of a notional airborne 3D sensor (not to scale). With a 1,000 m sensor range and 40°azimuth, at least three vertical obstacles (shown in blue) with spacing $\leq 709m$ will be present in the field of view.

Our approach starts by associating towers with a similar height and within a certain distance. Distance between towers is proportional to the tower height. In addition to tower height, the upper limit (z_{upper}) is defined by the uncertainty of the vertical measurement plus a global additional height buffer, α.

Algorithm 7 Tower association algorithm where ϵ represents height uncertainty and PL is a list of sublists of tower spatial hash keys.

$\nu \leftarrow$Height distance ratio
$\alpha \leftarrow$Additional height buffer
$z_{upper} \leftarrow$ Upper height limit
$z_{lower} \leftarrow$ Lower height limit
procedure ListBuilder(obs, IHT)
  for $i \in obs$ **do**
    $obs_{key} \leftarrow SpatialHasher(obs(i))$
    $Dist_{max} \leftarrow obs(i)_z \cdot \nu$
    $z_{max} \leftarrow obs(i)_z + obs(i)_\epsilon + \alpha$
    $z_{min} \leftarrow obs(i)_z - obs(i)_\epsilon - \alpha$
    if $obs(i)_z > z_{upper}$ AND $obs(i)_z < z_{lower}$ **then**
      for $j \in IHT$ **do** ▷ search subset in vicinity
        $Dist \leftarrow distanceBetween(obs(i), IHT(j))$
        if $Dist < Dist_{max}$ **then**
          if $IHT(j)_z \geq z_{min}$ AND $IHT(j)_z \leq z_{max}$ **then**
            $found \leftarrow False$
            for $k \in PL$ **do** ▷ k = sublist
              if $obs_{key} \in k$ **then** ▷ if observed tower exists on sublist
                $PL(k) \leftarrow IHT(j)_{key}$ ▷ add similar tower to sublist
                $found \leftarrow True$ break
              if $IHT(j)_{key} \in k$ **then** ▷ if similar tower is on sublist
                $PL(k) \leftarrow obs_{key}$ ▷ add observed tower to sublist
                $found \leftarrow True$ break
            if $not found$ **then** ▷ create new sublist with both towers
              $PL \leftarrow [obs_{key}, IHT(j)]$

Next, Algorithm 8 cycles through each list of associated towers, L. It starts by sorting the list entries so that they are arranged West to East, then North to South. Then it removes any tower list that contains less than three towers. Next, it checks if the interior angle (shown in Figure 4.7) between each set of three towers in the list has an interior angle greater than or equal to the minimum tower angle, θ_{min}.

Algorithm 8 also includes provisions for a more nuanced angle-based check. The blue text shows a secondary level of more restrictive height difference and distance if the towers are not aligned within a more conservative θ_{align}. This secondary approach is designed to be more permissive for closely aligned towers while maintaining scrutiny on towers that are not in a straight line.

Figure 4.7: Top down view of the interior dashed angle between three towers supporting an orange power line.

4.3 Setting

We use the Federal Aviation Administration's publicly available Digital Obstacle File (DOF) as a baseline. DOF contains multiple descriptive fields for each obstacle entry, including location, height Above Ground Level (AGL), height accuracy, horizontal accuracy and revision date. DOF files are separated by state and issued every 56 days. For our analysis, we chose Delaware. Delaware offers a variety of urban, suburban and rural infrastructure in a compact data set. A histogram of Delaware's March 19, 2023 DOF is shown in Figure 4.8. Table 4.2 shows the breakdown of horizontal accuracy attributed to each DOF entry. Of the 1,063 objects, more than half have a horizontal uncertainty greater than 250 ft with an overall average horizontal accuracy of 206 ft. On average, the last time a DOF entry was added, updated, or verified was 2015.

Algorithm 8 Transmission tower list checking algorithm. PL =list of sublists, L, of tower spatial hash keys. Additional alignment checks in blue.

$n \leftarrow$Minimum number of towers per list
$\theta_{min} \leftarrow$ minimum tower angle
$\theta_{align} \leftarrow$ maximum angle for alignment
$Ht_{max} \leftarrow$ maximum height for less-aligned series
$HDR_{series} \leftarrow$ height distance ratio for less-aligned series
procedure LISTSCRUBBER(PL)
  for $L \in PL$ **do**
    if $length(L) \geq n$ **then**
      $L \leftarrow sorted(L)$ ▷ order entries W-E, N-S
      for $i \in length(L-2)$ **do**
        $T_1 \leftarrow L(i)$
        $T_2 \leftarrow L(i+1)$
        $T_3 \leftarrow L(i+2)$
        $\theta_{series} \leftarrow 180 - angleBetween(T_1, T_2, T_3)$
        $Ht_{series} \leftarrow maxHeightBetween(T_1, T_2, T_3)$
        $Dist_{series} \leftarrow maxDistBetween(T_1, T_2, T_3)$
        $Dist_{max} \leftarrow maxHt(T_1, T_2, T_3) \cdot HDR_{series}$
        if i == 0 **then**
          if $\theta_{series} < \theta_{align}$ **then** continue
          if $\theta_{series} > \theta_{min}$ **then** continue ▷ no continue for alignment checks
            if $Ht_{series} < Ht_{max}$ **then** continue
            if $Dist_{series} < Dist_{max}$ **then** continue
            else$PL.append(L[i+1:])$
          else$PL.append(L[i+1:])$ ▷ add new L without first tower
        else
          if $\theta_{series} < \theta_{align}$ **then** continue
          if $\theta_{series} > \theta_{min}$ **then** continue
            if $Ht_{series} < Ht_{max}$ **then** continue
            if $Dist_{series} < Dist_{max}$ **then** continue
            else$PL.append(L[i:])$
          else$PL.append(L[i:])$

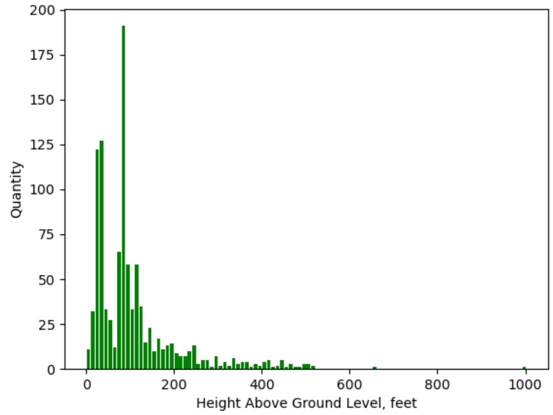

Figure 4.8: Histogram for vertical obstacles contained in Delaware's Digital Obstacle File.

Table 4.2: Distribution of horizontal accuracy for Delaware's Digital Obstacle File.

Horizontal Accuracy ($\pm$ feet)	**Quantity**	**Percent**
20	249	23.4
50	95	8.9
100	2	0.2
250	571	53.7
500	120	11.3
1000	1	0.1
undefined	25	2.4

The National Geospatial-Intelligence Agency has a similar vertical obstacle database, known as the Digital Vertical Obstacle File (DVOF). This database is unclassified but only available to United States government employees. For the state of Delaware, DVOF contains over 15,000 entries with an average horizontal accuracy less than 100 ft. For this analysis, we use Delaware's DVOF as observations against DOF.

4.4 Results

Delaware's DOF data set offers a variety of vertical obstacles. First, we ingest the DOF information into an Index Hash Table (IHT) and assign unique spatial hash values. Next, we build an Uncertainty Hash Table (UHT) which populates spatial hash cells within the radius of the existing DOF coordinate. This evaluation then finds whether a more recent DVOF observation lies within the UHT and proceeds to correlate and update the original DOF entry. We use the improved database

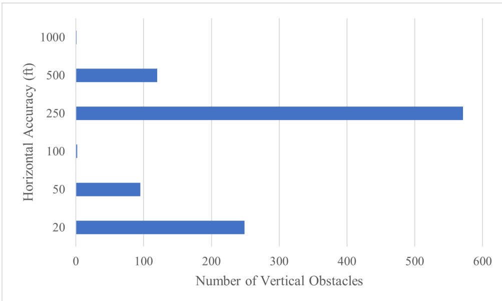

Figure 4.9: Original distribution of Delaware DOF horizontal uncertainty.

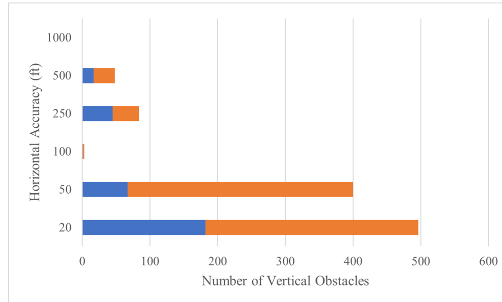

Figure 4.10: Distribution of Delaware horizontal uncertainty after correlation. Blue bars represent entries that were not changed, orange bars represent updated entries.

for power line inference.

4.4.1 Database updating

Figure 4.9 shows that, before implementing DVOF observations with Algorithm 6, the most frequent horizontal uncertainty value was 250 ft. The 1,063 entries had a average horizontal accuracy of 206 ft.

Figure 4.10 shows (in orange) how 719 of the original 1,063 entries were updated. Including the 344 uncorrelated entries (in blue), the overall horizontal uncertainty decreased to 56 ft.

We apply our method to all vertical structure entries in the refined list that results from the previously described database updating approach. This exhaustive comparison provides a sense of effectiveness with a variety of transmission tower arrangements. An entry in the Delaware DOF that was in the Transmission Line Tower category and was also inferred to be a transmission line tower was a true positive (TP). An entry that was correctly categorized as another type was a true negative (TN). False positives (FP) were DOF entries that were erroneously categorized as a transmission tower. False negatives (FN) were entries which were categorized as transmission towers in the DOF database, but were not labelled as a transmission tower with the previously described methods.

4.4.2 Power line inference

4.4.2.1 Rejecting false negatives

The most essential aspect while deducing the presence of power lines is avoiding false negatives. This requires the tower assembled transmission tower lists to include all potential transmission towers. The upper and lower height limit (z_{upper} and z_{lower}) in Algorithm 7 were set to 300 and 49 ft, respectively according to descriptions of recent transmission tower construction projects [147, 148].

Height distance ratio, ν, was evaluated from 1 to 15. The additional height buffer, α, was evaluated from 5 to 200 ft in 5 ft increments.

There were several examples of entries in the Delaware DOF that were categorized as transmission line towers that escaped the consolidation approach in Algorithm 7. Most of these exceptions were due to isolated entries which only show one or two transmission towers of a much longer series (Figure 4.11). Other transmission towers had significant height variations (Figures 4.11 and 4.13). Isolated entries were removed from further analysis since associated towers did not exist in Delaware's DOF.

Figure 4.11: An example of one isolated tower in the Delaware DOF. The 68 ft tall tower in the green rectangle is within 230 ft of the accompanying towers in yellow rectangles that support the dashed red power line.

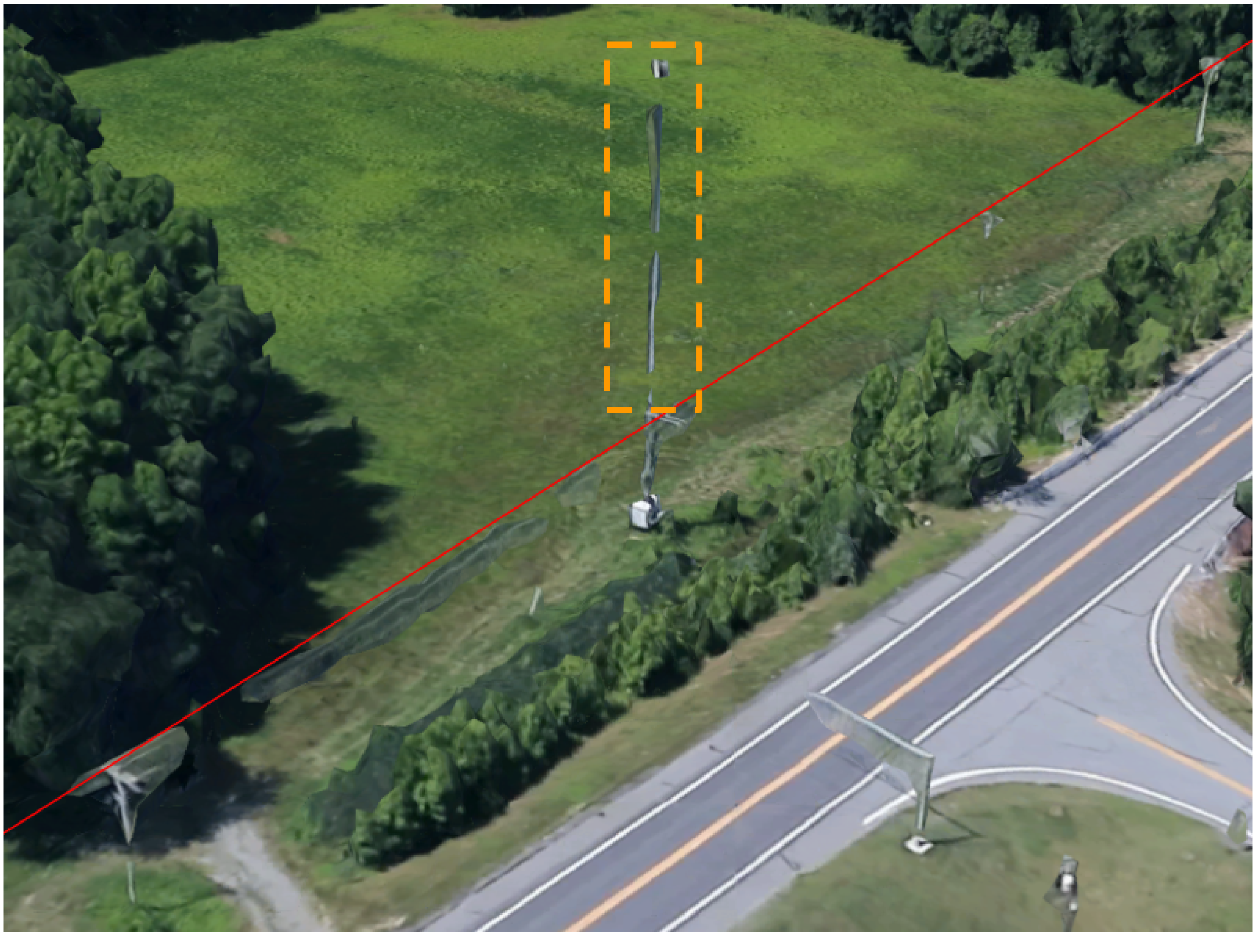

Figure 4.12: A 123 ft cell tower (dashed orange rectangle) co-located with a line of uncharted 30 ft utility poles supporting the red power line.

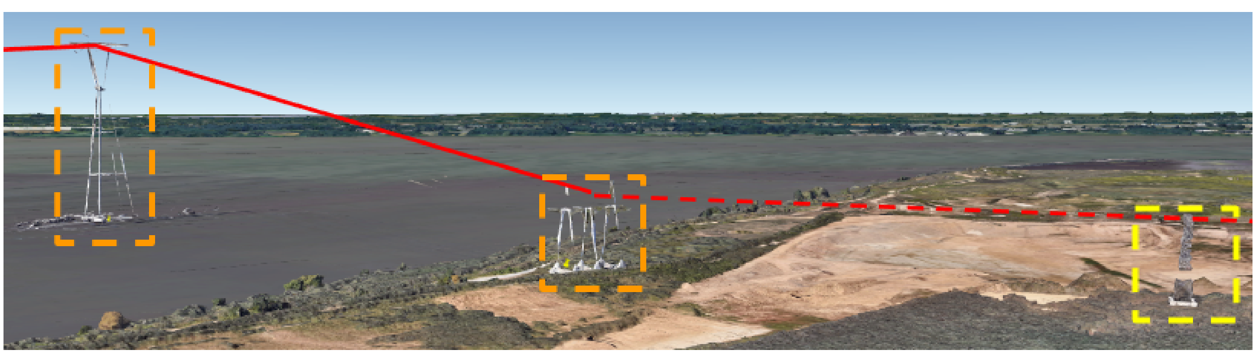

Figure 4.13: A tall 433 ft cell tower in the Delaware river (left dashed orange rectangle) is much higher than the 139 ft tower on the shoreline. The next land-based tower (yellow dashed rectangle) does not exist in Delaware's DOF.

Figure 4.14 shows how false positives increase with larger height buffer and height distance ratio. False negatives show the opposite trend, dropping sharply with a height distance ratio, ν, greater than 6. When $\nu = 85ft$ and $\alpha = 9$, the false negative rate is zero and the false positive rate is 14.9%.

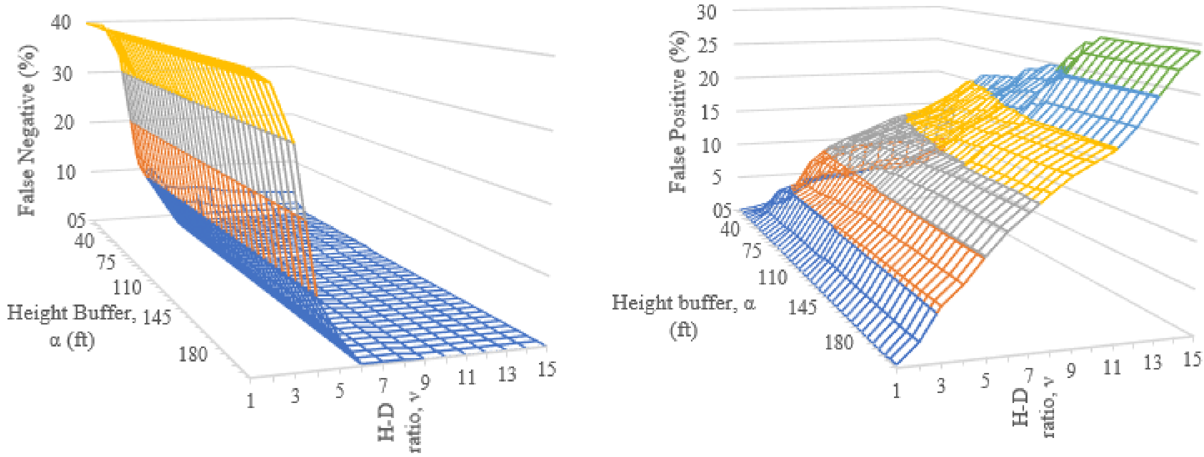

Figure 4.14: Results from evaluating tower association algorithm, false negative (left) and false positive (right).

4.4.2.2 Further reduction of false positives

We hypothesized that false positives would reduce when successive sets of three towers within a list L were closely aligned with a large interior angle. Continuing with $\nu = 85ft$ and $\alpha = 9$, the minimum tower angle, θ_{min} was incrementally evaluated according to the basic (black only) steps in Algorithm 8. Even with a conservative 90 degree angle between towers, the algorithm isolated and incorrectly rejected several entries that were categorized as transmission towers. Figure 4.15 shows an example of a tower (in red circle) that is rejected. If at least two more towers supporting the same dashed orange power line were charted, the span would be depicted as a separate span. Figure 4.16 is an example of a single tower at the beginning of a electrical transmission line. Figure 4.17 shows another example of towers that are not in the Delaware DOF past the shoreline.

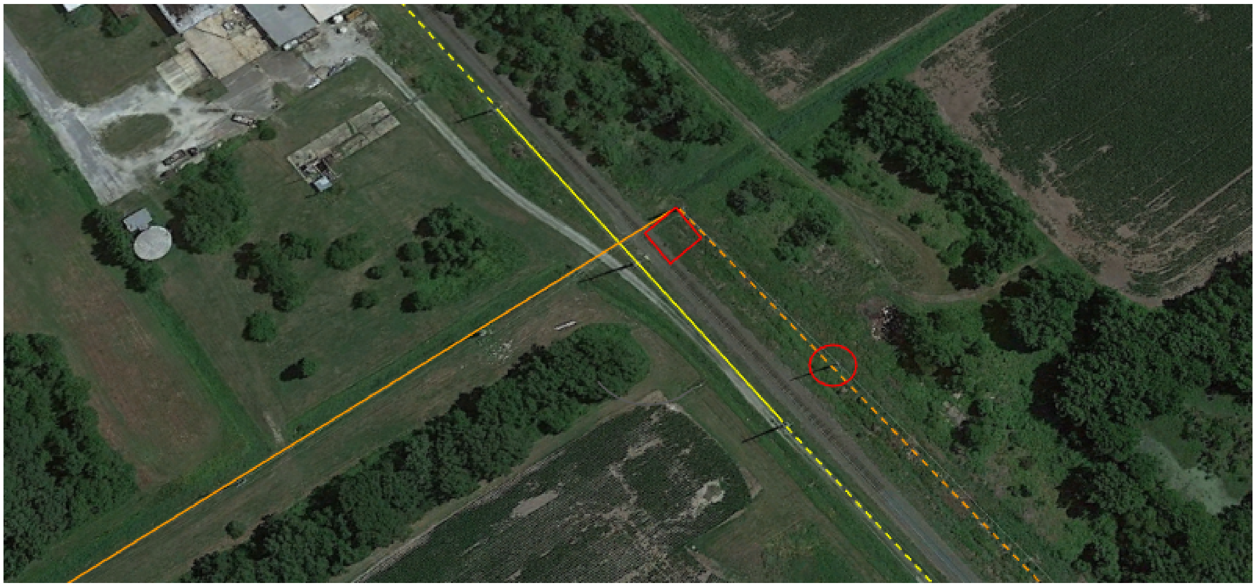

Figure 4.15: A line of charted towers supports an solid orange power line which takes a 90 degree turn at the tower in the red square. The tower in the red circle is rejected due to this sharp angle. Subsequent towers supporting the same power line (now a dashed yellow line) are not charted. Three towers in the Delaware DOF support a power line shown with a solid yellow line; the towers supporting the dashed yellow portion are not charted.

Figure 4.16: A transmission tower (red circle) is not categorized as a transmission tower supporting the dashed blue power line due to the sharp 90 degree turn towards the power plant. The power plant's 189 ft smokestack is an entry in Delaware DOF.

Figure 4.17: Four transmission towers (solid red circles) that support a power line (orange line) across a river are in the Delaware DOF and are correctly recognized by the List Builder algorithm. The angle checking algorithm rejects the assignment due to their rectangular arrangement. Uncharted towers (dashed red circles) that support the terrestrial portion of the power lines (orange dashed lines) are not in the Delaware DOF.

Even by ignoring these isolated entries, false negatives grew significantly while increasing the minimum tower angle, θ_{min}. False positives decreased slowly with increasing θ_{min}.

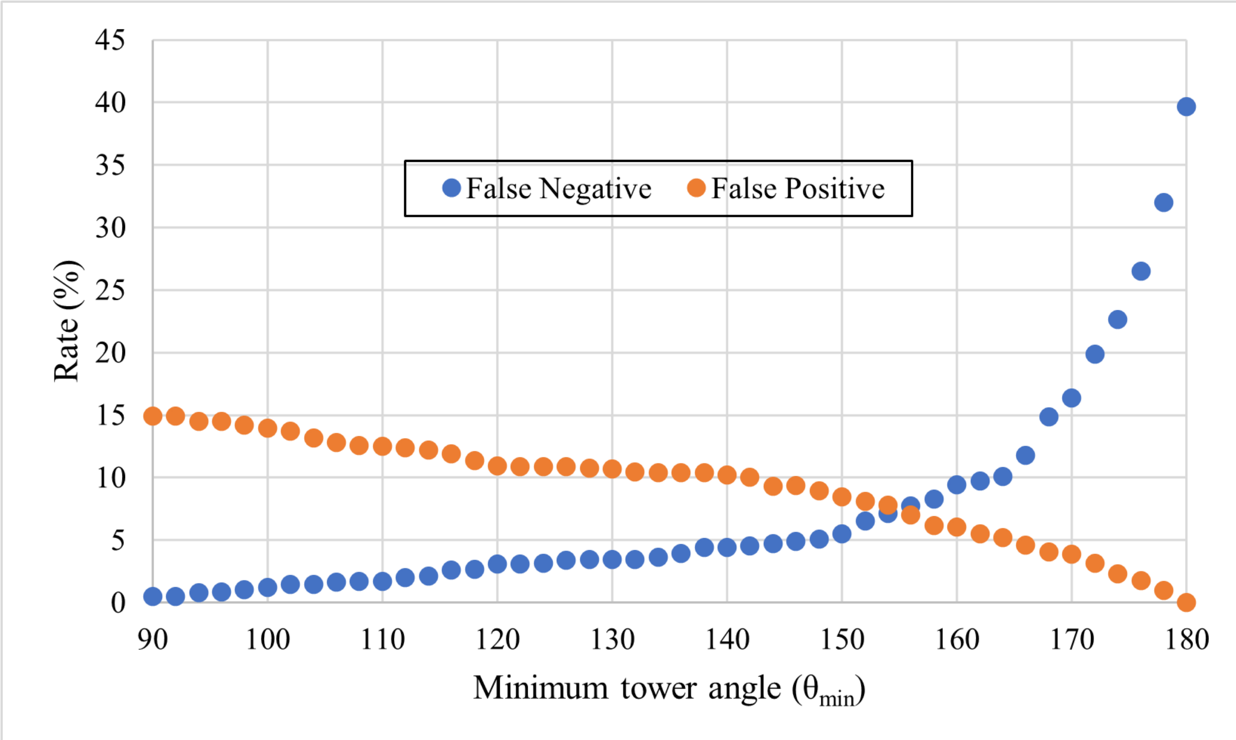

Figure 4.18: Concentration of false positive and false negative tower assignments with varying minimum tower angles, θ_{min}.

This rising number of false negatives led to the implementation of the secondary alignment checks in Algorithm 8. The secondary checks applied to sets of three tower series that had a angle of alignment, θ_{align}, less than than 4 deg. False negative and false positive rates for maximum series height differences (Ht_{max}) and height distance ratios that were more restrictive than the baseline $\nu = 85ft$ and $\alpha = 9$ resulted in Figure 4.19. As with the previous incremental checks, false negative rates increased when a smaller allowable height difference between towers. False negative rates also increases with a smaller height distance ratio, HDR_{series}. False positive rates showed the opposite tendency, increasing with larger distance and height allowances. Ultimately, the lowest false positive rate that maintained zero false negatives was 13%. This occurred with $HDR_{series} = 6$ and $Ht_{max} = 30ft$.

Overall results are presented in Table 4.3. Although adding the alignment filter described in Algorithm 8 reduces False Positives and increases True Negatives, it does not lead to the identification of any additional True Positives.

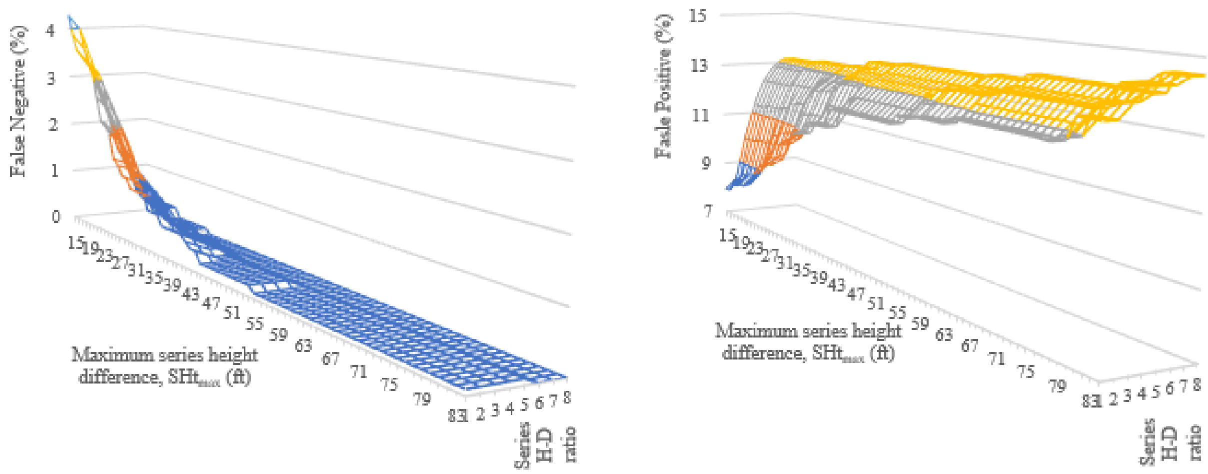

Figure 4.19: False Negative (left) and False Positive (right) rates obtained by incremental checking of maximum series height differences (SHt_{max}) from 15 to 84 ft and height distance ratios from 1 to 8.

Table 4.3: Optimum setting for power line inference. HDR: Height Distance Ratio, AHB: Additional Height Buffer, θ_{align}: minimum angle for alignment, Ht_{max}: maximum height difference among aligned tower series, HDR_{series}: Height Distance Ratio for aligned tower series, TN: True Negative, TP: True Positive, FN: False Negative (excludes previously described examples of isolated transmission towers), FP: False Positive.

HDR, ν	AHB, α (ft)	θ_{align} (deg)	Ht_{max} (ft)	HDR_{series}	TN	TP	FN	FP	FP (%)
9	85	-	-	-	471	408	0	155	14.99
		4	30	6	492	408	0	134	12.96

4.5 Discussion

This chapter explains and evaluates methods for efficiently correlating tower observations with existing, potentially less accurate, database entries. Creating the index and uncertainty hash tables allows the rapid updating of existing entries with a common spatial hash key while significantly increasing average horizontal accuracy. Next, the chapter explains a way to exploit the improved accuracy of this database to determine whether a vertical structure is likely to be an electrical transmission tower. Although the categorization of transmission line towers in the Delaware DOF was mostly correct, isolated towers had to be manually removed to avoid false negatives. Removing these outliers left us with more than 99% of the original dataset.

This evaluation treated entries from a larger vertical structure database (DVOF) as updated observations. DVOF contained over ten times the number of DOF for the state of Delaware. Despite the larger number of entries, more than 30% of DOF entries were not updated. Spot checking showed that most of these DOF entries that were not updated were not contained in DVOF. This indicates that even a large database (such as DVOF) can benefit from being compared to other vertical structure listings.

The methods were evaluated against the 1,063 vertical structures contained in a recent Delaware DOF. Although Delaware's DOF offered a variety of power line structures, the state does not contain a significant amount of terrain relief. Both the database and power line inference methods rely on similar heights above ground level to consolidate duplicate entries and associate sets of transmission towers. Although the height above ground might remain similar for towers in rolling terrain, tower height could also vary to clear hills or other elevated points. Rolling terrain may also cause less uniform tower spacing, which could challenge the proposed power line inference method.

The evaluation of these methods depended on an offline comparison of two databases. The approaches are also designed to incorporate piecemeal tower observations from airborne or ground sensors. The database updating method shows potential for incorporating distributed observations. Airborne platforms that could be gathering tower locations will likely also have significant own ship position uncertainty. The action model that would inform the amount of uncertainty was not included in the experimental results. The Uncertainty Hash Table construct is designed to accommodate large and overlapping horizontal uncertainties, but this ability should be verified in future work.

This chapter used tower height, height uncertainty, position, position uncertainty, and date of last action to compare, contrast, and associate tower structures. In addition to these properties, vertical structure sensors (such as LiDAR or video) also collect gather geometric and radiometric information. Geometric properties, such as width or the presence of cross bars, could be another way to correlate power line transmission structures. Also, sets of transmission towers are typically made of the same material, whether that is steel, wood or concrete. Each of these material types has a unique radiometric reflectance which could be yet another associative criteria.

Future work should examine the effectiveness of the vertical structure and power line finding methods in rolling and steep terrain. Other efforts should also investigate an architecture that accepts dispersed observations, as opposed to comparing offline databases. Raw observations will contain additional details about each structure which are not included in these databases. Incorporating these radiometric and structural properties could further increase the effectiveness of associating towers when inferring power lines or comparing database entries.

4.6 Summary

The database correlation approach addresses critical aspects of distributed mapping for the safety-critical aviation use case. This chapter also proposes a way of overcoming the difficult challenge of detecting power lines by considering the arrangement of associated towers which are much easier to detect. The simple and explainable process can be adapted to larger data sets supplemented with distributed observations. Comparing and correlating prior vertical obstacle databases between the wide variety of manual and automatic tower discoveries shows the potential for a comprehensive repository. This repository can be used for a variety of purposes, including reducing the risk of catastrophic obstacle collisions. The large scale and critical nature of vertical obstacle infrastructure demands an efficient way to manage continuing growth.

CHAPTER 5

Conclusion

Slow or failed detection of low salience vertical obstacles and associated wires is one of today's leading causes of fatal helicopter accidents. The risk of collisions with such obstacles is likely to increase as Advanced Aerial Mobility and broadening drone activity promises to increase the density of air traffic at low altitudes, while growing demand for electricity and communication will expand the number of vertical structures. The current 'see-and-avoid' detection paradigm relies on pilots to spend much of their visual attention looking outside for obstacles. This method is inadequate in low visibility conditions, cluttered environments and given the need for pilots to engage in multiple competing visual tasks. With the expected growing number of hazards and an increased traffic volume, the current approach to collision avoidance will become even less tenable. Although automated obstacle avoidance procedures are on the horizon, the current paradigm that relies on human operators to manually detect and avoid obstacles is likely to persist.

5.1 Intellectual Merit and Broader Impact

Our first goal was to find towers in raw sensor data. This required us to quickly find a minuscule number of returns that composed less than 1% of raw point clouds. The next goal was to gain insight on current challenges and mitigations and investigate the effectiveness of leading visualization and graphic options. We conducted a set of online focus groups. These discussions revealed that vertical structure detection remains a vexing problem. Current mitigation measures can be ineffective or distracting and are typically encumbered by incomplete and inaccurate databases. Our last two goals sought a way to efficiently consolidate vertical structure information, followed by a way of deducing power line presence. These graphic augmentations largely depend on accurate and complete vertical structure databases.

In the process of addressing the first goal, we contributed examples of simulated LiDAR points which portrayed the sparse and cluttered nature of returns from an airborne platform. We bench marked two leading clustering algorithms and compared their effectiveness to our novel overlap algorithm which showed an ability to find vertical structures despite their narrow cross section. For dense point clouds, we showed our mesh filter's ability to distill a small number of relevant points from large, real world point clouds. Next, we explained and evaluated a proportional height filter to deduce the presence of prominent towers. Additional filtering based on sphericity identified remaining vegetation protrusions, increasing the density of points from vertical structures.

Based on shortcomings confirmed in the focus groups, we varied expectancy and salience for a set of helicopter pilots in an immersive flight simulation study. We showed that increasing expectancy with graphic augmentations resulted in faster detection time during low visibility conditions. Altering scene salience by presenting different sensor visualizations did not have a significant effect on detection time in most cases.

In addressing the final set of goals, we contributed a method to efficiently update vertical obstacle databases and find power lines. A pair of hash tables allows for quick correlation of observations with existing entries and significantly increased average horizontal accuracy. In addition to vertical structures, associated power lines also present a collision hazard. Since the low cross section of power lines makes direct detection so difficult, we proposed a method that infers power line presence based on their arrangement. Using the updated database, the method was able to find most electrical transmission towers.

The overall aims of this dissertation were to show a way to quickly filter vertical obstacle information, efficiently correlate these data, and use the precise vertical obstacle information to predict potential wire strikes, while also considering ways to present the information in a palatable way to human operators. Conducting low-altitude flight operations will become even more hazardous without efficient data processing that carefully considers the information requirements of the air vehicle operator.

Judiciously placed graphics can increase awareness of towers, especially in low visibility conditions. Situations with distracted operators (who are likely to spend less time scanning out the window, which is itself a sort of self-imposed reduction in visibility) are also likely to benefit from graphic augmentation. The precise graphics also showed the potential of masking the outside world view while also fixating on the bright features. Other symbology approaches display significantly more details about aircraft state and navigation and must be mindful of this tendency for masking and fixation.

Visual scanning and future automated obstacle avoidance both rely on comprehensive and accurate databases to overcome annoying false positives (which could erode user confidence) or potentially catastrophic false negatives. Both the FAA's DOF and the labelled DALES LiDAR data set had errors which came to light during our analysis. This shows that even the current databases with a limited scope have room for improvement.

5.2 Future directions

The automatic vertical obstacle detection approaches should be evaluated against more data from representative perspectives. Although we showed effectiveness with 10 million point sections of the DALES data set, these tiles were gathered from the top down nadir perspective and had the benefit of greater than 400% coverage. More representative data would be gathered from the sideways perspective, but would have the challenge of many more returns from other objects and false returns. After confirming the ability to find towers in realistic data, the next iteration would be to combine all modules into one architecture to analyze this more representative data that contains even more observations of the same vertical structure. Multiple observations of the same vertical structure could have the potential to overcome normally distributed system navigation error. Repeated observations of vertical structures with decentralized sensors that informs an efficient, centralized database shows the greatest promise for overall information quality.

Aircraft operators, whether they are in the pilot seat or supervising from a distance, will continue to play a significant role in obstacle avoidance. Future human research studies should examine the issue of trust in automation by incorporating occasional augmentations that are not accurate. Augmentations of empty airspace or of non-hazardous objects could quickly become untrustworthy and annoying. Future work could also incorporate a wider variety of circle radii to determine if less horizontal accuracy affects trust, detection rate or detection time. The horizontal uncertainty of obstacles in the real world often overlap. It would be worthwhile to investigate the effect of overlapping circles when vertical structures are closer to one another. Finally, increasing ground speed to impart a faster closure rates or steeper turns to reduce time available to react to obstacles will be more realistic. It is also worth investigating a more engaging workload inject task, such as supervising multiple vehicles simultaneously, that decreases the time available for participants to detect obstacles prior to collision. Although the simulation provided state of the art visualizations and cockpit interfaces, there is no substitute for flight evaluations in an actual aircraft.

5.3 Lists Including the Appendices

```
\showlistofappendices
```

BIBLIOGRAPHY

[1] P. R. Flanigen, J. M. Drockton, and E. M. Atkins, “Airborne detection of vertical obstacles from sparse and cluttered point clouds,” *Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping VII*, vol. 12114, pp. 130–141, 2022.

[2] W. Zhang, J. Qi, P. Wan, H. Wang, D. Xie, and X. Wang, “Cloth simulation filter cloudcompare plugin.” `https://github.com/CloudCompare/CloudCompare/tree/master/plugins/core/Standard/qCSF`, 2022.

[3] R. Schnabel, R. Wahl, and R. Klein, “Efficient ransac for point-cloud shape detection,” in *Computer graphics forum*, vol. 26, pp. 214–226, Wiley Online Library, 2007.

[4] P. R. Flanigen, R. R. Copeland, N. Sarter, and E. M. Atkins, “Current challenges and mitigations for airborne detection of vertical obstacles,” in *Virtual, Augmented, and Mixed Reality (XR) Technology for Multi-Domain Operations III*, vol. 12125, pp. 57–69, SPIE, 2022.

[5] R. G. Fox, “The history of helicopter safety,” in *International helicopter safety symposium*, vol. 690, 2005.

[6] M. Couch and D. Lindell, “Study on rotorcraft safety and survivability,” tech. rep., DEFENSE ACQUISITION UNIV FT BELVOIR VA, 2010.

[7] CTIA - The Wireless Association, “2021 annual survey highlights.” `https://www.ctia.org/news/2021-annual-survey-highlights`, 2021.

[8] CTIA - The Wireless Association, “2022 annual survey highlights.” `https://www.ctia.org/news/2022-annual-survey-highlights`, 2022.

[9] Department of Energy, “Biden-Harris Administration Announces $13 Billion To Modernize And Expand America’s Power Grid.” `https://www.energy.gov/articles/biden-harris-administration-announces-13-billion-modernize-and-expand-americas-power-grid`, 2022.

[10] M. D. Takahashi, B. T. Fujizawa, J. A. Lusardi, M. S. Whalley, C. L. Goerzen, G. J. Schulein, N. L. Mielcarek, M. J. Cleary, J. P. Carr IV, and D. W. Waldman, “Autonomous guidance and flight control on a partial-authority black hawk helicopter,” *Journal of Aerospace Information Systems*, vol. 18, no. 10, pp. 686–701, 2021.

[11] J. Paduano, J. Wissler, G. Drozeski, M. Piedmonte, N. Dadkhah, J. Francis, J. Bold, and F. Langford, "Talos: An unmanned cargo delivery system for rotorcraft landing to unprepared sites," in *American Helicopter Society*, pp. 1–17, 2015.

[12] Defense Advanced Research Projects Agency, "ALIAS equipped Black Hawk helicopter completes first uninhabited flight." `https://www.darpa.mil/news-events/2022-02-08#:~:text=ALIAS%20is%20a%20flexible%2C%20extensible,of%20automation%20into%20existing%20aircraft.`, February 2022. (Accessed: 1 February 2023).

[13] J. M. Wolfe, K. R. Cave, and S. L. Franzel, "Guided search: an alternative to the feature integration model for visual search.," *Journal of Experimental Psychology: Human perception and performance*, vol. 15, no. 3, p. 419, 1989.

[14] C. Wickens, J. McCarley, and K. Steelman-Allen, "Nt-seev: A model of attention capture and noticing on the flight deck," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 53, pp. 769–773, SAGE Publications Sage CA: Los Angeles, CA, 2009.

[15] M. Yeh, J. L. Merlo, C. D. Wickens, and D. L. Brandenburg, "Head up versus head down: The costs of imprecision, unreliability, and visual clutter on cue effectiveness for display signaling," *Human factors*, vol. 45, no. 3, pp. 390–407, 2003.

[16] N. Zimmer, W. Lugsch, D. Ravenscroft, and J. Schiefele, "Aviation obstacle auto-extraction using remote sensing information," in *Proc. SPIE 7110, Remote Sensing for Environmental Monitoring, GIS Applications, and Geology VIII*, p. 71101C, 2008.

[17] M. Dobek, M. Mroz, and A. Ciecko, "Updating of the navigation charts with remote sensing techniques," in *17th International Multidisciplinary Scientific GeoConference*, pp. 829–836, 2017.

[18] E. Özdemir, F. Remondino, and A. Golkar, "An efficient and general framework for aerial point cloud classification in urban scenarios," *Remote Sensing*, vol. 13, no. 10, p. 1985, 2021.

[19] R. Huang, Y. Xu, and U. Stilla, "Granet: Global relation-aware attentional network for semantic segmentation of als point clouds," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 177, pp. 1–20, 2021.

[20] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise.," in *kdd*, vol. 96, pp. 226–231, 1996.

[21] C. Ordóñez, C. Cabo, and E. Sanz-Ablanedo, "Automatic detection and classification of pole-like objects for urban cartography using mobile laser scanning data," vol. 17, no. 7, pp. 1465–, 2017. doi: 10.3390/s17071465.

[22] B. Rodríguez-Cuenca, S. García-Cortés, C. Ordóñez, and M. C. Alonso, "Automatic detection and classification of pole-like objects in urban point cloud data using an anomaly detection algorithm," *Remote Sensing*, vol. 7, no. 10, 2015.

[23] C. Brenner, "Global localization of vehicles using local pole patterns," in *Pattern Recognition: 31st DAGM Symposium, Jena, Germany*, pp. 61–70, 2009.

[24] C. Brenner, "Extraction of features from mobile laser scanning data for future driver assistance systems," in *Advances in GIScience*, pp. 25–42, Springer, 2009.

[25] M. Lehtomäki, A. Jaakkola, J. Hyyppä, A. Kukko, and H. Kaartinen, "Detection of vertical pole-like objects in a road environment using vehicle-based laser scanning data," *Remote Sensing*, vol. 2, no. 3, pp. 641–664, 2010.

[26] Y.-Z. Chen, H.-J. Zhao, and R. Shibasaki, "A mobile system combining laser scanners and cameras for urban spatial objects extraction," in *2007 International Conference on Machine Learning and Cybernetics*, vol. 3, pp. 1729–1733, IEEE, 2007.

[27] M. Caros, A. Just, S. Segui, and J. Vitria, "Object segmentation of cluttered airborne lidar point clouds," *arXiv preprint arXiv:2210.16081*, May 2022.

[28] C. E. Parrish, G. H. Tuell, W. E. Carter, and R. L. Shrestha, "Configuring an airborne laser scanner for detecting airport obstructions," *Photogrammetric Engineering & Remote Sensing*, vol. 71, no. 1, p. 37–46, 2005.

[29] J. Zörner, J. R. Dymond, J. D. Shepherd, S. K. Wiser, and B. Jolly, "Lidar-based regional inventory of tall trees—wellington, new zealand," *Forests*, vol. 9, no. 11, p. 702, 2018.

[30] D.-A. Kwak, W.-K. Lee, J.-H. Lee, G. S. Biging, and P. Gong, "Detection of individual trees and estimation of tree height using lidar data," *Journal of Forest Research*, vol. 12, no. 6, pp. 425–434, 2007.

[31] J.-R. Roussel, D. Auty, N. C. Coops, P. Tompalski, T. R. Goodbody, A. S. Meador, J.-F. Bourdon, F. De Boissieu, and A. Achim, "lidr: An r package for analysis of airborne laser scanning (als) data," *Remote Sensing of Environment*, vol. 251, p. 112061, 2020.

[32] M. Awrangjeb and M. K. Islam, "Classifier-free detection of power line pylons from point cloud data," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 4, p. 81, 2017.

[33] L. C. Thomas and C. D. Wickens, "Eye-tracking and individual differences in off-normal event detection when flying with a synthetic vision system display," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 48, pp. 223–227, Sage Publications Sage CA: Los Angeles, CA, 2004.

[34] F. Viertler and M. Hajek, "Evaluation of visual augmentation methods for rotorcraft pilots in degraded visual environments," *Journal of the American Helicopter Society*, vol. 62, no. 1, pp. 1–11, 2017.

[35] A. C. Warden, C. D. Wickens, D. Mifsud, S. Ourada, B. A. Clegg, and F. R. Ortega, "Visual search in augmented reality: Effect of target cue type and location," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 66, pp. 373–377, SAGE Publications Sage CA: Los Angeles, CA, 2022.

[36] H. Davison and C. Wickens, "Rotorcraft hazard cueing: The effects on attention and trust (technical report arl-99-5/nasa-99-1)," *Savoy, IL: University of Illinois, Aviation Research Lab*, 1999.

[37] M. Yeh, *Attention and trust biases in the design of augmented reality displays*. University of Illinois at Urbana-Champaign, 2000.

[38] Y. Wang, Y. Wu, C. Chen, B. Wu, S. Ma, D. Wang, H. Li, and Z. Yang, "Inattentional blindness in augmented reality head-up display-assisted driving," *International Journal of Human–Computer Interaction*, vol. 38, no. 9, pp. 837–850, 2022.

[39] D. Friesen, C. Borst, M. D. Pavel, O. Stroosma, P. Masarati, and M. Mulder, "Design and evaluation of a constraint-based head-up display for helicopter obstacle avoidance," *Journal of Aerospace Information Systems*, vol. 18, no. 3, pp. 80–101, 2021.

[40] T. Münsterer, B. Singer, M. Zimmermann, and M. Gestwa, "Niag dve flight test results of lidar based dve support systems," in *Degraded Environments: Sensing, Processing, and Display 2018*, vol. 10642, pp. 151–169, SPIE, 2018.

[41] A. Kahana, B. T. Sweet, Z. Szoboszlay, and M. Rottem-Hovev, "Terrain and obstacle avoidance display for low-level helicopter operations in degraded visual environments," in *American Helicopter Society 70th Annual Forum, Montréal, Canada*, 2014.

[42] N. Peinecke, S. Schmerwitz, H.-U. Döhler, and T. Lüken, "Review of conformal displays: more than a highway in the sky," *Optical Engineering*, vol. 56, no. 5, pp. 051406–051406, 2017.

[43] B. Korn, S. Schmerwitz, B. Lorenz, and H.-U. Döhler, "Combining enhanced and synthetic vision for autonomous all-weather approach and landing," *The International Journal of Aviation Psychology*, vol. 19, no. 1, pp. 49–75, 2009.

[44] J. D. Miller, M. Godfroy-Cooper, and Z. P. Szoboszlay, "Augmented-reality multimodal cueing for obstacle awareness: Towards a new topology for threat-level presentation," in *Proceedings of the 75th Vertical Flight Society Annual Forum, Philadelphia, PA*, 2019.

[45] G. Hartnett, J. Hicks, D. Durbin, M. Godfroy-Cooper, J. Miller, K. A. Feltman, A. McAtee, P. St Onge, C. Aura, and J. Stewart, "Pilot cueing for 360 obstacle awareness during dve missions," tech. rep., ARMY AEROMEDICAL RESEARCH LAB FORT RUCKER AL FORT RUCKER United States, 2020.

[46] B. M. Davis, K. Dickerson, and S. C. Gillmore, "User centered design strategies for improving visualization of sensor data in rotorcraft cockpit displays for degraded visual environment operations," in *International Conference on Applied Human Factors and Ergonomics*, pp. 131–141, Springer, 2019.

[47] Federal Aviation Admin., "Advisory circular 70/7460-1l." `https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_70_7460-1L_-_Obstuction_Marking_and_Lighting_-_Change_2.pdf`, Aug. 2018.

[48] CTIA - The Wireless Association, "2021 Annual Survey Highlights." `https://api.ctia.org/wp-content/uploads/2021/07/2021-Annual-Survey-Highlights.pdf`, July 2021. (Accessed: 21 April 2023).

[49] P. Denholm, P. Brown, W. Cole, T. Mai, B. Sergi, M. Brown, P. Jadun, J. Ho, J. Mayernik, C. McMillan, *et al.*, "Examining supply-side options to achieve 100% clean electricity by 2035," tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2022.

[50] Gershenson, Dimitry and Rohrer, Brandon and Lerner, Anna, "A new predictive model for more accurate electrical grid mapping." `https://engineering.fb.com/2019/01/25/connectivity/electrical-grid-mapping/`, January 2019. (Accessed: 7 April 2023).

[51] C. Arderne, C. Zorn, C. Nicolas, and E. Koks, "Predictive mapping of the global power system using open data," *Scientific data*, vol. 7, no. 1, p. 19, 2020.

[52] C. Brenner, "Global localization of vehicles using local pole patterns," in *Pattern Recognition: 31st DAGM Symposium, Jena, Germany, September 9-11, 2009. Proceedings 31*, pp. 61–70, Springer, 2009.

[53] Y. Cai, Z. Lu, H. Wang, L. Chen, and Y. Li, "A lightweight feature map creation method for intelligent vehicle localization in urban road environments," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–15, 2022.

[54] A. Schaefer, D. Büscher, J. Vertens, L. Luft, and W. Burgard, "Long-term urban vehicle localization using pole landmarks extracted from 3-d lidar scans," in *2019 European Conference on Mobile Robots (ECMR)*, pp. 1–7, IEEE, 2019.

[55] H. Dong, X. Chen, and C. Stachniss, "Online range image-based pole extractor for long-term lidar localization in urban environments," in *2021 European Conference on Mobile Robots (ECMR)*, pp. 1–6, IEEE, 2021.

[56] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1689–1696, IEEE, 2020.

[57] G. Zhou, J. Yuan, I.-L. Yen, and F. Bastani, "Robust real-time UAV based power line detection and tracking," in *2016 IEEE International Conference on Image Processing (ICIP)*, pp. 744–748, IEEE, 2016. doi: 10.1109/ICIP.2016.7532456.

[58] A. A. V. Junior and P. S. Cugnasca, "Detecting cables and power lines in small-uas (unmanned aircraft systems) images through deep learning," in *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–7, IEEE, 2021.

[59] A. Stambler, H. Cover, and K. Strabala, "Benchmarking a lidar obstacle perception system for aircraft autonomy," in *Unmanned Systems Technology XX*, vol. 10640, pp. 159–169, SPIE, 2018.

[60] C. Ippolito, K. Krishnakumar, and S. Hening, "Preliminary results of powerline reconstruction from airborne lidar for safe autonomous low-altitude urban operations of small uas," in *2016 IEEE SENSORS*, pp. 1–3, IEEE, 2016.

[61] I. Golightly and D. Jones, "Corner detection and matching for visual tracking during power line inspection," vol. 21, no. 9, pp. 827–840, 2003. doi: 10.1016/S0262-8856(03)00097-0.

[62] C. Zhou, J. Yang, C. Zhao, and G. Hua, "Fast, accurate thin-structure obstacle detection for autonomous mobile robots," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1–10, 2017.

[63] F. Azevedo, A. Dias, J. Almeida, A. Oliveira, A. Ferreira, T. Santos, A. Martins, and E. Silva, "Lidar-based real-time detection and modeling of power lines for unmanned aerial vehicles," *Sensors*, vol. 19, no. 8, p. 1812, 2019.

[64] Epic Games, "Unreal Engine." `https://www.unrealengine.com`, 2021.

[65] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics: Results of the 11th International Conference*, pp. 621–635, Springer, 2018.

[66] G. Van Horn and P. Perona, "The devil is in the tails: Fine-grained classification in the wild," *arXiv preprint arXiv:1709.01450*, 2017.

[67] N. Varney, V. K. Asari, and Q. Graehling, "Dales: A large-scale aerial lidar data set for semantic segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 186–187, 2020.

[68] N. Gellerman, N. Kaabouch, and W. Semke, "A terrain avoidance algorithm based on the requirements of terrain awareness and warning systems," in *2015 IEEE Aerospace Conference*, pp. 1–6, IEEE, 2015. doi: 10.1109/AERO.2015.7119160.

[69] W. Zhang, C. Witharana, W. Li, C. Zhang, X. Li, and J. Parent, "Using deep learning to identify utility poles with crossarms and estimate their locations from google street view images," vol. 18, no. 8, pp. 2484–, 2018. doi: 10.3390/s18082484.

[70] S. Hrabar, "Reactive obstacle avoidance for rotorcraft uavs," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4967–4974, 2011.

[71] K. R. Schulz, S. Scherbarth, and U. Fabry, "Hellas: Obstacle warning system for helicopters," in *Laser Radar Technology and Applications VII*, vol. 4723, pp. 1–8, International Society for Optics and Photonics, 2002.

[72] T. Merz and F. Kendoul, "Beyond visual range obstacle avoidance and infrastructure inspection by an autonomous helicopter," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4953–4960, IEEE, 2011.

[73] S. Scherer, S. Singh, L. Chamberlain, and M. Elgersma, "Flying fast and low among obstacles: Methodology and experiments," in *Intl. J. Robotics Res.*, vol. 27, pp. 549–574, IEEE, 2008.

[74] Y. Jwa and G. Sohn, "A piecewise catenary curve model growing for 3d power line reconstruction," *Photogrammetric Engineering & Remote Sensing*, vol. 78, no. 12, pp. 1227–1240, 2012.

[75] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[76] I. Bogoslavskyi and C. Stachniss, "Efficient online segmentation for sparse 3d laser scans," *J. Photogrammetry, Remote Sensing and Geoinformation Science*, vol. 85, no. 1, pp. 41–52, 2017.

[77] W. Zhang, J. Qi, P. Wan, H. Wang, D. Xie, X. Wang, and G. Yan, "An easy-to-use airborne lidar data filtering method based on cloth simulation," *Remote sensing*, vol. 8, no. 6, p. 501, 2016.

[78] D. Girardeau-Montaut, *Détection de changement sur des données géométriques tridimensionnelles*. PhD thesis, Télécom ParisTech, 2006.

[79] R. Lumia, L. Shapiro, and O. Zuniga, "A new connected components algorithm for virtual memory computers," *Computer Vision, Graphics, and Image Processing*, vol. 22, no. 2, pp. 287–300, 1983.

[80] C. Hill, *Learning Scientific Programming with Python*. Cambridge University Press, 2016.

[81] S. Zolanvari, S. Ruano, A. Rana, A. Cummins, R. E. da Silva, M. Rahbar, and A. Smolic, "Dublincity: Annotated lidar point cloud and its applications," *arXiv preprint arXiv:1909.03613*, 2019.

[82] X. Li, C. Li, Z. Tong, A. Lim, J. Yuan, Y. Wu, J. Tang, and R. Huang, "Campus3d: A photogrammetry point cloud benchmark for hierarchical understanding of outdoor scene," in *Proceedings of the 28th ACM International Conference on Multimedia*, pp. 238–246, 2020.

[83] T. Hackel, N. Savinov, L. Ladicky, J. D. Wegner, K. Schindler, and M. Pollefeys, "Semantic3d. net: A new large-scale point cloud classification benchmark," *arXiv preprint arXiv:1704.03847*, 2017.

[84] Z. Ye, Y. Xu, R. Huang, X. Tong, X. Li, X. Liu, K. Luan, L. Hoegner, and U. Stilla, "Lasdu: A large-scale aerial lidar dataset for semantic labeling in dense urban areas," *ISPRS International Journal of Geo-Information*, vol. 9, no. 7, p. 450, 2020.

[85] M. Kölle, D. Laupheimer, S. Schmohl, N. Haala, F. Rottensteiner, J. D. Wegner, and H. Ledoux, "The hessigheim 3d (h3d) benchmark on semantic segmentation of high-resolution 3d point clouds and textured meshes from uav lidar and multi-view-stereo," *ISPRS Open Journal of Photogrammetry and Remote Sensing*, vol. 1, p. 11, 2021.

[86] Q. Hu, B. Yang, S. Khalid, W. Xiao, N. Trigoni, and A. Markham, "Towards semantic segmentation of urban-scale 3d point clouds: A dataset, benchmarks and challenges," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4977–4987, 2021.

[87] E. B. Fowlkes and C. L. Mallows, "A method for comparing two hierarchical clusterings," *Journal of the American statistical association*, vol. 78, no. 383, pp. 553–569, 1983.

[88] G. Rainieri, F. Fraboni, G. Russo, M. Tul, A. Pingitore, A. Tessari, and L. Pietrantoni, "Visual scanning techniques and mental workload of helicopter pilots during simulated flight," *Aerospace Medicine and Human Performance*, vol. 92, no. 1, pp. 11–19, 2021.

[89] J. Sullivan, J. H. Yang, M. Day, and Q. Kennedy, "Training simulation for helicopter navigation by characterizing visual scan patterns," *Aviation, space, and environmental medicine*, vol. 82, no. 9, pp. 871–878, 2011.

[90] J. H. Yang, Q. Kennedy, J. Sullivan, and R. D. Fricker, "Pilot performance: assessing how scan patterns & navigational assessments vary by flight expertise," *Aviation, space, and environmental medicine*, vol. 84, no. 2, pp. 116–124, 2013.

[91] M. G. Sanders, R. R. Simmons, and M. A. Hofmann, "Visual workload of the copilot/navigator during terrain flight," *Human factors*, vol. 21, no. 3, pp. 369–383, 1979.

[92] D. L. Woods, J. M. Wyma, E. W. Yund, T. J. Herron, and B. Reed, "Factors influencing the latency of simple reaction time," *Frontiers in human neuroscience*, vol. 9, p. 131, 2015.

[93] J. Leach, "Why people 'freeze'in an emergency: temporal and cognitive constraints on survival responses," *Aviation, space, and environmental medicine*, vol. 75, no. 6, pp. 539–542, 2004.

[94] AFS-800, "Advisory Circular 90-48D - Pilots' Role in Collision Avoidance." `https://www.faa.gov/documentlibrary/media/advisory_circular/ac_90-48d_chg_1.pdf`, June 2016. (Accessed: 2 March 2022).

[95] M. R. Endsley, "Measurement of situation awareness in dynamic systems," *Human factors*, vol. 37, no. 1, pp. 65–84, 1995.

[96] N. B. Sarter and D. D. Woods, "Situation awareness: A critical but ill-defined phenomenon," *The International Journal of Aviation Psychology*, vol. 1, no. 1, pp. 45–57, 1991.

[97] R. R. Copeland, "An analysis and classification process towards the qualification of autonomous functions in army aviation," in *Vertical Flight Society's 75th Annual Forum & Technology Display, Philadelphia, Pennsylvania, USA*, 2019.

[98] QSR International Pty Ltd., "Nvivo," 2020.

[99] R. A. Krueger, *Focus groups: A practical guide for applied research*. Sage publications, 2014.

[100] C. D. Wickens and J. S. McCarley, *Applied attention theory*. CRC press, 2019.

[101] N. Moacdieh and N. Sarter, "Display clutter: A review of definitions and measurement techniques," *Human factors*, vol. 57, no. 1, pp. 61–100, 2015.

[102] J. Minor, Z. Morford, and W. Harrington, "Sensor data/cueing continuum for rotorcraft degraded visual environment operations," in *Degraded Environments: Sensing, Processing, and Display 2017*, vol. 10197, pp. 307–323, SPIE, 2017.

[103] R. Feerst and F. Culver, "Utility helicopter operations flying in the wire environment tutorial for power line professionals," in *Proceedings of ESMO'95-1995 IEEE 7th International Conference on Transmission and Distribution Construction, Operation and Live-Line Maintenance*, p. 312, IEEE, 1995.

[104] E. M. Stelzer and C. D. Wickens, "Pilots strategically compensate for display enlargements in surveillance and flight control tasks," *Human factors*, vol. 48, no. 1, pp. 166–181, 2006.

[105] J. Kevin O'Regan, H. Deubel, J. J. Clark, and R. A. Rensink, "Picture changes during blinks: Looking without seeing and seeing without looking," *Visual Cognition*, vol. 7, no. 1-3, pp. 191–211, 2000.

[106] I. E. Hyman Jr, S. M. Boss, B. M. Wise, K. E. McKenzie, and J. M. Caggiano, "Did you see the unicycling clown? inattentional blindness while walking and talking on a cell phone," *Applied Cognitive Psychology*, vol. 24, no. 5, pp. 597–607, 2010.

[107] "Special Airworthiness Information Bulletin, Helicopter Terrain Awareness and Reporting System." Federal Aviation Administration, 18 March 2011 `https://ad.easa.europa.eu/blob/SW1124.pdf/SIB_SW-11-24_1`, 2011. (Accessed: 18 February 2022).

[108] "Advisory Circular 150/5300-18B - General Guidance and Specifications for Submission of Aeronautical Surveys to NGS: Field Data Collection and Geographic Information System (GIS) Standards Document Information." Office of Airport Safety and Standards - Airport Engineering Division, 21 May 2009 `https://www.faa.gov/regulations_policies/advisory_circulars/index.cfm/go/document.information/documentid/74204`, 2009. (Accessed: 14 February 2022).

[109] "Title 14, Chapter 1, Subchapter E, Part 77, Subpart B, Subsection 77.9." United States Code of Federal Regulations 15 February 2022 `https://www.ecfr.gov/current/title-14/chapter-I/subchapter-E/part-77/subpart-B/section-77.9`, 2022. (Accessed: 17 February 2022).

[110] R. Parasuraman and D. H. Manzey, "Complacency and bias in human use of automation: An attentional integration," *Human factors*, vol. 52, no. 3, pp. 381–410, 2010.

[111] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.

[112] T. B. Sheridan and W. L. Verplank, “Human and computer control of undersea teleoperators,” tech. rep., Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab, 1978.

[113] C. Wickens, “Attention: Theory, principles, models and applications,” *International Journal of Human–Computer Interaction*, vol. 37, no. 5, pp. 403–417, 2021.

[114] J. M. Wolfe, “Guided search 2.0 a revised model of visual search,” *Psychonomic bulletin & review*, vol. 1, pp. 202–238, 1994.

[115] C. D. Wickens and A. L. Alexander, “Attentional tunneling and task management in synthetic vision displays,” *The international journal of aviation psychology*, vol. 19, no. 2, pp. 182–199, 2009.

[116] J. M. Wolfe and T. S. Horowitz, “What attributes guide the deployment of visual attention and how do they do it?,” *Nature reviews neuroscience*, vol. 5, no. 6, pp. 495–501, 2004.

[117] J. M. Wolfe and T. S. Horowitz, “Five factors that guide attention in visual search,” *Nature Human Behaviour*, vol. 1, no. 3, p. 0058, 2017.

[118] J. M. Wolfe, “Guided search 6.0: An updated model of visual search,” *Psychonomic Bulletin & Review*, vol. 28, no. 4, pp. 1060–1092, 2021.

[119] C. D. Wickens, B. L. Hooey, B. F. Gore, A. Sebok, and C. S. Koenicke, “Identifying black swans in nextgen: Predicting human performance in off-nominal conditions,” *Human Factors*, vol. 51, no. 5, pp. 638–651, 2009.

[120] J. Duncan and G. W. Humphreys, “Visual search and stimulus similarity.,” *Psychological review*, vol. 96, no. 3, p. 433, 1989.

[121] National Oceanic and Atmospheric Administration, “Marginal Visual Flight Rules.” `https://forecast.weather.gov/glossary.php?word=MVFR#:~:text=Marginal%20Visual%20Flight%20Rules%20%2D%20in,3%20to%205%20miles%20visibility.`, 2023. (Accessed: 31 March 2023).

[122] FAA National Headquarters Publications and Administration, “FAA Order JO 7110.65Z - Air Traffic Control.” `https://www.faa.gov/air_traffic/publications/atpubs/atc_html/chap7_section_5.html`, 2022. (Accessed: 31 March 2023).

[123] T. X. B. Nguyen, K. Rosser, and J. Chahl, “A review of modern thermal imaging sensor technology and applications for autonomous aerial navigation,” *Journal of Imaging*, vol. 7, no. 10, p. 217, 2021.

[124] M. Kassner, W. Patera, and A. Bulling, “Pupil: An open source platform for pervasive eye tracking and mobile gaze-based interaction,” in *Adjunct Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp ’14 Adjunct, (New York, NY, USA), pp. 1151–1160, ACM, 2014.

[125] J. H. Goldberg and A. M. Wichansky, “Eye tracking in usability evaluation: A practitioner’s guide,” in *the Mind’s Eye*, pp. 493–516, Elsevier, 2003.

[126] P. Blignaut, "Fixation identification: The optimum threshold for a dispersion algorithm," *Attention, Perception, & Psychophysics*, vol. 71, pp. 881–895, 2009.

[127] RStudio Team, *RStudio: Integrated Development Environment for R*. RStudio, PBC., Boston, MA, 2020.

[128] D. Bates, "lme4: Linear mixed-effects models using s4 classes," *http://CRAN. R-project. org/package= lme4*, 2010.

[129] A. Kuznetsova, P. B. Brockhoff, and R. H. Christensen, "lmertest package: tests in linear mixed effects models," *Journal of statistical software*, vol. 82, pp. 1–26, 2017.

[130] V. T. Nagaraj and I. Chopra, "Safety study of wire strike devices installed on civil and military helicopters," 2008.

[131] C. D. Wickens and S. R. Dixon, "The benefits of imperfect diagnostic automation: A synthesis of the literature," *Theoretical Issues in Ergonomics Science*, vol. 8, no. 3, pp. 201–212, 2007.

[132] A. D. Andre and C. D. Wickens, "When users want what's not best for them," *Ergonomics in design*, vol. 3, no. 4, pp. 10–14, 1995.

[133] J. Irwin III, B. Brown, A. Rich, M. Schwerke, and E. Kocher, "Avms h-47 chinook with apas tactile cueing demonstration results," in *72th Annual Forum of the American Helicopter Society*, 2016.

[134] J. Driver and C. Spence, "Cross–modal links in spatial attention," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 353, no. 1373, pp. 1319–1331, 1998.

[135] A. M. Kelley, B. Cheung, B. D. Lawson, E. Rath, J. Chiasson, J. G. Ramiccio, and A. H. Rupert, "Tactile cues for orienting pilots during hover over moving targets," *Aviation, Space, and Environmental Medicine*, vol. 84, no. 12, pp. 1255–1261, 2013.

[136] M. Osmon, B. Fujizawa, Z. Szoboszlay, B. Davis, B. Thomas, , Z. Morford, A. Mouser, and K. Kuck, "RDECOM Degraded Visual Environment Mitigation (DVE–M) Science and Technology Program 2016 Flight Trials," Tech. Rep. AD1035322, Aviation Development Directorate, 2017.

[137] P. Flanigen, B. Fujizawa, and B. Davis, "RDECOM Degraded Visual Environment Mitigation (DVE-M) Program 2017 Flight Trials," Tech. Rep. AD1044917, Aviation Development Directorate, 2017.

[138] Air Transportation Division's 135 Flight Operations Section, "Terrain Awareness and Warning Systems (TAWS) Nuisance Alerts." `https://www.faa.gov/other_visit/aviation_industry/airline_operators/airline_safety/info/all_infos/info23003`, March 2023. (Accessed: 21 April 2023).

[139] Schwab, George, "Helicopter Terrain Awareness and Warning System Special Airworthiness Information Bulletin." `https://ad.easa.europa.eu/blob/SW1124.pdf/SIB_SW-11-24_1`, March 2011. (Accessed: 21 April 2023).

[140] R. Chandrasekaran, A. P. Payan, K. B. Collins, and D. N. Mavris, "Helicopter wire strike protection and prevention devices: Review, challenges, and recommendations," *Aerospace Science and Technology*, vol. 98, p. 105665, 2020.

[141] Aspen Environmental Group, "Transmission Structures Fact Sheet." https://ia.cpuc.ca.gov/environment/info/aspen/cltp/archive/Files_8_26_14/_4TransmissionStructuresFactSheet.pdf, 1995.

[142] J. Rivera, J. Leimhofer, and H.-A. Jacobsen, "Opengridmap: towards automatic power grid simulation model generation from crowdsourced data," *Computer Science-Research and Development*, vol. 32, pp. 13–23, 2017.

[143] Federal Aviation Admin., "Advisory circular 150/5300-17c." https://www.https://www.faa.gov/documentLibrary/media/Advisory_Circular/150-5300-17C-Chg1.pdf, Sept. 2017.

[144] "GeoPackage Encoding Standard." https://www.ogc.org/standard/geopackage/, November 2021. (Accessed: 7 April 2023).

[145] Federal Aviation Admin., "Order 8260.19i." https://www.faa.gov/documentLibrary/media/Order/Order_8260.19I.pdf, June 2020.

[146] L. Wiesmann, R. Marcuzzi, C. Stachniss, and J. Behley, "Retriever: Point cloud retrieval in compressed 3d maps," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 10925–10932, IEEE, 2022.

[147] Electrical Volt, "Electrical Transmission Tower: Types, Design and Parts." https://www.electricalvolt.com/2022/11/electrical-transmission-tower-types-design-parts/, November 2022. (Accessed: 21 April 2023).

[148] Western Renewables Link, "Fact sheet: transmission towers and conductors." https://www.westernrenewableslink.com.au/assets/resources/Transmission-towers-and-conductors-May-2022.pdf, May 2022. (Accessed: 21 April 2023).

APPENDIX A

Aircrew Mission Briefing

You are the pilot not on the controls (P, right seat) for a set of flight routes in a UH-60M helicopter flight simulator. The pilot on the controls (P*, left seat) will be responsible for managing the coupled aircraft. The P* will offer assistance as briefed during the first baseline route, but may be distracted during the subsequent routes. Ensure you are comfortable monitoring/adjusting the items in **bold** on the aircrew briefing. The P* will confirm your ability to complete these items in **bold** prior to the 6 evaluation routes. The standard aircrew mission briefing has additional *relevant details*:

1. Mission Overview: Safely transport personnel along several assigned flight routes in a congested urban environment
    a. Flight routes:
        i. *One "baseline" route*
        ii. *Six subsequent 5-7 minute evaluation routes*
        iii. *Each coupled to 200' AGL, 80 knots*
    b. Weather: *forecast weather will vary between routes and range from 700-1000' ceilings, 2-3 miles visibility*
    c. Highest risk associated with mission: *Simulated collision with other aircraft, terrain or obstacles due to imperfect planning data, low altitude, limited visibility*
2. Required items, mission equipment and personnel
    a. *Number of passengers will be specified before each route*
    b. *Out the window view will have three simulated visualizations:*
        i. *Unaided: naked eye*
        ii. *Image intensification: green night vision goggle (NVG) simulation*
        iii. *Thermal imaging: gray infrared (IR) simulation*
        iv. *Some structures will have graphic augmentation (yellow circle or bounding box)-you will encounter examples of each on the first "baseline" route*
3. Analysis of the aircraft: not applicable
4. Crew actions, duties and responsibilities
    a. Aircrew coordination-two challenge rule, standard terminology
    b. Transfer of flight controls: *flight will be coupled, either pilot may do momentary cyclic input if necessary but avoid large inputs*
    c. Airspace surveillance procedures:
        i. *Scan sectors: Left seat: 8 to 2 o'clock; Right seat: 10 to 4 o'clock*

Figure A.1: First page of Aircrew Mission Briefing

ii. *Immediately inform crewmembers of all air traffic or obstacles that pose a potential threat. Use clock-altitude-distance format ("tower, 12 o'clock, same altitude, 1 mile").* *Announce when seeing an obstacle, not just yellow box or circle graphical augmentation*

iii. Announce when focused inside and outside the aircraft

d. Brief emergency actions

(1) P* will initiate any immediate actions on flight controls, P will take immediate actions off flight controls and bring out checklist

5. General crew duties: Announce when focused inside
   a. Pilot on the controls (P*)
      (1) Fly the aircraft: primary focus outside, announce when inside
      (2) Cross-check systems and instruments
   b. Pilot not on the controls (P)
      (1) **Tune radios** and **adjust transponder**
      (2) Navigate
      (3) Copy clearances, ATIS and other information
      (4) Cross-check systems and instruments (**heading, barometric altitude, fuel**)
      (5) **Monitor/transmit on radios**: Callsign Army 0474
      (6) Read and complete checklist items as required
      (7) Set/adjust switches and systems as required (**altimeter setting**)
6. Time and place for crew-level after action review: *Co-pilot will take notes throughout. P will complete a post flight questionnaire after all 7 routes are complete.*
7. Crewmembers' questions, comments and acknowledgement of mission briefing

Your eye gaze position will be recorded along with any conversations on the intercom or radio(s). Prior to starting the baseline route in the simulator, we will calibrate the eye tracker by shifting your gaze between five dots on a computer screen. Ensure that the eye tracker fits comfortably and securely and that you can clearly hear the Air Traffic Control audio prior to starting the evaluation routes.

Figure A.2: Second page of Aircrew Mission Briefing

APPENDIX B

Post-flight Survey

Visualization Preferences

What is your participant ID?

How many unique towers do you estimate that you encountered (i.e. came within approximately two rotor discs of the aircraft)?

Figure B.1: First page of Post-flight Survey

For the Day-Unaided airspace surveillance task...

	Very low	.	.	.	.	.	Very High
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○

Figure B.2: Second page of Post-flight Survey

	Very low	.	.	.	.	.	Very High
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Day-Unaided airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

Figure B.3: Third page of Post-flight Survey

For the Day-Thermal/IR airspace surveillance task...

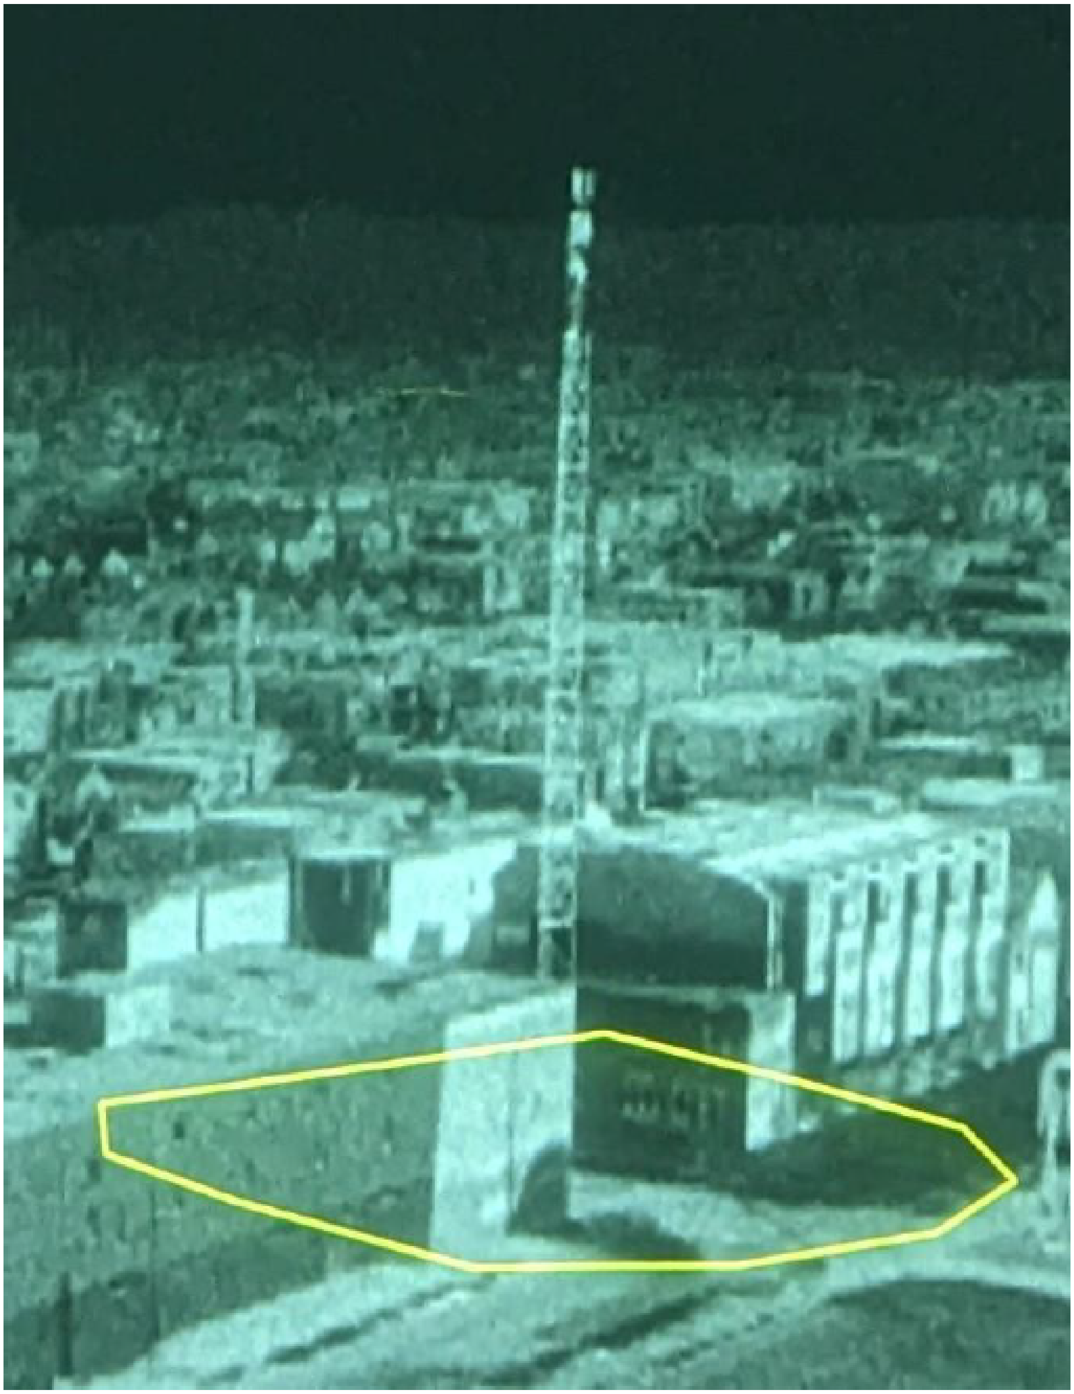

Figure B.4: Fourth page of Post-flight Survey

	Very low	.	.	.	.	.	Very high
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Day-Thermal/IR airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

Figure B.5: Fifth page of Post-flight Survey

For the Dusk-Thermal/IR airspace surveillance task...

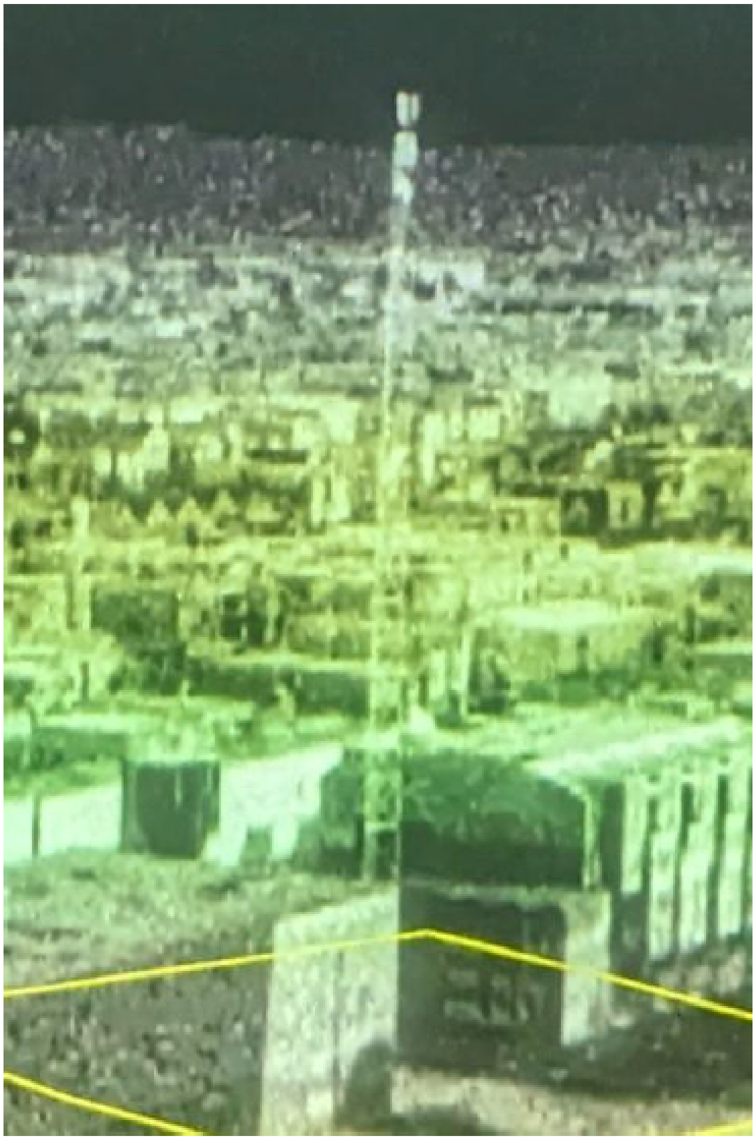

	Very low	.	.	.	.	.	Very high
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○

Figure B.6: Sixth page of Post-flight Survey

	Very low	.	.	.	.	.	Very high
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Dusk-Thermal/IR airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

Figure B.7: Seventh page of Post-flight Survey

For the Dusk-Unaided airspace surveillance task...

	Very low	.	.	.	.	.	Very high
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○

Figure B.8: Eighth page of Post-flight Survey

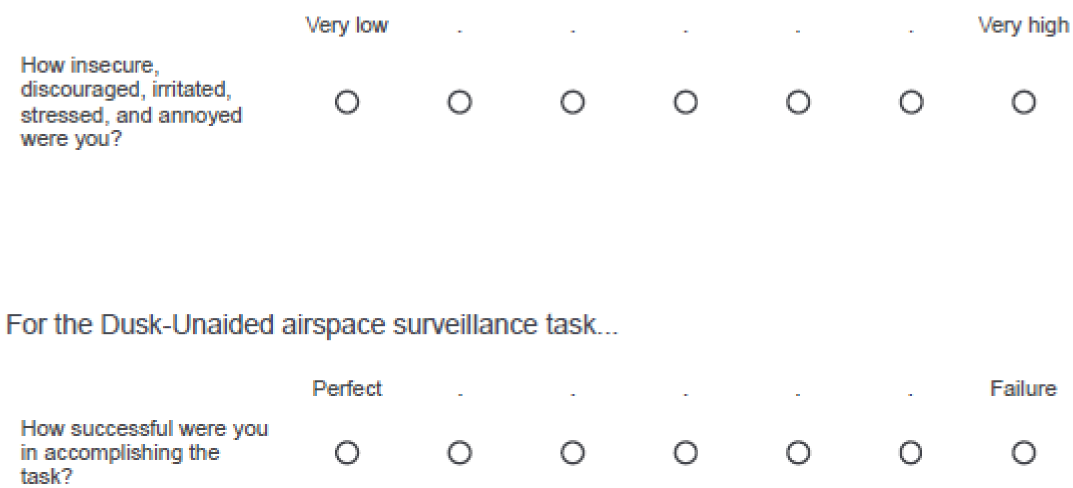

	Very low	.	.	.	.	.	Very high
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Dusk-Unaided airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

Figure B.9: Ninth page of Post-flight Survey

For the Night-NVG airspace surveillance task...

	Very low	.	.	.	.	.	Very high
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○

Figure B.10: Tenth page of Post-flight Survey

	Very low	.	.	.	.	.	Very high
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Night-NVG airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

Figure B.11: Eleventh page of Post-flight Survey

For the Night-Thermal/IR airspace surveillance task...

	Very low	.	.	.	.	.	Very high
How mentally demanding was the task?	○	○	○	○	○	○	○
How physically demanding was the task?	○	○	○	○	○	○	○
How hurried or rushed was the task?	○	○	○	○	○	○	○

Figure B.12: Twelfth page of Post-flight Survey

	Very low	.	.	.	.	.	Very high
How hard did you have to work to accomplish your level of performance?	○	○	○	○	○	○	○
How insecure, discouraged, irritated, stressed, and annoyed were you?	○	○	○	○	○	○	○

For the Night-Thermal/IR airspace surveillance task...

	Perfect	.	.	.	.	.	Failure
How successful were you in accomplishing the task?	○	○	○	○	○	○	○

What was your external awareness (proximity to terrain, buildings, etc.) during...

	1 (low)	2	3 (moderate)	4	5 (perfect)
Day-Unaided	○	○	○	○	○
Day-IR	○	○	○	○	○
Dusk-Unaided	○	○	○	○	○
Dusk-IR	○	○	○	○	○
Night-NVG	○	○	○	○	○
Night-IR	○	○	○	○	○

Which visualization mode provided most awareness of the external flight environment in Daytime conditions?

Unaided Thermal / IR

0 10 20 30 40 50 60 70 80 90 100

0 = unaided, 50 = same, 100 = IR

Figure B.13: Thirteenth page of Post-flight Survey

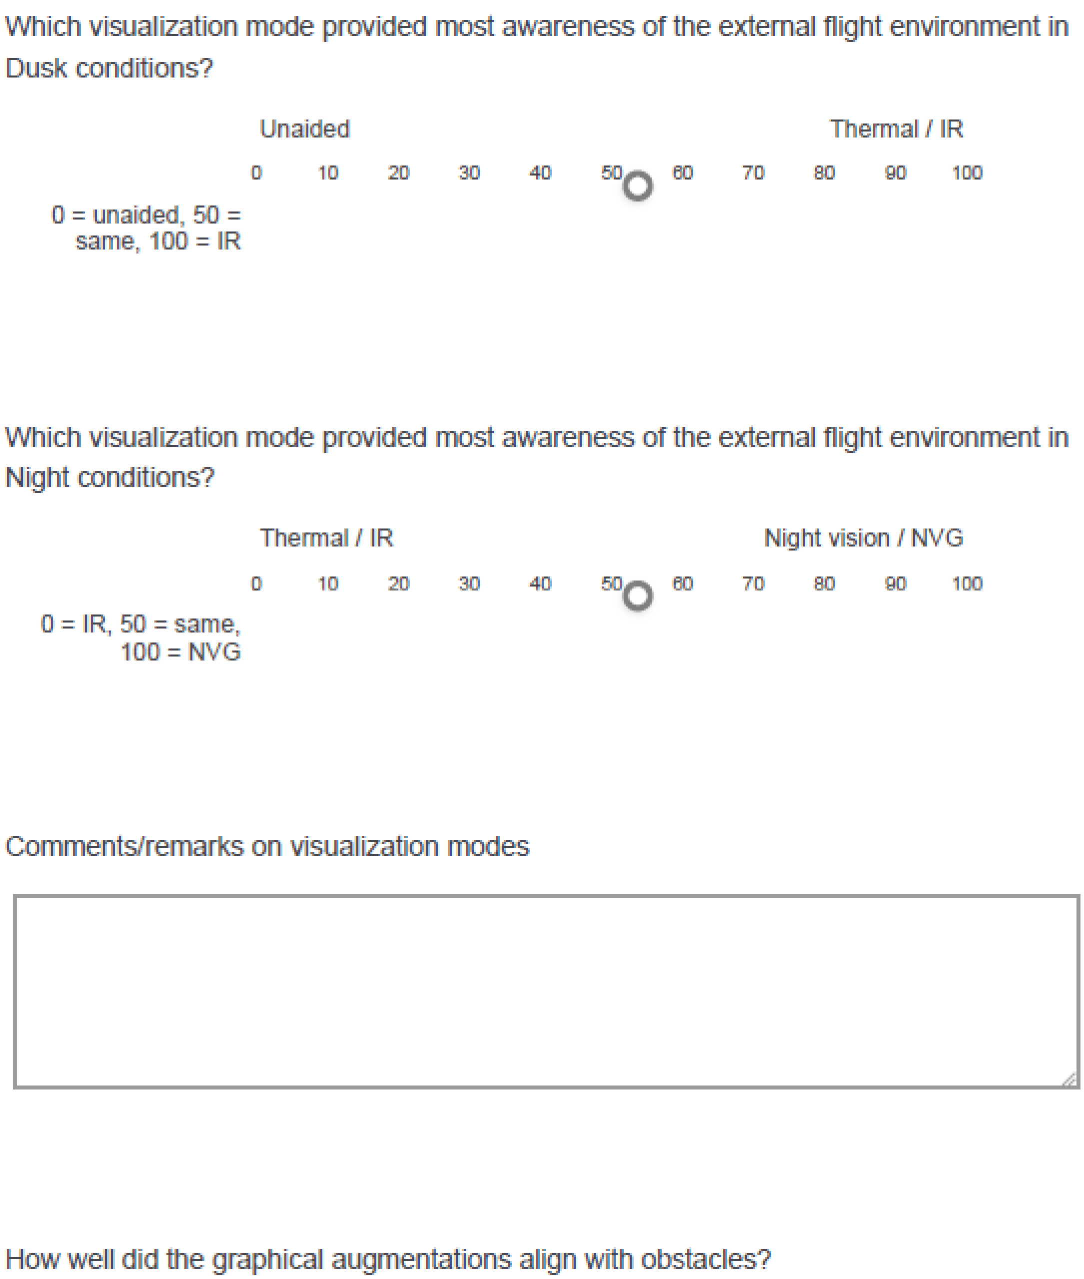
Which visualization mode provided most awareness of the external flight environment in Dusk conditions?

	Unaided							Thermal / IR			
	0	10	20	30	40	50	60	70	80	90	100
0 = unaided, 50 = same, 100 = IR											

Which visualization mode provided most awareness of the external flight environment in Night conditions?

	Thermal / IR							Night vision / NVG			
	0	10	20	30	40	50	60	70	80	90	100
0 = IR, 50 = same, 100 = NVG											

Comments/remarks on visualization modes

How well did the graphical augmentations align with obstacles?

Figure B.14: Fourteenth page of Post-flight Survey

	1 (barely)	2	3 (moderate)	4	5 (perfect)
Yellow circles	○	○	○	○	○
Red boxes	○	○	○	○	○

How did graphical augmentations affect your ability to locate vertical obstacles in flight?

What graphical augmentation was most useful for finding obstacles?

Yellow circles — Yellow boxes

0 10 20 30 40 50 60 70 80 90 100

0 = yellow circles, 50 = equal, 100 = red boxes

Mental Capacity: How much mental capacity did you have to spare...

	1 (plenty of capacity)	2	3 (some capacity)	4	5 (completely saturated)
Average across all routes	○	○	○	○	○
When responding to ATC	○	○	○	○	○
When identifying an obstacle	○	○	○	○	○

Figure B.15: Fifteenth page of Post-flight Survey

Information Quality: How good is the information portrayed by:

	1 (insufficient)	2	3	4	5 (very useful)
Unaided, no sensor	○	○	○	○	○
Gray IR	○	○	○	○	○
Green NVG	○	○	○	○	○
Yellow bounding box	○	○	○	○	○
Yellow circle	○	○	○	○	○

Comments on Situation Awareness

Figure B.16: Sixteenth page of Post-flight Survey