A Deep Learning Approach Using Social Media Data to Estimate Ground Risk of UAS in Urban Areas

Partnered with the Georgia Institute of Technology and Pattent, LLC

Primary Faculty Advisor: Dr. Dimitri Mavris
Other Faculty Advisors: Dr. Michael Balchanos
Team Members: Jeffrey Pattison-Aerospace Engineering PhD Candidate

June 2023

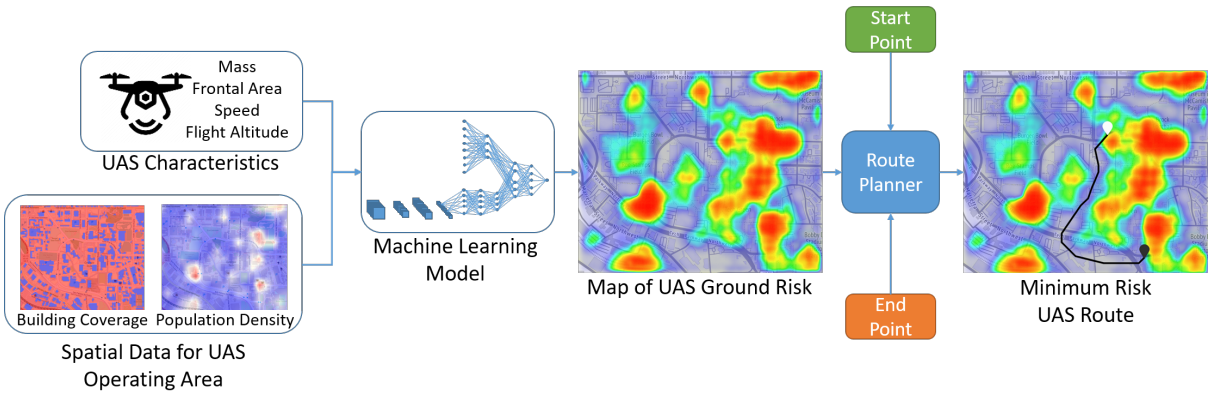

Contents

1 A Deep Learning Approach Using Social Media Data to Estimate Ground Risk of UAS in Urban Areas: Executive Summary

The use of Unmanned Aerial Systems (UAS) is growing in a variety of industries, including for inspections, law enforcement, and logistics. However, governmental regulatory agencies like the FAA still impose strict regulations around UAS to ensure safe operations. These regulations include payload and weight limits, restrictions on operating in airspace near airports, pilot requirements, and prohibitions on operating over people or beyond the visual line of sight. Although waivers can be obtained to bypass some of these restrictions, there is still a push from industry for the FAA to repeal some of the more restrictive regulations, allowing for more autonomous UAS operations. However, as UAS autonomy increases and the role of human pilots diminishes, more safety measures are required to ensure the UAS can still operate safely without pilots. These safety measures can include hardware or software changes, or can even include a revision in the UAS concept of operations and mission planning.

Current aviation practices requires risk assessment to determine the likelihood and severity of accidents. If the likelihood of an accident is too high, or if the consequences of an incident are too severe, then the risk level is deemed to be unacceptable. Due to the long history of manned aircraft, historical flight data can be used to estimate failure rates and safety metrics to determine probability and severity of failure. Consequently, this historical flight data can be leveraged to assess risk and determine if an aircraft is operating within an acceptable target level of safety. Unlike manned aircraft, UAS do not have decades of historical flight data to use to assess risk levels. As a result, researchers have turned to modeling and simulation methods in an effort to quantify the risk of UAS by approximating the expected fatality rates, measured in fatalities per flight hour. Common methods found in literature for estimating the ground risk of a UAS failure at a given location involve approximating the failure rate, determining the probable impact locations using the governing laws of physics given a failure event, and then approximating the probability of a UAS causing a fatality given the impact location. The probability of causing a fatality based on impact is dependent on the kinetic energy of the vehicle, the population density of the impact area, and the presence of possible shelter, like buildings and trees, that may provide some protection. This approach is probabilistic in nature due to uncertainties in the initial UAS state when estimating the impact locations, requiring several descent trajectories to be simulated. As a result, using this physics-based approach to approximate ground risk for a single location can be time consuming, which is why a Machine Learning approach is explored in this work relying on data created from the physics-based method.

Surrogate modeling is common in engineering as a way to approximate computationally expensive or time consuming processes. In recent years, Machine Learning methods have proven to be quite good surrogate models to replace computationally expensive models provided enough data for training. The focus of work in this paper is a novel approach for a new risk assessment method that quantifies UAS ground risk using Machine Learning, relying on social media activity data to supplement historical data on population density estimates. Using social media data as a way to estimate population density can be useful for identifying dynamic areas that could be high risk and unaccounted for using static historical data. With the physics-based approach, risk data can be generated with various UAS configurations at different locations, and this data can be used to train a Machine Learning model that takes UAS characteristics and the spatial distributions of population density and building coverage as inputs. The Machine Learning approach proposed allows for quick ground risk assessment to quantify the ground risk for given UAS flight conditions. Using this ground risk assessment, a UAS pilot can use the route planning tool proposed with the Machine Learning model to find a suitable route that has an acceptable level of risk based on the risk map created with the Machine Learning model.

2 Problem Statement and Background

The work proposed for the 2023 FAA Data Challenge is focused on assisting the rapidly evolving new and novel uses of the national airspace in relation to the introduction of small Unmanned Aerial Systems (UAS). Leveraging modern approaches of modeling and simulation to quantify the ground risk of UAS, machine learning can be used to enhance the risk assessment by decreasing the time required to estimate risk. This time reduction allows UAS pilots to quickly evaluate risk levels for a given UAS operation, or it can even be used to find the optimal operation that minimizes UAS ground risk.

Public interest in UAS has grown in the past few decades as they are being adopted to complete an increasing number of tasks, ranging from surveillance to mapping. Safety concerns around the use of UAS have slowed the adoption of the UAS into the National Airspace (NAS). Mr. N. A. Sabatini, the former Associate Administrator for the Aviation Safety before the House aviation Subcommittee, said in a talk, "There is a missing link in terms of technology today that prevents these aircraft from getting unrestricted access to the NAS" [1]. As a result, the process of UAS adoption has been incremental, where rules and regulations are slowly relaxed or repealed. Some of the more restrictive regulations forbid operating Beyond Visual Line of Sight (BVLOS) and largely forbid operating over groups of people. While obtaining waivers to operate BVLOS are possible, the ability to operate over people is still heavily regulated due to safety concerns. Hence, these requirements significantly hinder UAS use in urban areas that are dynamic and highly populated.

Urban areas can see large benefits from the capabilities UAS have to offer. This includes decreasing response time from first responders to improving logistics for shipping food and goods, or even delivering medical supplies in record times. However, the heavily populated areas have made UAS use very difficult with current regulations. Currently, the FAA says the ability to fly over people varies depending on the risk levels a UAS presents to those below, so a revision of current regulations would be required in order to exploit the full benefits of UAS. This change cannot occur until there is sufficient reason to believe the use of UAS in these areas is safe. Therefore, realistic and detailed risk analysis for UAS is required to ensure any UAS operation over people does not exceed some acceptable level of risk.

With safety being the primary goal of the FAA, risk evaluation and mitigation has long been required for manned aircraft. Due to the long history of manned aircraft, the risk assessment can be done using historical aviation accident and incident data [2]. Because the introduction of UAS has been more recent compared to manned aircraft, there is an insufficient collection of data related to operations like flight hours, number of accidents and incidents, and failure rates [3]. Consequently, the use of models and simulations has proven to be the next best option for UAS risk assessment. While the risk assessment for manned aircraft includes the risk to those onboard, UAS do not pose similar threats, so the primary risk for UAS is to those in other aircrafts and to those on the ground. This work will focus on the latter and explore ground risk.

There are several factors that can impact the ground risk levels of UAS, including the surrounding environment and the characteristics of the UAS and its trajectory. Several approaches can be found in literature outlining how to approximate risk for a UAS. From [1], the risk for a UAS is quantified as the expected number of fatalities per flight hour, and this metric can be determined using the kinetic energy of the UAS on impact, the probability of hitting a person, and the availability of shelter that might absorb some of the kinetic energy. The common ground risk assessment using this approach is to take a predefined UAS path and simulate different failure events along the path. For each failure event, a descent trajectory can be approximated using the governing laws of physics to estimate the most probable impact locations. For example, in the event of a power failure, a ballistic trajectory can be approximated to predict where the UAS will land. Finding the probable impact locations is required because there is some uncertainty in the UAS initial position, speed, and aerodynamic characteristics,

so several descent trajectories need to be simulated. From the impact location, the kinetic energy can be determined. Combining kinetic energy with the population density, the expected fatality rate can be calculated [1]. As mentioned, this process is repeated along various points of the UAS trajectory. Because the risk assessment is completed for a predefined path, this is of little service if the safety levels for the given path exceed the maximum level of safety. Hence, more information is required in order to assist in the route planning phase to ensure the path of the UAS does not exceed maximum safety levels.

In [4], the authors build on the ground risk assessment approach previously described, but use the approach to compile a risk map. This risk map identifies the high risk areas and is used to assist in the path planning of a UAS. Using the risk map with a path planning algorithm, the UAS can find an optimal path that minimizes the ground risk. The authors accomplish this by discretizing a flight zone and assessing the ground risk at every discretized location for a given cruising speed and UAS model. Using the high-fidelity probabilistic ground risk approach previously described at every location in the risk map can become quite time consuming and computationally expensive. This is problematic in urban areas where the population density is incredibly dynamic and the ground risk map needs to be updated frequently. To improve upon [4], this work offers a novel approach for ground risk assessment for unmanned rotor vehicles using a Machine Learning model in place of the physics-based model to enable rapid ground risk assessment. This risk assessment can then be used by the FAA and others to quantify just how dangerous a UAS operation over people is and enables better decision making when planning a UAS mission. To compare the physics-based approach with the Machine Learning approach, both methods will be used to create a risk map for the campus at the Georgia Institute of Technology to see how closely the Machine Learning method accurately estimates the physics-based model.

With a rapid way of generating a UAS risk map, UAS users can not only determine if a predetermined UAS route exceeds a maximum level of allowable risk, but it can also be used to find a route that minimizes risk. By minimizing the ground risk, an additional layer of safety is added to UAS operations in urban areas, which is required if the role of the pilot is to be reduced. Reducing the need for a pilot benefits all who are looking to use UAS by decreasing the cost of operation. However, the safety measure alone is not the only benefit, since there is a component of public trust as well to using UAS in urban areas. By demonstrating the UAS will be avoiding heavily populated areas, this may build trust in the use of UAS.

3 Solution Definition

3.1 The Physics-Based Model

As mentioned in the previous section, there is a need for risk evaluation for UAS operations to safely operate over people in urban areas, but with UAS being introduced into the airspace relatively recently, there is insufficient flight data for UAS to make similar risk assessments as for manned aircraft. For this reason, modeling and simulation has become the best way to estimate UAS risk. The modeling and simulation method used in this work to approximate UAS risk is a physics-based approach based on the state of the art methods outlined in [1] [4] [5] [6]. This physics-based approach will then be used to generate training data so a Machine Learning algorithm can approximate the UAS risk in a more time efficient manner.

In [1], the ground risk metric is defined as the expected rate of fatalities, with an acceptable level of risk being 10^{-7} fatalities per flight hour based on equivalent levels of safety seen in manned aircraft. The approach in [4] also uses this metric for ground risk assessment. The equation for the expected rate of fatalities can be seen below for a given location [4].

$$f_F = A_{exp} * D_p * P(fatality|exposure) * f_{GIA} \quad (1)$$

In the above equation, f_F is the expected rate of casualties, A_{exp} is the area exposed during the crash, D_p is the population density for the area of the crash, $P_{(fatality|exposure)}$ is the probability of a fatality given the exposure, and f_{GIA} is the rate of ground impact accidents.

3.1.1 Area Exposed During a Crash

The term A_{exp} from Equation 1 is the area exposed to a crash for a single person on the ground. From [4], this area can be found with Equation 2.

$$A_{exp} = \pi(r_p + r_{uas})^2 sin(\gamma) + 2(r_p + r_{uas})(h_p + r_{uas})cos(\gamma) \quad (2)$$

As seen above, r_p is the radius of the average person, r_{uas} is the radius of the UAS, γ is the glide angle, and h_p is the height of the average person. With this, the area exposed during a crash can be determined.

3.1.2 Population Density Estimation

The population density of an area plays a crucial role in the ground risk for a UAS. Highly populated areas will result in a higher probability of a UAS striking a person in the event of an unplanned and uncontrolled descent. Obtaining accurate estimates on the population density information is critical for enabling UAS ground risk assessment, otherwise the risk assessment is meaningless. Multiple methods for estimating population density have been mentioned in literature.

In [4], city census data is used to estimate the population density. Using census data is adequate for demonstrating proof of concept in an academic setting and is easy to obtain, but this data is static and may not be indicative to how humans move throughout the day. Additionally, census data may be stale and outdated by the time it is available. Another method mentioned in [4] relies on the use of mobile phone data. Mobile phone data may be a good resource for accurate estimates throughout the day, but obtaining this data is difficult. Accessing mobile phone data for a given area may take days to obtain depending on the size of the area of interest. In [14], the LandScan Global Population Database [15] is used from the Oak Ridge National Laboratory. The LandScan database is open-source and provides high resolution (about 90m x 90m) averages of population density throughout the day and contains population density averages for daytime as well as night time. Because the LandScan database offers averages at different times of the day, it was the chosen databased for this work to assist in population density estimates. The below images show a heatmap of the population density distribution at the Georgia Institute of Technology for daytime and night time created using the 2021 LandScan database. The units for the heatmap are people per square meter with each cell being approximately 10m x 10m.

(a) Population Density Heatmap During Daytime

(b) Population Density Heatmap During Night Time

Figure 1: Population Density Distribution from 2021 LandScan Database

The LandScan database has one of the same disadvantages as the census data in that it is relatively static. Because LandScan is historical data, it fails to account for anomalous events that might result in a change in the normal expected population density. In urban areas, these anomalous events include festivals, concerts, or parades. As a result, any risk assessment solution used in urban areas will need to account for these occurrences that LandScan cannot in order to obtain a more complete picture on the population density as it changes. This work proposes to supplement the use of historical data with a new method for monitoring population density by using social media activity.

The popular social media site Snapchat has accessible information that shows how many posts have been made within a given radius for any given GPS location. Using this method, one can request this information for any location of interest to generate a heatmap of the Snapchat activity. This can become useful to gain insight on events that might draw large numbers of people, which would not be accounted for in the historical data for population densities. To test the validity of this, the social media activity was observed while the Dogwood Festival occurred at Piedmont Park in Atlanta, Georgia from April 15,2023 to April 17, 2023. The two images below show the social media activity at Piedmont Park during different times of the day while the festival took place.

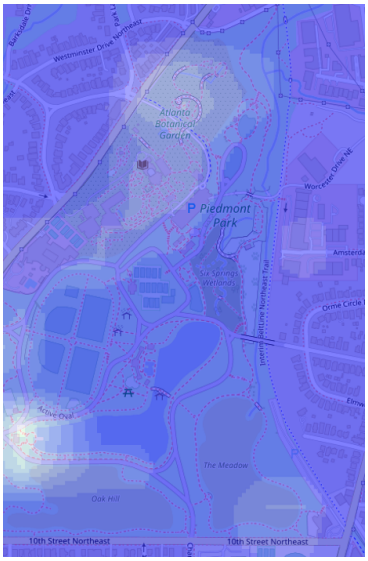

(a) Social Media Activity on 4/15 Around 10:40am

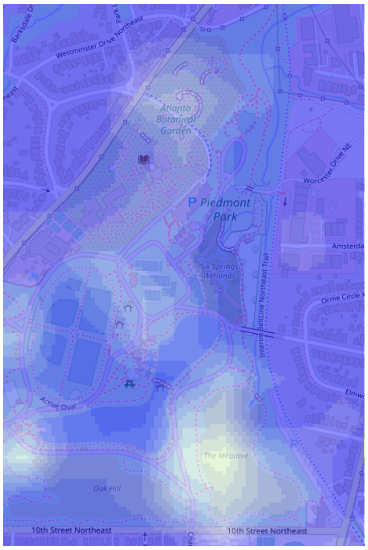

(b) Social Media Activity on 4/15 Around 8:30pm

Figure 2: Social Media Activity Generated During the Dogwood Festival

From the images above, it can be seen that there was more social media activity as the day progressed, suggesting there was a growing population. To show this change in population was abnormal, the social media data was compared with the LandScan data for the same location. This comparison is shown in the images below.

(a) LandScan Population Density Estimates

(b) Social Media Activity Population Density Estimates

Figure 3: A Comparison Between LandScan Database and the Social Media Activity

The discrepancy between the LandScan population density and the social media data collected confirms the historical data does not properly account for anomalous events like the Dogwood Festival. However, the social media data collected cannot be used by itself because not everyone uses social media, nor is it used everywhere. That is why the two combined can be used to paint a better picture of population densities by taking the maximum value between the two resources. This way the large events that draw crowds and generate social media activity can be taken into consideration while the LandScan database can provide average estimates at every other location that does not generate much social media activity. As better methods for estimating population densities are derived, they can be used in place of the methods proposed in this paper.

3.1.3 Probability of Fatality Given Exposure

The probability of fatality given exposure is the probability of a UAS strike resulting in death if it were to impact a person. One approach outlined in [1] is to map the kinetic energy of the UAS on impact to the probability of resulting in a fatality given the UAS impacts a person. This model takes into account not only the kinetic energy, but it also includes a shelter factor that can provide some protection to people on the ground. This shelter factor may be different depending on if the shelter provided is a building, tree, or if there is no shelter at all. The more protection a shelter may provide, the higher this shelter factor, and the less likely a UAS accident is to result in a fatality. No shelter would have a shelter factor of 0 while a building would have a shelter factor of 5 [4]. The equation for finding probability of fatality given exposure, $P_(fatality|exposure)$, can be found below with Equation 3 from [1].

$$P(fatality|exposure) = \frac{1}{1+\sqrt{\frac{\alpha}{\beta}}[\frac{\beta}{E_{imp}}]^{\frac{1}{4ps}}} \tag{3}$$

In the above equation, E_{imp} is the kinetic energy at impact and p_s is the sheltering factor to take into account surrounding structures that may absorb some of the energy. The α parameter is the impact energy required for a fatality probability of 50% with a sheltering factor of 0.5 while β is the impact energy required for a fatality as the sheltering factor goes to 0. According to [4], acceptable values for α and β are 100 kJ and 34 J, respectively. Also from [4], common values for the sheltering factor for different types of shelter include 0 for no shelter, 2.5 for sparse trees, and 5 for low buildings. The open-source database OpenStreetMap contains information on the location of buildings that can be used for finding the shelter factor [16]. The image below shows the building coverage layout for the Georgia Institute of Technology, with blue representing the buildings. Every location that did not have any building coverage was assumed to provide no shelter. Each cell in the image below is approximately 10m x 10m.

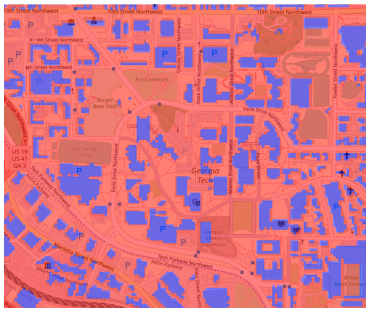

Figure 4: Recorded Building Layouts for Georgia Institute of Technology

The final component of Equation 1 is f_{GIA}, which is the rate at which ground impacts occur. This value is measured as number of occurrences per hour, and ideally would be based on flight history data. In [1], the value is estimated to be between 10^{-6} to 10^{-9} and is based on the average accident rate involving unmanned aircraft. However, this failure rate is dependent on each vehicle specifically, and is subject to change as vehicles become safer. For this work, a constant value of 10^{-6} incidents per flight hour is used as the conservative estimate within the range from [1]. This value can be updated and changed as better estimates are collected on the true value of the probability of UAS failure.

The above approach outlined the different components of Equation 1 to quantify the ground risk as expected fatality rates for a UAS in a given initial location. With the sheltering factor and population density playing important roles in the ground risk calculation, it is important to know where a UAS is going to land. If a UAS experiences a failure event at a given location, its descent trajectory could land the vehicle at a location far from where the incident occurred depending the UAS altitude and speed. As a result, estimating the probable impact locations of the UAS is required for proper ground risk assessment to identify the quantities of its impact location, like shelter and population density. There are countless ways that a UAS could fail resulting in an unplanned or uncontrolled descent. Some of these failures include a power outage, a loss of one or more propellers for multi rotor vehicles, or a loss of control from a pilot. Each of these different failure types result in different descent trajectories. According to [5], the ground risk of a UAS is dominated by a ballistic descent type that might occur if power is lost when compared to other descent trajectories from other failure types, so for this work only a ballistic descent is modeled to find the probable impact locations. The governing laws of physics can be used to find the probable impact locations of a UAS that loses power at some location with some velocity, altitude, and physical characteristics. The governing equations to find the impact locations are shown below.

$$m\ddot{x} = \frac{-1}{2}\rho|\dot{x}|\dot{x}C_dA \tag{4}$$

$$m\ddot{y} = \frac{-1}{2}\rho|\dot{y}|\dot{y}C_dA \tag{5}$$

$$m\ddot{z} = \frac{-1}{2}\rho|\dot{z}|\dot{z}C_dA - mg \tag{6}$$

In the governing equations, m is the UAS mass, $\ddot{x}$, $\ddot{y}$, $\ddot{z}$ are the acceleration of the UAS in the global frame with z being the altitude, ρ is the air density, C_d is the UAS drag coefficient, and A is the frontal area of the UAS. For some given initial location x_0, y_0, z_0 with speeds $\dot{x_0}$, $\dot{y_0}$, $\dot{z_0}$, the governing equations can be solved for when the final altitude z_f is zero to find the values of x_f and y_f, the UAS impact location. However, there is likely going to be some uncertainty in the initial position, velocity, and drag coefficient [7]. In order to account for the uncertainties of the initial conditions, several descent trajectories are required to find the most likely impact locations. Therefore, for a given descent trajectory i, the initial conditions x_{0i}, y_{0i}, z_{0i}, $\dot{x_{0i}}$, $\dot{y_{0i}}$, $\dot{z_{0i}}$, and C_{di} are pulled from a normal distribution with a given mean. For example, if the probable impact locations are needed for a UAS flying with a recorded initial speed of 5 m/s at some given location, then the velocities used for simulating the probable impact locations would be taken from a normal distribution centered around 5 m/s. The table below illustrates the normal distributions used for the position, velocity, and drag coefficient, taken from [7].

Parameter	Distribution
x_{0i}	N(x_0,0.5)
y_{0i}	N(y_0,0.5)
z_{0i}	N(z_0,0.5)
$\dot{x_{0i}}$	N($\dot{x_0}$,2.0)
$\dot{y_{0i}}$	N($\dot{y_0}$,2.0)
$\dot{z_{0i}}$	N($\dot{z_0}$,2.0)
C_{di}	N(C_d,0.2)

Table 1: UAS Condition Parameters

To fully account for the uncertainty, 500 descent trajectories were simulated to find the probable impact locations for a UAS traveling with some given initial conditions. Because the objective of the risk assessment is to create a risk map to assist in route planning, no route for the UAS has been determined yet, so the direction the UAS travels has not been specified. To account for this, each of the 500 different descent trajectories is given a different heading, with the heading determined by sampling from a uniform distribution ranging between 0 and 360 degrees. The image below shows the results of 500 simulated descent trajectories to find the most probable impact locations for a UAS with some initial conditions. In the image, it is assumed the UAS initially starts at (0,0) with each blue marker identifying one of the probable impact locations.

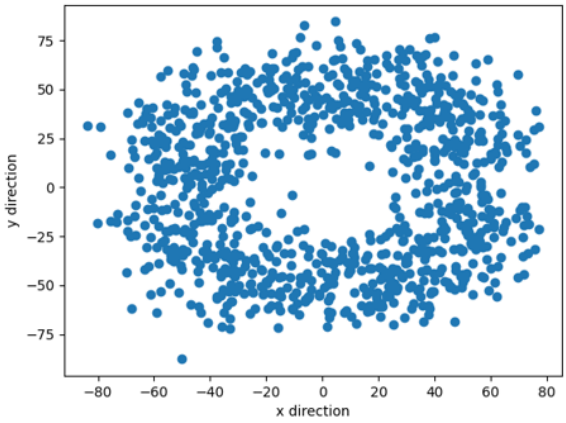

Figure 5: 500 Probable Impact Locations for a UAS

For each of the 500 trajectories simulated at a given location, the kinetic energy of the UAS, population density and sheltering factor were recorded at the location of impact and used to find the expected fatality rate with the equations above. After calculating 500 expected fatality rates for a given UAS starting location, the mean fatality rate of the 500 trajectories was recorded as the fatality rate for that location with the given initial conditions. To create a risk map used to assist in route planning, this process needs to be repeated at various locations. For every location within the map, probable impact locations and expected fatality rates need to be calculated. Because of this, the process to generate a risk map can become quite time consuming as the size of the map increases. The research in [4] and [7] attempt to make simplifying assumptions for determining the probable impact location to reduce the computation time with the trade-off of losing some accuracy. However, adding simplifying assumptions can only go so far before the credibility of the analysis is lost. For this reason, machine learning methods are explored to estimate the risk in a more time efficient manner compared to the high fidelity physics-based model. As a result, the physics-based model can be made to be as high fidelity as possible with little concern for the computation time because the machine learning model can be used to approximate it within some reasonable degree of accuracy in a fraction of the time required.

4 Methodology

4.1 The Machine Learning Approach

From predictive text and language processing to computer vision and autonomous vehicles, Machine Learning has become quite popular with the rise of Artificial Intelligence (AI). Machine Learning algorithms have shown tremendous capability in learning to recognize patterns in data for problems that are complex for traditional approaches [8]. For this reason, this work explores how well Machine Learning methods would be able to learn and estimate the ground risk of UAS based on the data collected from the high fidelity physics-based model.

There are many different types of algorithms that fit into the Machine Learning category, with some being better suited for certain applications over others. For the application of estimating UAS ground risk, the objective is to feed the UAS initial condition parameters and spatial data around a GPS location into the Machine Learning model as inputs, which would then output an expected fatality rate for that given location. This process could be repeated at every location in an area of interest to generate a risk map. One way to do this would be to generate data using the physics-based model to create

a database of fatality rates mapped to the input conditions of the UAS. This is a supervised learning approach where the physics model generates the output data, and the Machine Learning model can learn the pattern between the input parameters and the output fatality rate. Because the fatality rates are continuous, the Machine Learning algorithm will need to be used for regression rather than classification. Therefore, a supervised learning algorithm for regression narrows down the list of suitable Machine Learning algorithms to use.

The type of input data will also affect which type of Machine Learning algorithm is suitable. Based on the description of the physics-based model, the type of input data required is mixed between numeric and spatial data. Numeric data would include the mass, speed, frontal area, and altitude of the UAS. All of these components affect the kinetic energy upon impact, and therefore affect the expected fatality rate. Beyond the parameters just mentioned, the sheltering factor from the surrounding coverage and the population density also play a role in the fatality rate as well. These values are scalar, but based on the probable impact locations seen in Figure 5, it is difficult to pinpoint a single value to use since the population density and building coverage will change based on where the UAS lands. For this reason, the entire area encapsulating the probable impact locations is required as input to accurately estimate the expected fatality rates. Therefore, the Machine Learning model used needs to be able to account for both scalar values like the UAS characteristics and initial conditions and also the spatial inputs like shelter and population density. While some Machine Learning algorithms can handle either numeric data or spatial data, there is no one Machine Learning algorithm well suited for both. The approach proposed in this paper combines two popular Machine Learning algorithms to account for the different input data types. Based on the work in [9], it can be shown that a Multilayer Perceptron (MLP) and a Convolutional Neural Network (CNN) can be combined to account for numeric as well as spatial data. The MLP accounts for the numeric data while the CNN accounts for the spatial data. The outputs of each of these models can then be combined with additional layers to produce a single output, predicting the expected fatality rate. A simple summary of the model can be seen in the image below.

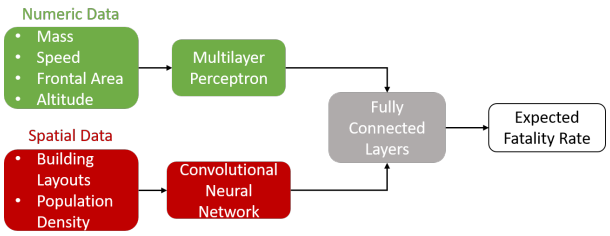

Figure 6: Flowchart for the Machine Learning Model

4.1.1 The Multilayer Perceptron

The Multilayer Perceptron is one of the simplest Artificial Neural Network (ANN) architectures and is comprised of one or mode hidden layers between an input layer and an output layer [8]. Each layer consists of a number of neurons, and each neuron in a layer takes as input all the values of the neurons in the layer before it and then transforms the inputs using an activation function. These values are then passed to all the neurons in the next layer. For traditional regression applications using only a MLP, the final layer is called the output layer and consists of a single neuron, which is the predicted value. The diagram below illustrates a MLP with one input layer and two hidden layers used for this work. The first hidden layer has eight neurons and the second hidden layer has four neurons, which is the same architecture as found in [9]. The input parameters are the UAS characteristics and initial conditions.

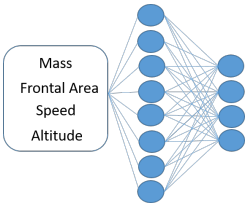

Figure 7: The Multilayer Perceptron of the Machine Learning Model

4.1.2 The Convolutional Neural Network

Convolutional Neural Networks are specialized for processing data with grid-like topology to learn patterns and spatial relationships, which makes them popular when images are inputs since images are grid-like arrays of pixels [10] [11]. The main components of the CNN are the convolutional layers and the pooling layers [12]. Most CNN architectures are a combination of convolutional and pooling layers with the final layer being a fully connected layer comprised of all the nodes of an input array unraveled into a single layer.

For this work, the grid-like topology input for the CNN is the building coverage and population density for a given location of interest. This input data takes the form of two arrays, population density and building layouts, of the surrounding area of the location of the UAS. These arrays are intended to only encompass the area of land where the UAS is likely to land. The cells in the array for the building coverage take the value of either 0 and 1 with 0 representing shelter and 1 representing no shelter. The values in the population density array are the number of people per square meter. Based on the CNN architecture in [9], a summary of the CNN can be seen in the image below. The CNN has an initial convolutional layer with 16 filters, then a max pooling layer, then another convolutional layer with 32 filters and max pooling layer, then the CNN is flattened and connected to a hidden layer with 16 neurons, and then finally one last layer with four neurons.

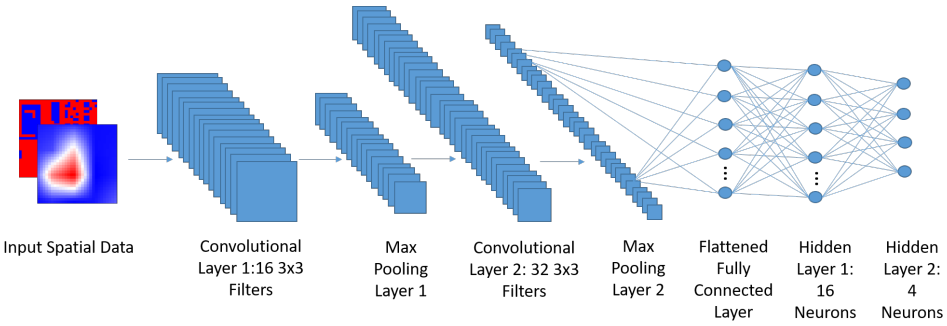

Figure 8: Illustration of the CNN Used

To estimate risk using the MLP and the CNN, the two output layers of each model were then concatenated together. Three additional layers were then added to the combined output. Two hidden layers, one with 10 neurons and one with five neurons, and finally an output layer with one neuron, were added to the model. The final output layer with one neuron is the predicted fatality rate. The image below shows the entire model.

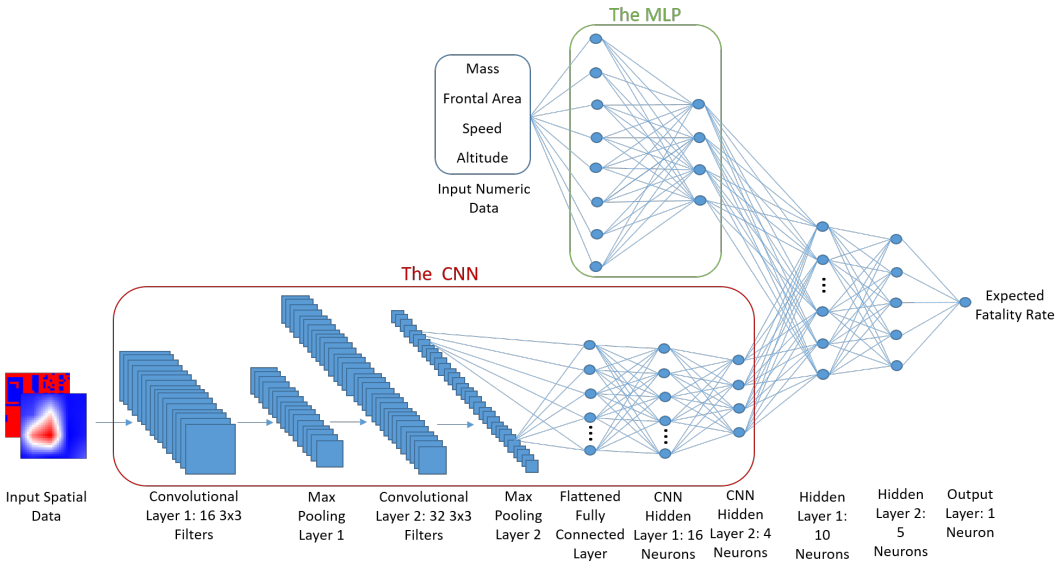

Figure 9: Combination of the MLP and the CNN to Predict Fatality Rates

Incorporating a ground risk assessment is vital to UAS operations to protect people on the ground as UAS become more ubiquitous in the rapidly evolving national airspace, especially in highly populated areas. The benefit of using a Machine Learning model like the one proposed over the physics-based models currently used is the reduction in computation time required. However, this benefit comes with the trade-off that there will be some error between the results of the Machine Learning model and the physics-based model. The main challenge with Machine Learning applications is how to reduce the error. Often times this error reduction can be accomplished by increasing the data used for training. By leveraging the physics-based model, a significant supply of data can be generated to train the Machine Learning model.

4.2 Generating Training Data

To train the Machine Learning model, data is required mapping the input conditions to the output fatality rate. Fortunately, the physics-based model can be used to generate this data by using the initial conditions of the UAS and the spatial data at the UAS location to calculate an expected fatality rate. An ideal Machine Learning model used for UAS risk assessment would be able to handle any combination of UAS speed, altitude, mass, frontal area, and velocity along with any distribution of shelter and population density, so the training data would also ideally reflect several of these combinations. However, this is not practical in a real world, so the objective is to include a large number of these combinations by observing the common operating conditions of UAS today. A summary of the ranges used for the UAS characteristics can be seen below and was created using off-the-shelf UAS characteristics.

Parameter	Minimum	Maximum
Mass (kg)	2.0	9.0
Frontal Area (m^2)	0.347	0.81
Speed (m/s)	5	35
Altitude (m)	15	140

Table 2: UAS Parameter Ranges

The minimum limit for the mass was based on the findings from [17] that show a significant difference in the chance of causing a neck injury between the 1.2 kg DJI Phantom 3 and the 3.1 kg DJI Inspire 1. The DJI Phantom 3 had little chance to cause a neck injury while the DJI Inspire 1 had a much higher chance, so a mass of 2 kg between the DJI Phantom and DJI Inspire was used as the minimum value for this work.

With the ranges for the UAS characteristics determined, finding a way to sample a large number of combinations of parameters is still required to train a Machine Learning model to be as robust as possible. For this, a Design of Experiments was used to efficiently explore the design space. A Latin Hypercube Design of Experiments was chosen because it is a space-filling design that creates design points evenly spread throughout the design space [18]. To ensure the design space was being thoroughly sampled, 248,000 data points of different combinations of UAS parameters were generated. However, the design points of the UAS parameters is only a portion of the data required, and spatial data for each of the design points was still needed.

To generate the input training spatial data, the building coverage and population density for the Georgia Institute of Technology were used. In Figure 1 and Figure 4, the population density and building coverage for the Georgia Institute of Technology are shown as 2D arrays, with each 10 m x 10 m cell being a unique location. The total size of the array was 113 rows and 135 columns, resulting in 15,504 unique locations. For each unique location of interest, a 29 cells x 29 cells window of the building coverage and population density was extracted and used as the spatial input data for that location of interest. The figure below demonstrates this concept. In the figure, the expected fatality rate is being estimated for the location identified by the marker.

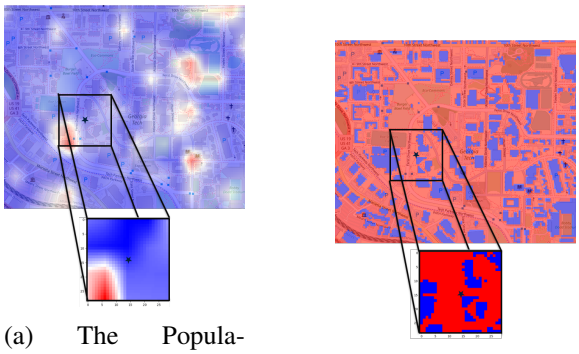

(a) The Population Density Data Extracted for One Location of Interest

(b) The Building Data Extracted for One Location of Interest

Figure 10: The Spatial Data Extracted for a Single Location of Interest

In an attempt to capture the effect of population density and shelter coverage on the fatality rate with the effect of the UAS conditions, 16 different UAS conditions were used at each location of interest to utilize the full data set generated using the Latin Hypercube Design of Experiments. With the combination of the input spatial data and the input numeric data, the physics-based model described in the paper was then used to estimate the expected fatality rate for each of the 248,000 cases.

A significant part of training any machine learning algorithm is cleaning the data used for training. After running each of the 248,000 input data points in the physics-based model to estimate the expected fatality rates, the output data was inspected to see if there were any outliers that might make training difficult. The raw data distribution for the resulting fatality rates can be seen below.

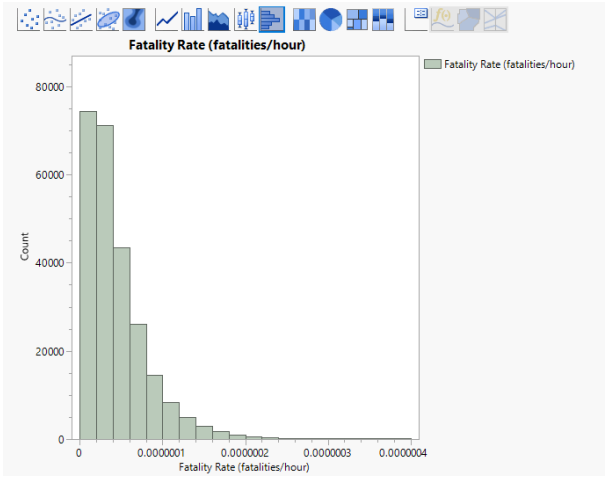

Figure 11: The Raw Data of Expected Fatality Rates Collected from the Physics-Based Model for 248,000 Cases

As seen in the figure above, the fatality rates calculated are heavily skewed with a range from 0 to $3.8x10^{-7}$ and values reaching as low as $2x10^{-10}$. This is problematic for a couple of reasons. Machine Learning algorithms have difficulty training when the data is not equally represented. In this case, there are few values at the high end of the fatality rate calculation and plenty data points with low risk. These high risk data points are not sufficiently represented. One step taken to fix this was to sample data points that resulted in a more uniform distribution. Part of this also included trimming some of the outlier data points at the high end as well in order to obtain enough points for training. A distribution of the data points used can be seen below. The number of data points in this distribution is 15,524, which is a considerable down sizing from the original 248,000. However, the full range of the UAS input characteristics is still represented in this data set to ensure there were no gaps in the design space for training. After cleaning the data, it was then used for the Machine Learning model training.

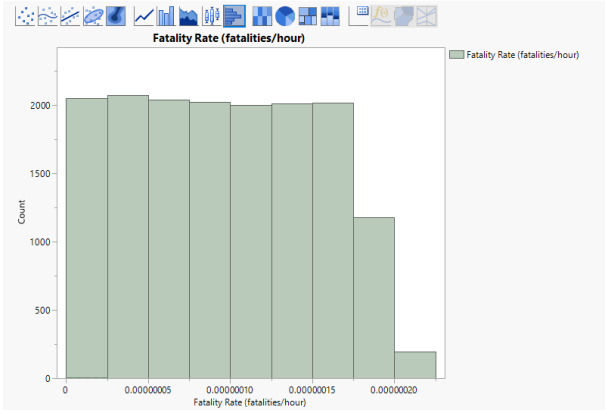

Figure 12: The Filtered Data of 15,524 Data Points

4.3 Machine Learning Model Training

A key aspect to any machine learning model training is the hyperparameters. Similar to how the model architecture can alter the performance, so can the hyperparameters set for training. These hyperparameters include the type of loss function used, the learning rate, batch size, and even the number of epochs used for training. The loss function used for this work was the Mean Absolute Percent Error

(MAPE). The batch size and number of epochs were 128 and 200, respectively. These values were the same as used in [9]. However, an early stop condition was implemented to prevent the model from overfitting the training data if the testing validation accuracy stopped improving.

Aside from the hyperparameters mentioned above, other hyperparameters that affect the performance of the model are those associated with the architecture of the model, like the number of layers or the number of neurons in each layer of the MLP. Changing the architecture can alter the performance, so several different architectures should be explored to find the best one. However, with an unlimited number of possibilities for architectures, it is infeasible to try them all. For this reason, it is recommended to find a suitable model that has decent performance as the baseline and fine tune the baseline model from there. The model described above in Section 4.1 is the baseline model used for this work. Parameters in the baseline model architecture subject to change include the number of neurons in each of the MLP layers, the number of filters in the CNN convolutional layers, the number of neurons in the CNN fully connected layer, and the number of neurons in the final two fully connected layers after the outputs of the CNN and MLP have been connected. A range of each of these parameters is summarized in the table below.

Model Parameter	Minimum	Maximum
MLP Hidden Layer 1 Neurons	4	100
MLP Hidden Layer 2 Neurons	4	100
Number of Filters in Convolutional Layer 1	4	32
Number of Filters in Convolutional Layer 2	4	128
CNN Hidden Layer 1 Neurons	4	32
CNN Hidden Layer 2 Neurons	4	32
Hidden Layer 1 Neurons	4	32
Hidden Layer 2 Neurons	4	32

Table 3: Hyperparameter Ranges for Machine Learning Model Training

The Machine Learning model training was done in Python using TensorFlow. Compatible with TensorFlow is another Python package, Keras. One of the functionalities of Keras is to automate the random search to find the optimal model configuration, so 200 random model combinations were generated from the ranges above and tested. Because there is some randomness in the actual training process, each model combination was trained three different times. The model with the best performance out of the 200 random combinations was saved.

4.4 Challenges Faced and Overcome

Often times in Machine Learning applications, brute force trial and error is the best way to find a model that works. When a model does not perform adequately it could be the result of poor data quality, poor hyperparameter selection, a poor selection of the most important factors, or a combination of these and many more. Pinpointing the exact cause can be quite challenging. Some of the problems faced in this work included problems with data quality and sub-optimal model architecture.

Right away the training data quality was problematic due to distribution of the raw data for the expected fatality rates. With how skewed the distribution was, there were several attempts at downsizing the data in a way to find enough data points to use without significantly reducing the total range of the data points due to the outliers. A sample of data that could have been chosen could have been those data points ranging from 10^{-9} to 10^{-8}, which would have been almost 27,000 data points for training, but then the Machine Learning model would only be suitable at estimating risk whose true value is within

that range. The higher the upper limit on the maximum value of the range, the less data points there were. A solution was found by removing the top 10% of the raw data to remove the outliers and taking data points that covered whole range of the data to make a uniform distribution.

Other challenges that came with the data were the zero risk values collected from the physics-based model. The loss function for the Machine Learning model training was the Mean Absolute Percent Error (MAPE), which has the truth value for the estimated risk in the denominator. With the truth value being zero, this resulted in very high MAPE. The solution to this was to put a lower limit on the fatality rates used for training. The floor value used was 10^{-9} fatalities per flight hour. Anything less than this value was deemed to be safe anyway since it was two orders of magnitude lower than the acceptable value of 10^{-7} fatalities per flight hour, and in fact it would be more conservative to assume this higher risk value anyway. It should be worth noting that even though the magnitude of the risk values were incredibly small, this did not result in the MAPE exploding like when the risk was zero because all the risk values were scaled by the maximum value prior to being used for training.

The raw data generated from the physics model not only appeared to be heavily skewed, but it also took the apppearance of having a logarithmic distribution. There were some efforts to scale the output data with a natural logarithmic and logarithmic with base 10. This scaling gave the data a more uniform distribution, and as a result more of the data points could be used. Initially, this data transformation resulted in very low errors on both the training dataset and the validation dataset, sometimes with a less than 2% MAPE. However, it was later revealed that this low error provided a false sense of accuracy. When the data was log shifted to train the model, the model would do poorly when deployed to create a risk map. This is because the predicted output from the model would need be transformed back to account for the log transformation done during training, so the same data point that would be 2% during training would be around 30%-40% off when the predicted value was transformed back. Ultimately, the approach to scale the output data using logarithmic transformations was scrapped.

Initially, training a model on the dataset was very difficult when the orders of magnitude were very different between the maximum and minimum points. Decreasing the range was not viewed as a good solution, so several attempts to increase the performance model were tried. These methods included discretizing the entire dataset range into bins and turning the regression problem into a classification problem. Instead of trying to predict the raw value for the risk, the model would attempt to estimate the class a datapoint belonged, and this class would be associated with the approximate risk value for the given datapoint. However, this approach failed to produce a model with satisfactory accuracy regardless of the number of classes the data was split into for training, so eventually this approach was scrapped as well. Other attempts to improve the model performance for regression included using pretrained CNNs available through TensorFlow. However, these models proved to be slow and did not boost the overall performance.

There were two changes that ultimately lead to greater success training the Machine Learning model. First, there was a reduction in the numeric input parameters. Choosing the right input parameters for a model is crucial, since too few can inhibit the model to learn on some information it might need, and too many might require a more complex model, making training more difficult. Initially, the drag coefficient was included in the numeric data as an input for the model. However, after further review of the data, it did not seem as if the drag coefficient played a significant role in the fatality rate. With the small range for the drag coefficient used, this could have attributed to the low affect the drag coefficient had on the risk values. Ultimately, the drag coefficient was removed as an input for the model, and this helped boost model performance.

The second big change that lead to an improvement in the model performance was the change in the structure of the input data and the CNN. With the spatial data input needing to capture the entire area a UAS might impact upon descent, the initial size for the input spatial data was 65 cells x 65 cells,

which would be equal to a plot of land about 650 m x 650 m. Upon further review of the input parameter range of the UAS, the maximum distance the UAS was predicted to land was less than 200 m away from the location of failure, suggesting the input spatial data was far too big. With the input spatial data being so large, this was proved difficult for the model to learn. This could be similar to how too many numeric parameters hindered the ability for the model to learn. Therefore, the input spatial data size was reduced from 65 cells x 65 cells to 29 cells x 29 cells to still ensure the entire area of UAS impact was captured.

5 Results and Discussion

5.1 Evaluating the Machine Learning Model

The baseline Machine Learning model and the Machine Learning model optimized with the hyperparameter random search both showed promising results when observing their MAPE on the training and validation dataset. A summary of their two architectures and their performance can be seen in the table below.

Model Parameter	Baseline Model	Optimized Model
MLP Hidden Layer 1 Neurons	8	54
MLP Hidden Layer 2 Neurons	4	54
Number of Filters in Convolutional Layer 1	16	24
Number of Filters in Convolutional Layer 2	32	4
CNN Hidden Layer 1 Neurons	16	32
CNN Hidden Layer 2 Neurons	4	32
Hidden Layer 1 Neurons	10	54
Hidden Layer 2 Neurons	5	12
Training MAPE (%)	16	15
Validation MAPE (%)	22	17

Table 4: Summary of Machine Learning Models

Based on the table above, it is shown that both the baseline model and the optimized model performed similarly with the training data, but the optimized model performed a little better on the validation data. However, the two models still need to be compared to the physics-based model.

To compare the physics-based model with the Machine Learning models, a risk map was created for the campus of the Georgia Institute of Technology using the daytime population density information from LandScan combined with the social media activity collected. A UAS was assumed to have a mass of 6 kg, a frontal area of 0.6 m^2, a flight altitude of 35 m and a flight speed of 25 m/s operating over the campus. With this UAS configuration, a risk map was finally created using the physics-based model after hours required for completion while the Machine Learning models were able to complete the risk map in a matter of seconds with the same UAS conditions. Although the optimized model seemed to do better during training, the baseline model performed better when compared to the risk map created by the physics-based model. When compared to the physics-based risk map, the MAPE for the baseline was 19.3% while it was 22.3% for the optimized model, so the baseline model was selected as the superior model. A comparison between the risk map created using the baseline Machine Learning model and the physics-based risk map can be seen in the image below. As mentioned, it took several hours to create the physics-based risk map and a few seconds to create the Machine Learning risk map. The heat map in each figure represents the expected fatality rate, measured as fatalities per flight hour.

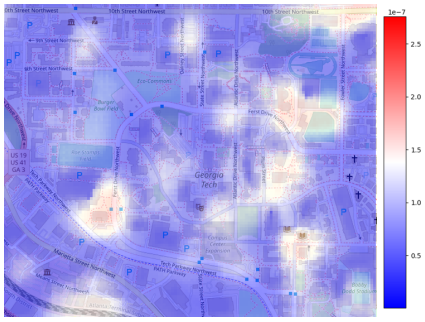

(a) Risk Map Created Using the Machine Learning Model

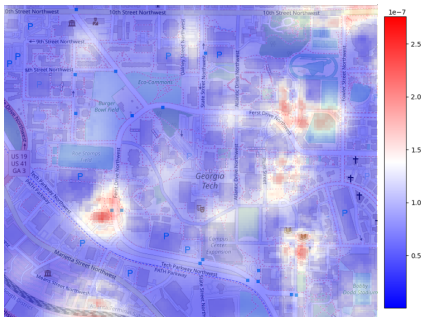

(b) Risk Map Created Using the Physics-Based Model

Figure 13: Comparison Between the Machine Learning Model and Physics-Based Model

The comparison above shows the Machine Learning can identify the same high risk areas as the physics-based model, although the predicted risk from the Machine Learning model in these high risk areas appears to be less than that from the physics-based model. However, the MAPE for the Machine Learning model creating this risk map was still only about 19%. This can be interpreted as if the physics-based model determines the risk for a single location is $1x10^{-6}$ fatalities per flight hour, then the Machine Learning model might predict the risk value to be $0.81x10^{-6}$ fatalities per flight hour. That is one fatality every 114 years compared to a predicted 0.81 fatalities every 114 years. With how low this frequency is, the 19% MAPE was deemed as reasonable, although future work will include trying to minimize the MAPE further.

One benefit of using the Machine Learning model is the rapid generation of risk maps for UAS users to identify how the UAS flight conditions might affect risk. The image below shows two different risk maps created for two different UAS configurations.

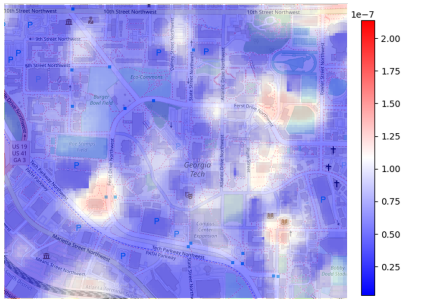

(a) Risk Map Created Using the Machine Learning Model for Low Risk UAS: Mass=2 kg, Frontal Area=0.4 m^2, Speed=25 m/s, Altitude=35 m

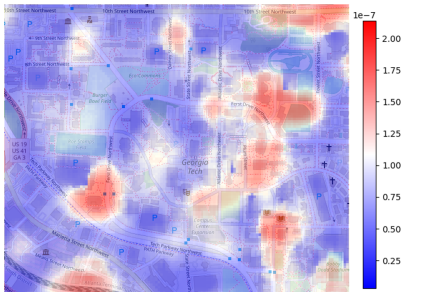

(b) Risk Map Created Using the Machine Learning Model for High Risk UAS:: Mass=8 kg, Frontal Area=0.75 m^2, Speed=30 m/s, Altitude=35 m

Figure 14: Comparison of Risk Maps for Different UAS Conditions

By altering the UAS flight conditions and UAS parameters, UAS pilots can observe how the risk map changes. As expected, a larger and faster UAS increases the risk compared to a smaller and slower UAS. This ability to rapidly create risk maps with changing UAS flight parameters is meant to assist UAS pilots in the route planning to ensure the maximum level of allowable safety is not exceeded.

5.2 Risk-Informed Route Planning

Pattent, the industry partner for this work, is a small startup specializing in creating low risk routes for UAS in urban areas. Their route planning solution uses population density and building coverage to minimize the time spent over people and maximize the time spent over buildings that can provide coverage to people on the ground. The route algorithm they use also accounts for physical obstacles like tall buildings and no fly zones specified by the FAA facility maps to find suitable routes. Pattent has provided their expertise and route planning technology to assist in the current efforts of using the risk map to find low risk routes.

The route planning algorithm provided by Pattent is a proprietary algorithm based on a modified version of the popular A* algorithm. This modification changes the heuristic function used in A* to be a weighted combination of safety and distance. As a result, the algorithm can be used to find a combination of the safest route or the fastest route as desired. For this work, the Machine Learning model can be used to create a risk map to use with the route planning algorithm. This allows for the algorithm to find the safest route based on minimizing the risk of the path using the risk map. This route planning algorithm was used to find the safest route between two points at the Georgia Institute of Technology using both risk maps created by the physics-based model and the Machine Learning model to compare the resulting paths. The comparison of the routes created can be seen below where the red route is the route created using the physics-based risk map and the black route is the route created using the Machine Learning model. The heat map in the figure is the risk map associated with the Machine Learning model.

Figure 15: Comparison of the Routes Created Using the Machine Learning Risk Map and the Physics Risk Map

The two routes created using each of the risk maps are strikingly similar, suggesting the Machine Learning model is adequate at replacing the physics-based model for assisting in route planning. Additional confirmation of this is found when comparing the predicted risk and actual risk of the route created using the Machine Learning risk map. The table below summarizes the risk for this route. The predicted risk values and actual risk values are the risk values obtained using the Machine Learning risk map and physics-based risk map, respectively.

Risk	Predicted	Actual	Absolute Percent Error (%)
Maximum Risk (fatalities/hour)	$8.55x10^{-8}$	$7.13x10^{-8}$	19.83
Average Risk (fatalities/hour)	$1.72x10^{-8}$	$1.79x10^{-8}$	4.19

Table 5: Summary of Risk Using the Machine Learning Risk Map for Route Planning

From the table above, the error between the predicted maximum risk level and the actual risk level is greater than the error of the predicted average risk and the actual average risk. As stated previously, the Machine Learning model has a higher error with the higher risk areas, but with the model still able to identify the high risk regions, it correctly avoids these areas when assisting in route planning. The error between the predicted average risk and the actual average risk is very low, suggesting that as long as the route planner attempts to avoid the high risk areas, the predicted risk values will be very close to the actual risk values. It should also be noted that both routes have a maximum risk value less than the acceptable risk level of 10^{-7} mentioned previously.

One of Pattent's products, Route Scout, is a web-based application to create routes for UAS users. Using the web-based user interface, a UAS user enters the desired start and endpoint for the UAS, and the user interface shows the resulting route and allows the user to download the list of GPS waypoints for the UAS to follow. The Machine Learning model created in this work was implemented in a prototype version of the next iteration of the Route Scout, allowing users not only enter in the desired start and end point, but to use the desired UAS flight characteristics to create a more informed risk assessment. The user can also specify the maximum level of safety required. The figure below displays the user interface and routes created using two different UAS configurations with a desired maximum risk level less than 10^{-7} fatalities per flight hour.

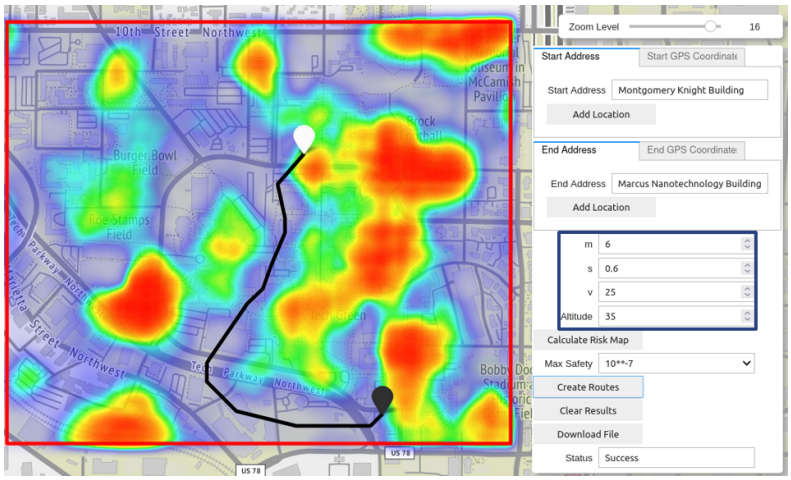

(a) Route Planner for Low Risk UAS

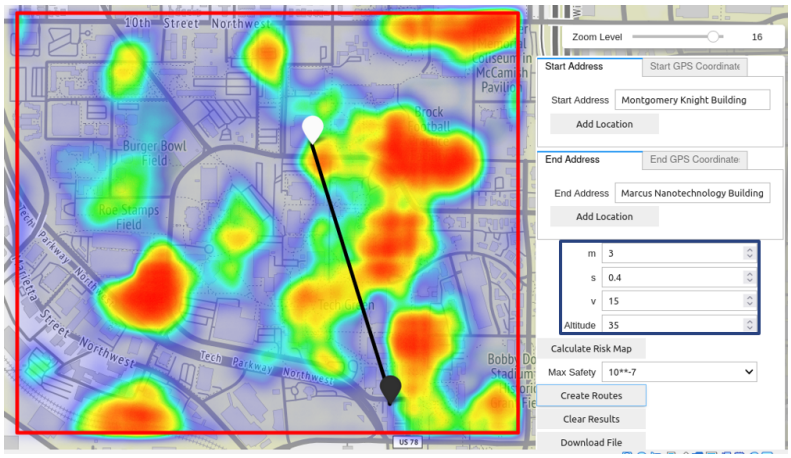

(b) Route Planner for High Risk UAS

Figure 16: Comparison of Routes Generated for Different UAS Conditions

Aside from altering the UAS characteristics to find a route, the UAS user can adjust the target

level of safety based on the application. Some applications, like law enforcement and first responders, may prioritize response time over safety level, so there is justification to relax the maximum level of safety. Using the user interface, the UAS user can make these changes accordingly to emphasize arrival time. The images below show the resulting routes created for the same UAS configuration with different levels of safety required. As expected, the route that allows for higher rates of fatality takes the more direct path over the high risk areas.

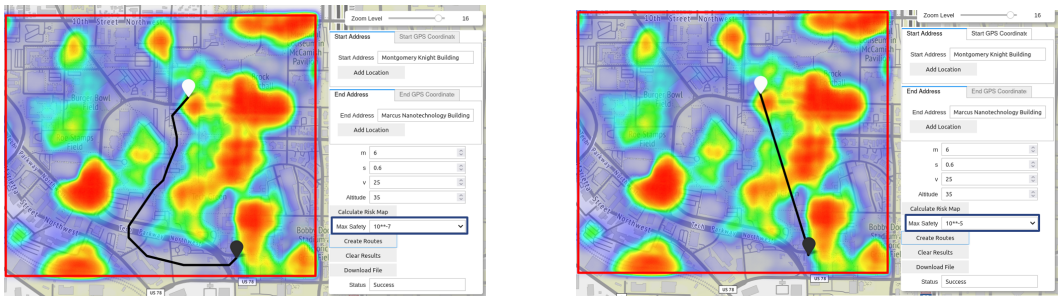

(a) Route Planned with Maximum Risk Level 10^{-7} Fatalities per Flight Hour

(b) Route Planned with Maximum Risk Level 10^{-5} Fatalities per Flight Hour

Figure 17: Comparison of Routes Generated for Different Levels of Desired Safety

The objective of the Route Scout was to provide UAS pilots an easy way to plan risk-informed routes flexible to their need. The user interface was designed to be intuitive and easy to use. Whether it is law enforcement or delivery companies, all UAS users can utilize this technology by using it to create safe routes. The key enabler for this is the Machine Learning model that allows for quick risk assessment to ensure any route planned does not exceed a maximum level of allowable safety without taking significant time to determine expected risk values.

6 Conclusion

Similar to manned aircraft, UAS require risk assessment to ensure any UAS is operating with some acceptable level of risk. Because the UAS has been introduced relatively recently compared to manned aircraft, the UAS do not have the comparable historical flight data that would be required for such risk assessment. As a result, modeling and simulation methods have become the next best option to approximate UAS risk levels. The state of the art methods for estimating UAS ground risk rely on physics-based models to determine the descent trajectory for a UAS given a failure and then determine the likelihood of causing a fatality if the UAS were to strike a person. This process can become computationally expensive, and is not suitable for dynamic environments like cities. This work shows that Machine Learning methods can be used to replace the slow physics-based methods. Using the physics-based methods to generate training data, a Machine Learning model was trained with a 16% error on the training data and a 22% error on the validation data. The Machine Learning model allows for UAS users to rapidly generate a ground risk map based on their desired UAS flight conditions. With this risk map, it was shown that a route planner can be used to find a path that does not exceed some allowable risk value set by the user. This solution is presented as a flexible web-based application that can be used by any UAS pilot operating in highly populated areas. Future work would include increasing the fidelity of the physics-based model by incorporating additional descent trajectories and the effect of wind. Additionally, more diverse population densities and building layouts should be added to the training to increase robustness.

References

[1] K. Dalamagkidis, "On Integrating Unmanned Aircraft Systems into the National Airspace System". SpringerLink, 2009, ISBN: 978-90-481-2096-3

[2] T. Perez, R. A. Clothier, and B. Williams, “Risk-management of UAS robust autonomy for integration into civil aviation safety frameworks.”

[3] J. Fortes, R. Fraga, K. Martin, "Safety Analysis for UAS Operation Using Stochastic Fast-Time Simulation"

[4] S. Primatesta, "An Innovative Algorithm to Estimate Risk Optimum Path for Unmanned Aerial Vehicles in Urban Environments" International Conference on Air Transport. 2018.

[5] S. Primatesta, "Ground Risk Map for Unmanned Aircraft in Urban Environments" Journal of Intelligent and Robotic Systems. 2019.

[6] A. Cour-Harbo, "Quantifying Risk of Ground Impact Fatalities for Small Unmanned Aircraft" Journal of Intelligent and Robotic Systems. 2018.

[7] A. Cour-Harbo "Ground Impact Probability Distribution for Small Unmanned Aircraft in Ballistic Descent" 2020 International Conference on Unmanned Aircraft Systems. 2020.

[8] A. Geron, "Hands-On Machine Learning with Scikit-Learn and TensorFlow." Published by O’Reilly Media, Inc. 2017.

[9] A. Rosenbrock, "Keras: Multiple Inputs and Mixed Data" pyimagesearch. 2019. https://pyimagesearch.com/2019/02/04/keras-multiple-inputs-and-mixed-data/

[10] I. Goodfellow, "Deep Learning" MIT Press. 2016.

[11] K. Sangvhi, "Image Classification Techniques" Analytics Vidhya. 2020.

[12] "Convolutional Neural Networks" IBM. https://www.ibm.com/topics/convolutional-neural-networks

[13] M. Mishra, "Convolutional Neural Networks, Explained" Towards Data Science. 2020.

[14] J. Breunig et. al. "Modeling Risk-Based Approach for Small Unmanned Aircraft Systems" The MITRE Corporation. 2018.

[15] Oak Ridge National Laboratory LandScan Database URL:https://web.ornl.gov/sci/landscan

[16] OpenStreetMap contributors, "Planet dump retrieved from https://planet.osm.org" 2017. URL: https://www.openstreetmap.org

[17] E. Campolettano et. al. "Ranges of Injury Risk Associated with Impact from Unmanned Aircraft Systems" Annals of Biomedical Engineering. 2017.

[18] R. Myers. "Response Surface Methodology Third Edition" John Wiley and Sons, Inc. 2009.