

Risk-Aware Planning and Assignment for Ground Vehicles using Uncertain Perception from Aerial Vehicles

Vishnu D. Sharma[†], Maymoonah Toubeh[†], Lifeng Zhou[†], Pratap Tokekar

Abstract—We propose a risk-aware framework for multi-robot, multi-demand assignment and planning in unknown environments. Our motivation is disaster response and search-and-rescue scenarios where ground vehicles must reach demand locations as soon as possible. We consider a setting where the terrain information is available only in the form of an aerial, georeferenced image. Deep learning techniques can be used for semantic segmentation of the aerial image to create a cost map for safe ground robot navigation. Such segmentation may still be noisy. Hence, we present a joint planning and perception framework that accounts for the risk introduced due to noisy perception. Our contributions are two-fold: (i) we show how to use Bayesian deep learning techniques to extract risk at the perception level; and (ii) use a risk-theoretical measure, CVaR, for risk-aware planning and assignment. The pipeline is theoretically established, then empirically analyzed through two datasets. We find that accounting for risk at both levels produces quantifiably safer paths and assignments.

I. INTRODUCTION

Consider scenarios where an autonomous ground vehicle must navigate in an unknown environment. Examples include search-and-rescue, space exploration, and disaster response. For instance, consider a disaster response scenario where ground vehicles must supply resources to specific demand locations as soon as possible. In such settings, prior GPS or satellite maps of the environment may no longer be valid. Instead, an aerial robot may be employed to take live aerial images which can then be used to plan the paths of the ground vehicles towards the demand locations. However, due to the inherent uncertainty of aerial images, the paths that are found may not actually represent the situation on the ground. Therefore, there is a risk of the vehicles taking longer to reach the demand positions than planned. In safety-critical situations, one way to mitigate the risk is to assign multiple vehicles to the same demand, with the earliest arriving one actually responding to the demand.

Motivated by this scenario, we study the problem of how to find risk-aware paths for multiple vehicles to serve multiple demand locations. There are two problems to be solved — assigning the vehicles to the demand locations and finding risk-aware paths from start to demand locations.

We present a risk-aware framework to solve both problems simultaneously.

The environment where the vehicles navigate is captured by an overhead image. We implement a deep learning technique for semantic segmentation of the overhead image. Due to the uncertainty from segmentation, the travel cost of the vehicle turns out to be a random variable. Built on our previous work [1], our first contribution is to show how to utilize Bayesian deep learning techniques to handle the risk from the planning and perception level. After risk-aware planning and perception, we generate a set of candidate paths corresponding to different risk levels from each vehicle’s start position to each demand location. Our second contribution is to assign each vehicle to a risk-aware path from its candidate path set to a demand. We utilize a risk measure, Conditional-Value-at-Risk (CVaR), to manage the risk from the uncertainty. Our assignment framework provides the flexibility to trade off between risk and reward, which builds on our previous work [2], with risk here being assessed at multiple levels of the algorithm. Our empirical results show that this risk-aware framework results in safer path planning and assignment.

Related Work: Deep learning has shown significant improvements in perception capabilities for many robotics applications. However, the potential of the positive impact deep learning may have on real-world scenarios is inevitably proportionate to their interpretability and applicability to imperfect environments. In these cases, deep neural networks can even misrepresent data outside the training distribution, giving predictions that are incorrect without providing a clear measure of certainty associated with the result [3]. The extraction of uncertainty information, as opposed to the reliance on point estimates, is crucial in safety-critical applications, such as autonomous navigation in an urban, unstructured setting. Methods like Natural-Parameter Networks [4] propose modeling the network parameters and inputs as Gaussian distributions. However, these modifications impose huge computation cost on the model due to the increment in the number of trainable parameters. Lightweight Probabilistic Deep Networks [5] alleviates this concern to some extent by making the weights deterministic. Large size networks are also unsuitable for real-time robotics applications which may have constraints on inference time and memory. [6] and [7] propose methods that allow uncertainty extraction from deep learning models, specifically those that do not interfere with the overall structure or training process. In this work, we leverage [6], which shows that dropouts, which are often used as a regularization enhancement in neural networks,

M. Toubeh, L. Zhou are with the Department of Electrical & Computer Engineering, Virginia Tech, Blacksburg, VA 24061, USA. {may93, lfzhou}@vt.edu. V.D. Sharma and P. Tokekar were with Virginia Tech during completion of this work and are now with the Department of Computer Science, University of Maryland, College Park, MD 20742, USA (email: {vishnuds, tokekar}@umd.edu).

[†]These authors contributed equally.

This work is supported by the Office of Naval Research under Grant No. N000141812829.

can provide approximate Bayesian inference and thus help in uncertainty estimation for the deep learning models.

In addition to considering uncertainty at the perception level, we also utilize some popular risk measures to handle uncertainty at the planning and assignment levels. A typical measure for optimization under uncertainty is the expectation of a stochastic function. However, the expectation is a risk-neutral measure and may not be desirable, especially in critical tasks, like search-and-rescue operations [8]. For example, we may find a path with a lower expected cost but with high variance. It is quite likely that the vehicle may encounter a much larger cost (as compared to the expected one) when traveling on this path. Thus, instead of using expected cost, we utilize some other risk-aware measures, such as mean-variance [9], [10] and conditional-value-at-risk (CVaR) [8], [11].

In particular, we use mean-variance as the risk-aware cost measure in the A* planner [12] for planning paths for the vehicles. The mean-variance measure allows us to balance the mean cost and uncertainty (variance) when planning paths. Also, for the path assignment, we use CVaR to deal with the uncertainty on the path level. CVaR_α explicitly takes into account the risk associated with bad scenarios [8], [11]. Specifically, CVaR_α measures the expectation of a random variable in the 100α -percentile worst scenarios. Here, the user-defined risk-level, $0 < \alpha \leq 1$, provides a user with the flexibility to choose a risk that they would like to take. Setting $\alpha = 1$ makes CVaR_α equal to the expectation whereas $\text{CVaR}_\alpha \approx 0$ is akin to worst-case optimization.

Risk in autonomous systems ultimately involves the risk of failure, whether due to damage to the agents or their environment or mission failure. Several risk-aware planning frameworks have been proposed for collision avoidance, such as the dynamic risk density function in [13] which infers congestion cost from sensor inputs. Our navigation cost function observes the risk associated with noisy sensor perception and do not assume sensor perception is always correct. Uncertain semantic maps have also been proposed for perception-based planning. Both [14] and [15] model an unknown environment as a semantic map by assuming Gaussian distributions over landmark positions for simultaneous localization and planning (SLAM). These address the uncertainty in planning, but not in perception, i.e. uncertain landmark classification.

Contributions. In this paper, we have three main contributions.

- We present a framework that plans and assigns risk-aware paths for robots that navigate in unknown environments.
- We utilize the Bayesian deep learning technique to learn an unknown environment whose information is only available by an overhead, georeferenced image.
- We deal with the uncertainty at both path planning and assignment levels by optimizing the corresponding risk measures. In the end, we assign each vehicle a risk-aware path to a demand location and the path assignment is guaranteed to have a bounded approximation

performance of the optimal assignment.

This work builds on our prior work where we studied these two problems (uncertainty extraction from deep learning and CVaR optimization) individually. Here, we investigate the joint problem. We find that utilizing the uncertainty extraction from deep learning and managing the risk from uncertainty by CVaR optimization provide the vehicles with safer and risk-aware paths in unknown environments.

II. PRELIMINARIES

We start by defining the notations used in the paper. We give background on a risk measure, conditional-value-at-risk (CVaR). We then provide a formal definition of the joint problem of planning and assignment.

A. Conditional Value at Risk

Let X be a random variable. $\text{CVaR}_\alpha(X)$ denotes the expectation on the α -worst scenarios of the utility or cost function f with $\alpha \in (0, 1]$. More specifically, if X indicates reward or benefit, $\text{CVaR}_\alpha(X)$ denotes the expectation on the left α -tailed scenarios. While, if X represents loss or penalty, $\text{CVaR}_\alpha(X)$ is the expectation on the right α -tailed cases. Here, α is the confidence level or the risk level. If α is close to 0, CVaR_α is close to the worst-case. If α is equal to 1, CVaR_α is same as the expectation.

In this paper, the utility function $f(\mathcal{S}, y)$ defined on set \mathcal{S} is a random variable with randomness induced by parameter y . Since $f(\mathcal{S}, y)$ is utility that indicates benefit, $\text{CVaR}_\alpha[f(\mathcal{S}, y)]$ denotes the expectation on the left α -tailed cases, as shown in Figure 1.

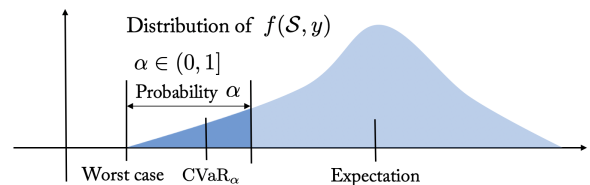


Fig. 1. CVaR_α of function $f(\mathcal{S}, y)$.

We generally maximize $\text{CVaR}_\alpha[f(\mathcal{S}, y)]$ by solving

$$\max_{\mathcal{S}, \tau} \tau - \frac{1}{\alpha} \mathbb{E}[(\tau - f(\mathcal{S}, y))_+], \quad (1)$$

where \mathcal{S} is a decision set (or solution set), $(t)_+ = \max(t, 0)$, and $\tau \in \mathbb{R}_+$ is an auxiliary parameter. For the ease of expression, we define

$$H(\mathcal{S}, \tau) = \tau - \frac{1}{\alpha} \mathbb{E}[(\tau - f(\mathcal{S}, y))_+] \quad (2)$$

B. Problem Formulation

We consider the problem of finding paths for multiple vehicles to serve multiple demand locations. In particular, we are given N vehicles' start positions, $\mathcal{V} = \{1, \dots, N\}$ and M demand locations, $\mathcal{D} = \{1, \dots, M\}$ in the environment. The environment is represented by an overhead, georeferenced RGB image as shown in Figure 2. The goal is to find offline paths for each vehicle such that they collectively

serve all the demands using navigation costs derived from the overhead images.

The cost of a path in the environment can be estimated by first performing a semantic segmentation of the overhead image. However, semantic segmentation is typically imperfect [3], and as such the estimated cost of a path may not be accurate. The problem we address in this paper is that of finding paths for vehicles to collectively serve all demands under travel-cost uncertainty.

We are motivated by tasks that are urgent and time-critical, such as fighting fires [16] and delivering medical supplies in emergencies [17]. When the number of vehicles is more than the demands, assigning multiple redundant vehicles to demands helps counter the effect of uncertainty [18]. The waiting time at a demand location is the time taken by the earliest vehicle to arrive at that location. If the travel times are deterministic, then it is known in advance which vehicle will arrive first. When travel times are uncertain, as in this work, the arrival time of the earliest vehicle itself is a random variable. The goal is to assign vehicles to demand locations and find corresponding paths for the vehicles from the start to the assigned demand locations.

For convenience, we convert the minimization problem into a maximization one to make it submodular by taking the reciprocal of the travel cost. Specifically, we use the travel *efficiency*, the reciprocal of travel cost, as the measure. Thus, the travel efficiency of a demand location is the maximum of the travel efficiencies of the vehicles that reach this demand location. The overall travel efficiency, denoted by f is the sum of the travel efficiencies of all demand locations. Notably, f is also a random variable.

Our goal is to find risk-aware paths from vehicles' start positions to demand locations given a user-defined risk level α . We formulate a risk-aware path finding problem by maximizing CVaR_α on the travel efficiency (Problem 1).

Problem 1 (Risk-Aware Path Finding)

$$\max_{\mathcal{S} \subseteq \mathcal{X}} \text{CVaR}_\alpha[f(\mathcal{S}, y)] \quad (3)$$

where \mathcal{S} is a set of paths for vehicles (“per path per vehicle”), \mathcal{X} is a ground set of paths from which \mathcal{S} is chosen, and $f(\mathcal{S}, y)$ is the travel efficiency on the path set \mathcal{S} , with randomness induced by y .

III. FRAMEWORK AND ANALYSIS

The framework we propose consists of three main parts: semantic segmentation, candidate path generation, and risk-aware assignment. The overall pipeline of the framework and its parts are shown in Figure 2. The inputs to the framework are a single overhead aerial two-dimensional image, the vehicles' start positions, and the demand locations. The output is a risk-aware assignment of paths from vehicles' start positions to demand locations.

The input is first semantically segmented into per pixel labels. These labels are assigned a cost proportionate to the risk involved in traversing them. The cost map and the

uncertainty associated with the segmentation are then used as inputs to a path planner which generates candidate paths for assignment. Finally, the candidate paths from each vehicle's start position to each demand location are computed by maximizing CVaR , for risk-aware path assignment.

A. Semantic Segmentation

To plan a path for a vehicle in the environment, we need to first recognize features in the image, such as road, person, car, and so on. Given an input aerial image provided by an unmanned aerial vehicle (UAV), a deep learning model is used to provide a semantically segmented map. We then utilize an approximation technique, using dropout, to learn the uncertainty in the semantic labels provided by the model [6]. The advantage of using uncertainty has been studied [3]. After semantic segmentation and uncertainty extraction, we assign different cost values for the robot traveling on different terrain features, similar to the cost function in [1]. For example, the cost of traveling on the road is less than that of traveling on vegetation. Traveling on a person and cars is impossible, and thus, the corresponding cost can be set to infinity. Notably, the cost can be the energy or time spent by a vehicle traveling on the terrains. Since there exists uncertainty in the terrain features provided by the deep learning model, the cost of traveling on them is also a random variable with some uncertainty.

B. Candidate Path Generation

Once a risk-aware cost map is generated based on the semantic segmentation of the terrain, augmented by the confidence in prediction, candidate paths are generated from each vehicle location to each demand location by an A* planner. For the different combinations of vehicles and demands, a candidate path represents a potential feasible route given on the map. Relying on deep learning segmentation alone can be risky. Given confidence information, we expect paths to avoid regions of high uncertainty which could involve out-of-distribution data or misclassifications, as shown in Figure 3.

For the A* planner, a risk-aware cost function $C(x)$ is defined on each pixel x . It combines the cost associated with the classified terrain type and the variance.

$$\hat{C}(x) := C(l_x) + \lambda \text{var}(l_x), \quad (4)$$

where C refers to the risk-neutral mean cost associated with the most likely predicted label l_x for a pixel x given multiple softmax outputs. The risk-aware function \hat{C} assigns a cost to each pixel x , characterized by a user-defined cost mapping and a weight parameter λ to quantify emphasis on the uncertainty. We use the variance in the prediction as a measure of uncertainty.

C. Risk-Aware Path Assignment

Consider a set of K candidate paths $\mathcal{P} = \{1, \dots, K\}$ generated from each vehicle's start position to each demand location in Section III-B. We then assign each vehicle a path to a demand location. As mentioned in Section II-B, we follow a redundant assignment setting [18], where each

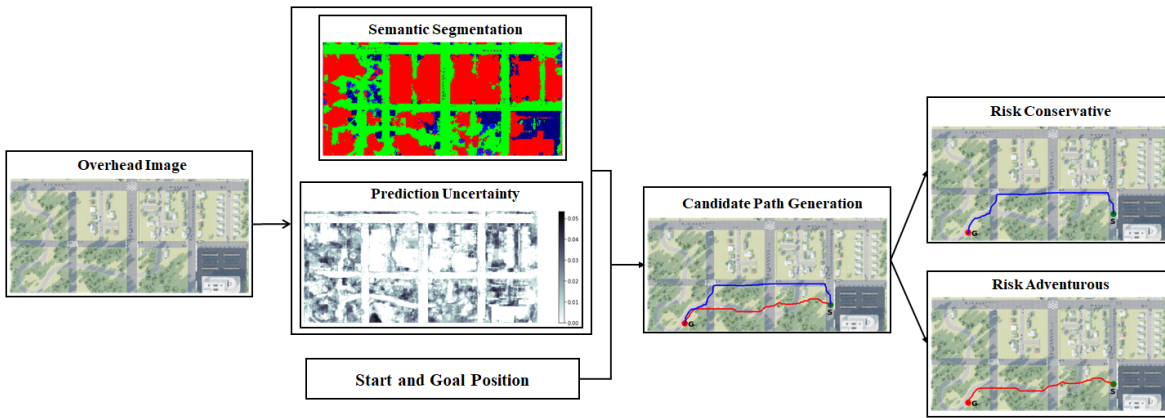


Fig. 2. The breakdown of the framework's parts. Given an overhead image input, the algorithm provides a semantic segmentation and uncertainty map, then generates candidate paths, and finally performs the risk-aware path assignment of vehicles to demands.

vehicle can be assigned to at most one demand location and multiple vehicles can be assigned to the same demand location. Only the vehicle that arrives through a path with with the highest efficiency is chosen at each demand location.

The travel efficiency is denoted for a vehicle starting at $i \in \mathcal{V}$ arriving at demand location $j \in \mathcal{D}$ through path $k \in \mathcal{P}$ as e_{ijk} . Correspondingly, we take the tuple (i, j, k) as an assignment where the vehicle-path pair (i, k) is assigned to the demand location j . The total efficiency is denoted at all M demand locations as

$$f(\mathcal{S}, y) = \sum_{j \in \mathcal{M}} \max_{(i, j, k) \in \mathcal{S}_j} e_{ijk} \quad (5)$$

where \mathcal{S}_j denotes the assignment set $\{(i, j, k)\}$ to the demand location j . The total assignment set $\mathcal{S} := \bigcup_{j=1}^M \mathcal{S}_j$ is a collection of the assignment sets at all demand locations. Notably, since each vehicle-path pair (i, k) can be assigned to at most one demand location, all \mathcal{S}_j (s) are disjoint, i.e., $\mathcal{S}_j \cap \mathcal{S}_{j'} = \emptyset, \forall j \neq j', j, j' \in \mathcal{D}$. This is called a partition matroid constraint in the literature [19], denoted by \mathcal{I} .

We use the “max” operator to capture the selection rule that only the vehicle-path pair (i, k) with the maximum efficiency is chosen at each demand location. Due to the “max” operator, the total efficiency $f(\mathcal{S}, y)$ turns out to be monotone submodular in \mathcal{S} . If there is no assignment, we set $f(\emptyset, y) = 0$ to normalize f . Here, y indicates the randomness of $f(\mathcal{S}, y)$ due to the uncertainty in travel efficiency.

Then, by utilizing Equation 1 and the candidate paths generated in Section III-B, we transform Problem 1 to a risk-aware path assignment problem below.

Problem 2 (Risk-Aware Path Assignment)

$$\begin{aligned} & \max_{\mathcal{S} \subseteq \mathcal{X}} \tau - \frac{1}{\alpha} \mathbb{E}[(\tau - \sum_{j \in \mathcal{M}} \max_{(i, j, k) \in \mathcal{S}_j} e_{ijk})_+] \\ & \text{s.t. } \mathcal{S} = \bigcup_{j=1}^M \mathcal{S}_j, \mathcal{S} \in \mathcal{I} \quad \text{and} \quad \tau \in [0, \Gamma], \end{aligned} \quad (6)$$

with \mathcal{S} the path assignment set (“per path per vehicle”), \mathcal{I} the partition matroid constraint [19], \mathcal{X} the ground set of all

Algorithm 1: Risk-Aware Path Assignment

Input:

- Vehicles’ initial positions \mathcal{V} , demand locations \mathcal{D} , and path set \mathcal{P}
- User-defined risk level $\alpha \in [0, 1]$
- Range of the parameter $\tau \in [0, \Gamma]$ and discretization stage $\Delta \in (0, \Gamma]$
- An oracle \mathcal{O} that approximates $H(\mathcal{S}, \tau)$ as $\hat{H}(\mathcal{S}, \tau)$

Output:

- Path assignment \mathcal{S}

```

1:  $\mathcal{M} \leftarrow \emptyset$ 
2: for  $i \in \{0, 1, \dots, \lceil \frac{\Gamma}{\Delta} \rceil\}$  do
3:    $\tau^i = i\Delta$ 
4:    $\mathcal{S}^i \leftarrow \emptyset$ 
5:   for  $l = 1 : |\mathcal{D}|$  do
6:      $(i^*, j^*, k^*) =$ 
        $\operatorname{argmax}_{i \in \mathcal{V}, j \in \mathcal{D}, k \in \mathcal{P}} \hat{H}(\mathcal{S}^i \cup (i, j, k), \tau^i) - \hat{H}(\mathcal{S}^i, \tau^i)$ 
7:      $\mathcal{S}^i \leftarrow \mathcal{S}^i \cup (i^*, j^*, k^*)$ 
8:      $\mathcal{V} = \mathcal{V} \setminus i^*$ 
9:   end for
10:   $\mathcal{M} = \mathcal{M} \cup \{(\mathcal{S}^i, \tau^i)\}$ 
11: end for
12:  $(\mathcal{S}, \tau^*) = \operatorname{argmax}_{(\mathcal{S}^i, \tau^i) \in \mathcal{M}} \hat{H}(\mathcal{S}^i, \tau^i)$ 

```

possible assignments, $\{(i, j, k)\}, k \in \mathcal{P}, i \in \mathcal{V}, j \in \mathcal{D}$, and $\Gamma \in \mathbb{R}^+$ the upper bound of the parameter τ . Γ can be set as an upper bound on $f(\mathcal{S}, y)$ (Eq. 5).

Building on the sequential greedy algorithm (SGA) from our previous work [2], we present Algorithm 1 for solving Problem 2.

There are four stages in Algorithm 1:

a) *Initialization (line 1):* We initialize a storage set \mathcal{M} to be empty. We use \mathcal{M} to store the assignment \mathcal{S} with the corresponding τ when searching all possible values of τ .

b) *Searching for τ (for loop in lines 2–11):* We sequentially search for all possible values of τ within $[0, \Gamma]$ by a user-defined separation Δ (line 3). Γ is an upper bound on

τ and can be set based on the specific problem at hand. We show how to compute Γ in specific scenarios in Section IV.

c) *Greedy algorithm (lines 4–9)*: For a specific τ , e.g., τ_i , we use the greedy approach [19] to choose the corresponding assignment set \mathcal{S}^i . We first initialize set \mathcal{S}^i to be the empty set (line 4). Then we execute the greedy algorithm in $|\mathcal{D}|$ rounds (line 5), since the total number of demand locations to be served is \mathcal{D} . At each round, the assignment (i^*, j^*, k^*) which gives the maximum marginal gain of $\hat{H}(\mathcal{S}^i, \tau^i)$ is selected (line 6) and added into set \mathcal{S}^i (line 7). Then, we remove the vehicle position i^* from \mathcal{V} (line 8) to make sure vehicle position i^* and the corresponding paths will never be selected in the following rounds.

d) *Finding the best assignment (line 12)*: From the collection of (\mathcal{S}, τ) pairs, \mathcal{M} (line 10), we pick the one that maximizes $\hat{H}(\mathcal{S}^i, \tau^i)$ as the output \mathcal{S} with the corresponding τ denoted by τ^* (line 12).

Oracle Design: We use an oracle \mathcal{O} to calculate the value of $H(\mathcal{S}, \tau)$. The oracle uses a sampling based method to approximate $H(\mathcal{S}, \tau)$ [2].

It has been shown in [2] that Algorithm 1 generates an assignment that has the bounded approximation performance of the optimal assignment.

IV. EVALUATION

In this section, we report the results from empirical studies evaluating the proposed risk-aware perception, planning, and assignment framework. We start by describing the experimental setup and then describe the results.

Setup. We use the AirSim [20] simulator as it offers photo-realistic images along with the semantically segmented ground truth in various pre-defined environments. We use the *CityEnviron* environment which contains city-like and suburban landscapes. We collect the downward-looking aerial images at an altitude value of 200m in the simulator. The dataset thus generated has 480 images, which are divided in the ratio of 10:3:3 as training, validation, and testing data for the Bayesian SegNet [21]. We use the PyTorch implementation of a basic version of this network [22] and add dropout layers in accordance with the original architecture. In our ground truth, we reduce the ground truth to 12 classes consistent with the original model [21].

Due to limitations in the simulator, roads and grass patches are indistinguishable in the ground truth as depicted in Figure 3. We keep such images limited to the test dataset. Hence, the grass patches act as unknown objects (similar to out-of-distribution) to the model and provide interesting observations on the uncertainty in prediction for such objects.

We produce 20 outputs (or *stochastic samples*) for each image using the trained Bayesian SegNet. For each pixel, the predicted label is found as the most frequent label among all of the most likely labels for each pixel, i.e., $l_x = \text{Mode}_{\text{samples}}(\arg \max_c P(c|x))$ where x is the pixel input to the model, c is the class label, and Mode stands for the statistical mode. Uncertainty for each pixel is defined as the average variance in the probabilities of labels i.e. $\text{uncertainty}(x) = \frac{1}{c} \sum_c \text{Var}(P(c|x))$. The cost for each

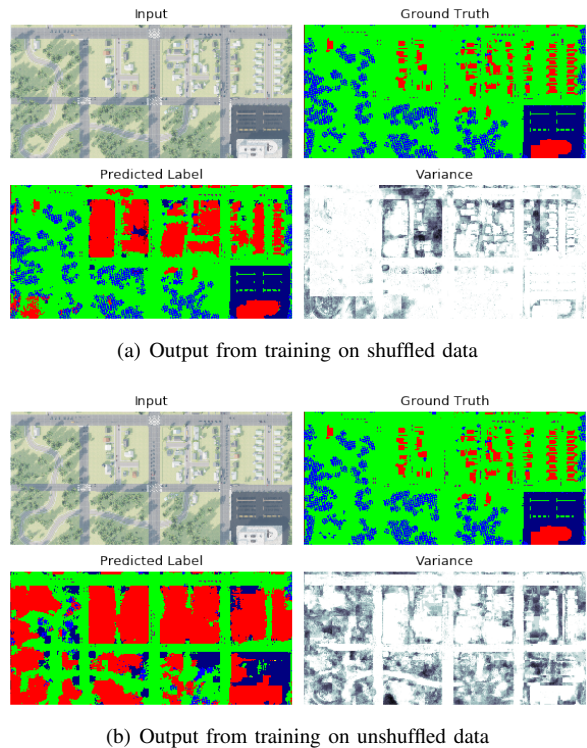


Fig. 3. Difference in variance due to data distribution.

pixel is calculated using these two quantities, which is further passed to the path planning algorithm (Sec. III-B). To provide an orthographic view in the results, we make predictions over partially overlapping images with perspective projections. Then we combine subsets around the center of the images.

Results. a) *Out-of-Distribution Data.* The Bayesian SegNet model used in our path planning algorithm has dissimilar data distribution for training and testing. The effect of this is shown in Figure 3. When the data distributions are the same across training and test dataset (shuffled data), the model is able to make predictions with very low variance (implied by brighter shades). However, in the case of dissimilar distribution, some part of the image acts as unseen data and thus the model prediction has high variance. This highlights the importance of uncertainty. For example, the prediction is precise for already observed data and thus the decision involving such input would be less risky, whereas the opposite is true for unobserved data.

b) *Average Cross-Entropy on the Test Dataset.* In order to understand the correlation between the quality of prediction and the training data, we look at the cross-entropy of the model prediction (over unshuffled data) averaged over the 20 samples and the number of pixels in the training data for each class in Figure 4. The model performs comparatively well in identifying obstacles like buildings. Objects which are rarely observed in the ground truth have very high cross-entropy. The lack of difference between the vegetation and the road in the ground truth also affects the performance of relevant classes. As expected, high cross-entropy is observed for classes where the pixel count is low.

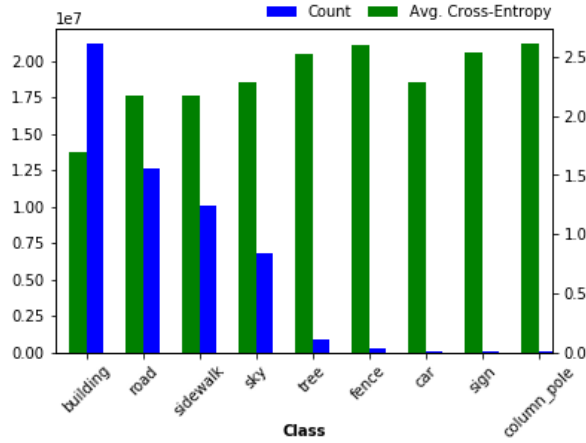


Fig. 4. Average Cross-Entropy for each class/label.

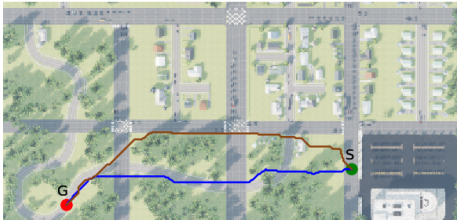


Fig. 5. Effect of λ on path planning. Smaller λ ($=10$) gives a shorter path with high uncertainty; larger λ ($=50$) gives a longer path with low uncertainty.

c) Risk-Aware Path Planning and Assignment.

The value of λ decides the risk-awareness of path planning. This effect is shown in Figure 5 where the cost is 1 for the navigable classes and 3 for non-navigable classes (except for the tree class, where it is 2). For a small value of λ , the algorithm plans a short path passing through the vegetation. For a high value of λ , this area is avoided due to the high uncertainty in the prediction for this part of the scene.

By varying the values of λ , we generate $K = 2$ candidate paths for each pair of start and demand location. We consider assigning $N = 3$ supply vehicles to $M = 2$ demand locations in this unknown environment.

Due to the imperfect semantic segmentation, the cost or efficiency of the path is a random variable. We show the efficiency distributions of the paths from the vehicles' start positions to the demand locations in Figure 6 and Figure 7.

We use Algorithm 1 to assign each vehicle a risk-aware path to a demand location. For example, in Figure 6-(a), vehicle v_2 is assigned path p_1 for demand location d_1 when the risk level is small, e.g., $\alpha = 0.01$. In contrast, as seen in Figure 7-(a), when the risk level is high, e.g., $\alpha = 1$, the assignment is more adventurous, and thus the path with a larger mean efficiency and a larger variance is selected. As shown in Figure 7-(b), vehicle v_2 switches to path p_2 for demand location d_2 , because the efficiency of this path has a larger mean. The path assignment changes follow the comparison of CVaR values, i.e., $\text{CVaR}_{0.01}[e(p_1)] > \text{CVaR}_{0.01}[e(p_2)]$ and $\text{CVaR}_1[e(p_1)] < \text{CVaR}_1[e(p_2)]$ with

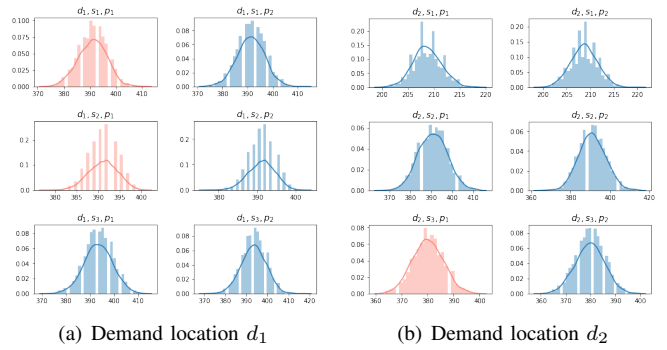


Fig. 6. Efficiency distributions of paths and the path assignment when $\alpha = 0.01$. The assigned path for each robot is marked in red.

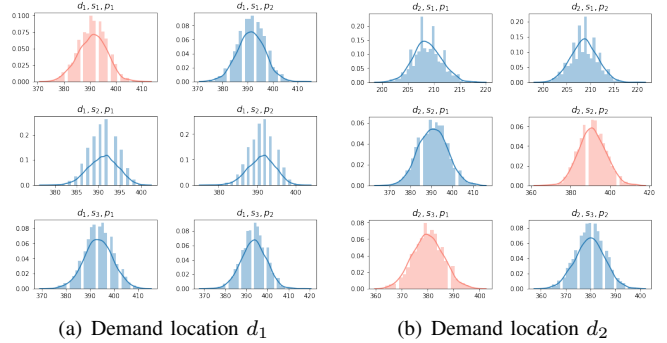


Fig. 7. Efficiency distributions of paths and the path assignment when $\alpha = 1$. The assigned path for each robot is marked in red.

$e(p_1)$ and $e(p_2)$ denoting the efficiency of path p_1 and p_2 , respectively. Thus, the risk level, α , provides the flexibility to trade off between risk and total efficiency (reward).

d) *Quantitative results.* In order to quantify the effect of λ , we define a metric called *surprise* for a path as the difference between the cost of the ground truth labels and predicted labels. We take 6 combinations of start and demand positions given in Figure 8 and find the average value of *surprise* as shown in Figure 9. Ideally, we expect the *surprise* to reduce with emphasis on variance. However, for large values of λ , even a few pixels with high variance may greatly increase the cost. In general, a higher value of λ may cause the robot to choose a longer path, which may increase the cost of traversal. However, sometimes the path may pass through only navigable regions, resulting in a small value of *surprise*. This causes the *surprise* to have a larger variance. Thus, the value of λ should be chosen after careful consideration of the variance in predictions and the range of the cost mapping.

We also plot the distribution of the total travel efficiency (Eq. 5) in Figure 10. \mathcal{S} is the path assignment selected by Algorithm 1. With small risk levels α , paths with low efficiencies (equivalently, low uncertainty) are mostly selected. This is because a small risk level indicates the assignment is conservative and only willing to take a small amount of risk. Thus, the path with lower efficiency and lower uncertainty is assigned to avoid the risk induced by the uncertainty. In contrast, when α is large, the assignment is riskier. In such a case, the paths with high efficiencies (equivalently, high

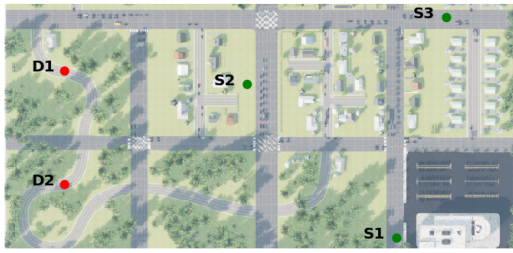


Fig. 8. Start and demand positions for *surprise* calculations.

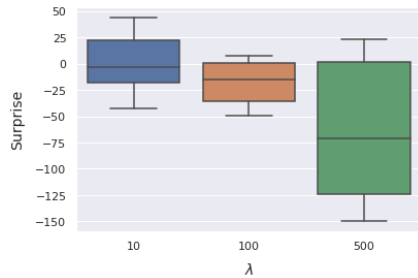


Fig. 9. *Surprise* vs λ .

uncertainty) are selected.

V. CONCLUSION

In this paper, we propose a risk-aware path planning and assignment framework for vehicles navigating in an unknown environment. We consider a scenario in which the information of the unknown environment is only available by an overhead, georeferenced image that is taken by an aerial robot. To deal with this challenge, we utilize Bayesian deep learning to learn the environment through semantic segmentation. Since the output of this segmentation is noisy, the cost of the vehicle traversing in the environment is uncertain. To deal with the cost uncertainty, we optimize some popular risk measures to generate and assign risk-aware paths for the vehicles. We use extensive simulation results to show the effectiveness of our risk-aware strategy.

Future work will focus on the online coordination of aerial and ground vehicles to achieve long-term, real-time risk-aware navigation in unknown environments.

REFERENCES

- [1] M. Toubeh and P. Tokekar, "Risk-aware planning by extracting uncertainty from deep learning-based perception," in *AAAI 2018 Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*, 2018, pp. 82–87.
- [2] L. Zhou and P. Tokekar, "An approximation algorithm for risk-averse submodular optimization," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2018, pp. 144–159.
- [3] Y. Gal, "Uncertainty in deep learning," Ph.D. dissertation, University of Cambridge, January 2017.
- [4] H. Wang, X. Shi, and D.-Y. Yeung, "Natural-parameter networks: A class of probabilistic neural networks," 2016.
- [5] J. Gast and S. Roth, "Lightweight probabilistic deep networks," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018. [Online]. Available: <http://dx.doi.org/10.1109/CVPR.2018.00355>
- [6] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*, 2016, pp. 1050–1059.

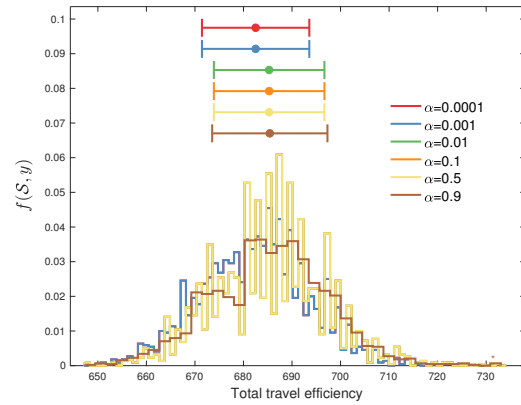


Fig. 10. Distribution of the total travel efficiency $f(S, y)$ by Algorithm 1.

- [7] A. Loquercio, M. Segù, and D. Scaramuzza, "A general framework for uncertainty estimation in deep learning," *IEEE Robotics and Automation Letters*, p. 11, 2020. [Online]. Available: <http://dx.doi.org/10.1109/LRA.2020.2974682>
- [8] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research*. Springer, 2020, pp. 75–84.
- [9] S. I. Marcus, E. Fernández-Gaucherand, D. Hernández-Hernández, S. Coraluppi, and P. Fard, "Risk sensitive markov decision processes," in *Systems and control in the twenty-first century*. Springer, 1997, pp. 263–279.
- [10] J. J. Chung, A. J. Smith, R. Skeelee, and G. A. Hollinger, "Risk-aware graph search with dynamic edge cost discovery," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 182–195, 2019.
- [11] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [12] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," 1998.
- [13] A. Pierson, C.-I. Vasile, A. Gandhi, W. Schwarting, S. Karaman, and D. Rus, "Dynamic risk density for autonomous navigation in cluttered environments without object detection," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5807–5814.
- [14] Y. Kantaros and G. J. Pappas, "Optimal temporal logic planning for multi-robot systems in uncertain semantic maps," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4127–4132.
- [15] J. Fu, N. Atanasov, U. Topcu, and G. J. Pappas, "Optimal temporal logic planning in probabilistic semantic maps," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3690–3697.
- [16] K. Harikumar, J. Senthilnath, and S. Sundaram, "Multi-uav oxyrrhis marina-inspired search and dynamic formation control for forest fire-fighting," *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 2, pp. 863–873, 2018.
- [17] E. Ackerman and E. Strickland, "Medical delivery drones take flight in east africa," *IEEE Spectrum*, vol. 55, no. 1, pp. 34–35, 2018.
- [18] A. Prorok, "Redundant robot assignment on graphs with uncertain edge costs," in *Distributed Autonomous Robotic Systems*. Springer, 2019, pp. 313–327.
- [19] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, "An analysis of approximations for maximizing submodular set functions—ii," in *Polyhedral combinatorics*. Springer, 1978, pp. 73–87.
- [20] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and service robotics*. Springer, 2018, pp. 621–635.
- [21] A. Kendall, V. Badrinarayanan, and R. Cipolla, "Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding," *arXiv preprint arXiv:1511.02680*, 2015.
- [22] M. P. Shah, "Semantic segmentation architectures implemented in pytorch." <https://github.com/meetshah1995/pytorch-semseg>, 2017.