Risk-Aware Path Planning and Assignment with Uncertainty Extraction from Deep Learning

Maymoonah Toubeh†1, Lifeng Zhou†1, Pratap Tokekar2

***Abstract*— We propose a risk-aware framework for multi-robot task 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 is typically noisy, so 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 account for risk at the perception level; and (ii) use a risk-theoretical measure, CVaR, at the planning and assignment level. The pipeline is theoretically established, then empirically analyzed through simulations. We find that accounting for risk at both levels produces quantifiably safer paths and assignments.**

I. Introduction

Many scenarios still exist where the environment that autonomous vehicles must navigate in is unknown to them. Such cases include search and rescue, space exploration, and military, among many others. Consider a disaster response scenario where ground vehicles need to supply resources at specific demand locations as soon as possible. In such settings, prior GPS or satellite maps of the environment may no longer be valid. The motivating setup for our work is an assisted planning and assignment problem involving multiple ground vehicles to be given risk-aware paths to their assigned tasks. It can be useful to employ multiple robots in a search and rescue operation where it is critical to minimize risks and maximize the utility of the vehicles as they can be helpful.

We consider finding paths for multiple vehicles to serve multiple demand locations. 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. Build 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. By utilizing CVaR, 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.

II. 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, $\mathcal{V} = \{v_1, \cdots, v_N\}$ and M demand locations, $\mathcal{D} = \{d_1, \cdots, d_M\}$ in the environment. The environment is represented by an overhead, georeferenced, RGB image as shown in Figure. 1. The goal is to find offline paths for each vehicle such that they collectively serve all the demands using navigation cost 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 utilize a measure, CVaR_α, that explicitly takes into account the risk associated with bad scenarios [4]. Specifically, CVaR_α measures the expectation of a random variable in the 100α–percentile worst scenarios. Here, $0 < \alpha \leq 1$ is a user-defined risk-level, which 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.

We are motivated by tasks that are urgent and time-critical, such as fighting fires [5] and delivering medical supplies in emergencies [6]. When the number of vehicles is more than the demands, assigning multiple redundant vehicles to demands helps counter the effect of uncertainty [7]. 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 by taking the reciprocal of the travel cost. Specifically, we use the travel *efficiency*, the reciprocal of travel cost, as the measure. The overall travel efficiency,

[1]The authors are with the Department of Electrical & Computer Engineering, Virginia Tech, Blacksburg, VA 24061, USA. {may93, lfzhou}@vt.edu.

[2]The author is with the Department of Computer Science, University of Maryland, College Park, MD 20742, USA (email: tokekar@umd.edu).

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

†Both authors contributed equally.

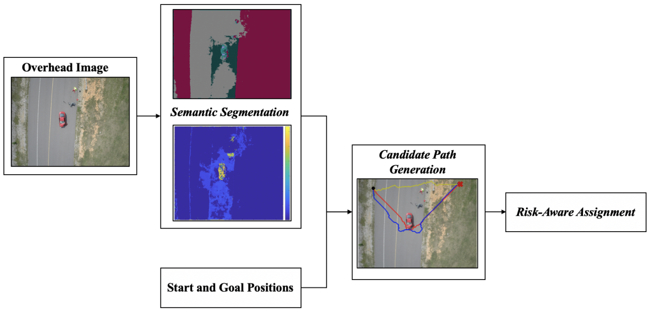

Fig. 1. The breakdown of the algorithm's parts. Given an overhead image input, the algorithm provides a semantic segmentation (and uncertainty map), then generates candidate paths, and finally performs the risk-aware assignment of vehicles to paths.

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)

$$\max_{\mathcal{S}\subseteq\mathcal{X}} \text{CVaR}_\alpha[f(\mathcal{S}, y)] \tag{1}$$

where $\mathcal{S}$ is a path set for vehicles with "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. Algorithm and Analysis

The algorithm we propose consists of three main parts: semantic segmentation, candidate path generation, and risk-ware assignment. The overall pipeline of the framework and its parts are shown in Figure 1.

The inputs to the algorithm are a single overhead aerial two-dimensional image, vehicles' start positions and 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 input 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 (Problem 1), for risk-aware path assignment.

IV. Simulations

We consider assigning $N = 3$ supply vehicles to $M = 2$ demand locations in a 2D environment, which is represented by an overhead image (Fig. 1). By path generation technique, each vehicle has $K = 3$ candidate paths from its start position to each demand location.

Due to the imperfectness of semantic segmentation, the efficiency of the path is a random variable. We show the efficiency distributions of the paths from vehicles' start positions to demand locations in Figure 2 and Figure 3.

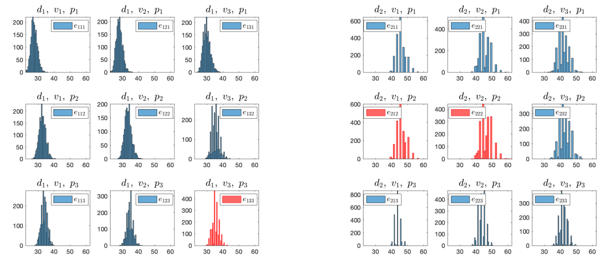

(a) Demand location d_1 (b) Demand location d_2

Fig. 2. Efficiency distributions of paths and the path assignment when $\alpha = 0.01$. The assigned path for each robot is marked in red.

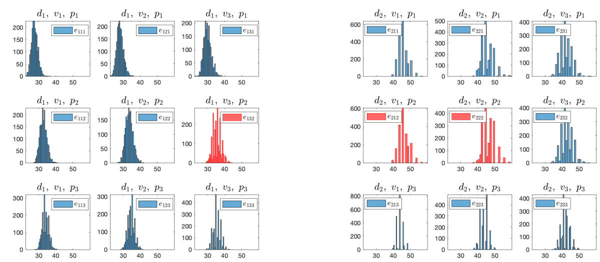

(a) Demand location d_1 (b) Demand location d_2

Fig. 3. Efficiency distributions of paths and the path assignment when $\alpha = 1$. The assigned path for each robot is marked in red.

We use SGA [2] to assign each vehicle a path to a demand location. For example, in Figure 2-(a), vehicle's v_3 is assigned path p_3 for demand location d_1. In contrast, when the risk level is high, e.g., $\alpha = 1$, the assignment is more adventurous, since the paths with a larger mean and a larger variance are selected. As shown in Figure 3-(a), vehicle's v_3 switches to path p_2 for demand location d_1. Thus, the risk level, α, provides a user with the flexibility to trade off between risk and total efficiency (reward).

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 (WAFR)*, 2018, in press.

[3] Y. Gal, "Uncertainty in deep learning," Ph.D. dissertation, University of Cambridge, January 2017.

[4] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.

[5] K. Harikumar, J. Senthilnath, and S. Sundaram, "Multi-uav oxyrrhis marina-inspired search and dynamic formation control for forest firefighting," *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 2, pp. 863–873, 2018.

[6] E. Ackerman and E. Strickland, "Medical delivery drones take flight in east africa," *IEEE Spectrum*, vol. 55, no. 1, pp. 34–35, 2018.

[7] A. Prorok, "Redundant robot assignment on graphs with uncertain edge costs," in *Distributed Autonomous Robotic Systems*. Springer, 2019, pp. 313–327.