CEAS-GNC-2022-045

CEAS EuroGNC 2022
"Conference on Guidance, Navigation and Control"
3-5 May 2022 @ Technische Universität Berlin, Germany

Energy-aware Route Planning for Drone Delivery Systems

Ruifan Liu PhD candidate, Cranfield University, School of Aerospace, Transport and Manufacturing, MK43 0AL, Cranfield, UK. ruifan.liu@cranfield.ac.uk

Hyo-sang Shin Professor, Cranfield University, School of Aerospace, Transport and Manufacturing, MK43 0AL, Cranfield, UK. h.shin@cranfield.ac.uk

Minguk Seo Research fellow, Cranfield University, School of Aerospace, Transport and Manufacturing, MK43 0AL, Cranfield, UK. minguk.seo@cranfield.ac.uk

Antonios Tsourdos Professor, Cranfield University, School of Aerospace, Transport and Manufacturing, MK43 0AL, Cranfield, UK. a.tsourdos@cranfield.ac.uk

ABSTRACT

This paper investigates a stochastic route optimization problem for drone delivery systems, with a focus on managing energy risk while optimizing service qualities. In drone delivery, couriers, i.e., small commercial unmanned aerial vehicles (UAVs), face a very limited power and load capacity, which needs to be carefully settled throughout the route optimization process. Additionally, due to its wind-sensitive property, the energy cost of a UAV is fiercely affected by the airflow and might spread across a wide range. To optimize the route while managing its energy use properly, we propose an Energy-aware Planning Framework (EaPF), which is embedded with an energy prediction model and adapts to any existing route optimizer. Instead of sampling using either analytical or numerical formulations, the proposed energy prediction model directly yields the distribution of energy consumption via a Mixture Density Network (MDN). Based on the statistical model, an energy risk criterion and an objective function in the form of expectation are both developed for route optimization. In addition, an event-driven routing simulator is devised with the aim to accommodate the energy model and incorporate it with an optimizer. Finally, we implement the proposed planning framework to the simulation test on a medical delivery mission, demonstrating its superiority in terms of energy risk management and solution qualities.

Keywords: Drone Delivery; Route Optimization; Energy Risk Management; Mixture Density Network

1 Introduction

With the rapid development of UAV technology, UAVs are now increasingly being adopted by the logistics industry, with intensive experimental tests being conducted to open up drone delivery techniques to real-world deployment [1]. The introduction of small commercial UAVs into logistics markets is expected to bring substantial economic and environmental benefits. Their instinctive flexibility and efficiency also provide an alternative solution for time-critical deliveries and the 'last mile' problems.

Despite its advantages, the introduction of a new courier always implies the need for new logistical techniques. Delivering goods with small UAVs faces very strict battery and load limitations compared to conventional ground transportation. To that end, renewable facilities are set in practical deployments to extend the range of services and increase delivery efficiency. In that case, planning routes in an energy-aware manner tends to maximize the quality of services while being eco-friendly. Furthermore, in the aerial logistic system, energy risk management is very important for mitigating the potential loss of UAVs. It involves the process of pinpointing, evaluating, and prioritizing risks associated with uncertainties in UAV's energy consumption. Controlling energy risk can enhance the stable operation and reliability of the delivery service.

Intensive research has been conducted on the subject of drone delivery planning, in which a certain objective is pursued while taking into account the energy constraint. Existing planning methods, however, are often based on simple, abstract approximations of the underlying system, such as using a pre-defined flight duration to imply the energy budget. This may benefit from computational tractability and allows for quick analysis, but comes at the cost of ignoring real-world complexities. The wind, for example, might cause the battery to drain quicker than expected. Intending to produce more reliable solutions, a more sophisticated problem model is worth studying.

To this end, this paper gives an in-depth discussion of UAV energy modelling and addresses route planning for UAV deliveries as a stochastic optimization problem within an energy-aware planning framework. Specifically, we first identify five basic components of the green aerial logistics system and its route optimization problem. Based on the conceptional definition, an energy-aware planning framework is then proposed to address uncertainties in drone delivery problems, which incorporates a statistical model of energy consumption and a heuristic optimizer.

To summarize, the contribution of this paper is listed as follows:

- Via statistical modelling, we approximate UAV's energy consumption using a Mixture Density Network (MDN) in consideration of winds. This allows for a view into the distribution of energy costs without the need for sampling. Furthermore, the parameters of the energy consumption model could be updated online with the aim of reducing the discrepancy between the planner model and the real world.
- An energy risk metric is introduced based on the MDN model of drone energy cost. The metric is used to assess the risk for a route and is calculated analytically, also avoiding the tedious sampling process.
- Combined with simulation techniques, an Energy-aware Planning Framework (EaPF) is suggested to address the energy risk issue for stochastic route optimization problems. It integrates the above-mentioned energy model and forms a general scheme for energy-aware simulation-based route planning. This framework adapts to any route optimizer and provides more reliable solutions for energy risk-sensitive applications.

The rest of the paper is organized as follows. Section II provides an overview of previous works with regard to energy-aware planning and simulation-based optimization. Section III gives a conceptional definition of green aerial logistics systems and then introduces the route optimization problem for the system. An energy-ware planning framework for the logistic systems is proposed in IV. This is followed in Section V by the energy cost statistical modelling and the criterion of energy risk assessment. Section VI and Section VII respectively present the details about the implementation of the routing simulator and heuristic optimizer. A case study is illustrated in Section VIII. Section IX concludes this study and outlines a brief plan for future research.

2 Related Works

2.1 Energy-aware Planning

Regarding the limited capacity of batteries carried by small UAVs, energy consumption is an important factor to consider in planning, routing and scheduling. To predict the energy demand, a variety of energy consumption models have been established, either based on the dynamics of UAVs or via fitting the practical experiment dataset. A rotary-wing vehicle's power model is investigated in [2], considering both hovering and forward flight. Based on this study, Kirschstein *et al.* [3] propose an energy consumption model focusing on the drone delivery task, where four flight patterns are distinguished: take-off, level flight, hovering and landing. More aerodynamic-based drone energy models can also be found in [4–6]. In addition to analytical models, the energy model has recently been established on regression techniques and machine learning. Through partitioning the flight of the UAV into segments of different movements, Prasetia *et al.* [7] predict the energy consumption via collecting data, preprocessing data, and regressing models for these movements separately. Choudhry *et al.* [8] develop a deep energy model with temporal convolutional networks, which is trained on a real-world dataset [9] and shows a 29% improvement in power predictions.

Further, the limited battery capacity of the small commercial UAV always poses a significant risk to its operation that the energy might run out before the drone lands at the next destination. Quantifying the risk of using up batteries is very important for path planning and task routing. Risk assessment for aerial vehicles in terms of energy is relatively a new research domain and it stems from the difficulty in identifying and qualifying various factors that affect UAV energy consumption. Inspired by the risk assessment metric proposed in previous work, a Conditional-Value-at-Risk (CVaR) is employed to qualify the risk based on the Monte Carlo forward simulation of the energy distribution [8]. Risk assessment also facilitates planning in a risk-aware manner where the risk of each route is quantified to meet specific safety regulations. For instance, the possibility of not completing a plan is studied for a stochastic orienteering problem where the travel time for each edge follows a pre-known gamma distribution [10]. In [10], a risk-aware criterion with respect to the probability of violating deadlines is considered in addition to the reward-collecting objective, which is then used in a local research method. For the battery-bounded drone delivery, energy prediction is also used to extend the range of delivery [11], in which energy cost is regarded as a block box and trained via comparing the set of predicted destinations.

2.2 Simulation-based Optimization for Routing Problems

Simulation-based optimization extends deterministic techniques to deal with stochastic optimization problems. One typical application, referred to as the term *sim-heuristics*, is combined with metaheuristic methods. By integrating simulations into an existing metaheuristic framework, sim-heuristic allows modellers to deal with uncertainties in a natural way. Moreover, sim-heuristics also facilitate the introduction of risks and reliability analysis/assessment of solutions in stochastic combinatorial problems. A general scheme of sim-heuristic is given in [12] for solving stochastic combinatorial optimization problems. For its specific application on routing problems, Mazzuco *et al.* [13] investigates the problem of daily good deliveries and presents the first outcomes of a simulation-based optimization approach for vehicle routing and transport scheduling. In terms of the stochastic character of real-world VRPs, Mayer *et al.* [14] offers a Rich VRP simulation model based on the discrete event simulation (DES) workflow. The simulation model is then used to evaluate the solution candidates for metaheuristics.

Besides combinations with metaheuristics, researchers also expect the combined use of system simulation and machine learning. A trial model of applying reinforcement learning (RL) in the AnyLogic model is investigated in [15]. They demonstrate how to implement RL in the AnyLogic simulation model using the Pathmind Library. The combination of DES and deep RL has also been studied on large logistics networks [16], for which the authors first employ meta-heuristic approaches and then replace them with a reinforcement learning agent to find the most promising action sets. Event-driven simulation is also beneficial in the domain of multi-agent reinforcement learning (MARL). Due to the curse of dimensionality, researchers tend to consider sparse actions or macro actions in MARL, which are probably asynchronous especially when agents act in a decentralized manner. To address this asynchronicity, a possible way is to model multi-agent decision processes as event-driven processes [17]. It makes the DES a promising tool to simulate the multi-agent environment and train optimal policies.

3 Green Aerial Logistics Systems

3.1 Motivation

Suppose a mission scenario where there is a variety of delivery requests $j \in \mathscr{T}$ and one fleet of UAVs $a \in \mathscr{A}$ is sent to accomplish these requests maximizing customer services qualities without violating their energy budget. It is identified as a route optimization problem for specified logistics systems where the battery-driven UAVs serve as logistic carriers. Different from ground vehicles, these aerial vehicles always face a highly strict energy constraint. If an energy prediction model is accessible, planning with the awareness of energy tends to increase the quality of routes by reducing risks and extending the flight range. However, in many real-world settings, the energy cost of UAVs is hard to predict and has shown to be stochastic due to winds. Modelling it as a statistic distribution is much more reasonable than making single-output predictions. At the same time, planning in an unpredictable environment is required to be robust and dynamic. The robustness indicates the ability to systematically mitigate the sensitivity of the policy to ambiguity in the underlying transition probabilities, while the dynamic property allows the decision-maker to efficiently compute a good overall strategy by succinctly encoding the evolving information state. To conclude, how to establish the above-described energy model and incorporate it into a robust and dynamic planner is the main motivation behind this paper.

3.2 Green Aerial Logistics System

We build the green aerial logistics system (GALS) based on the architecture of the green logistics system (GLS) from [18]. Similar to GLS, there are also five essential components for GALS, i.e., transportation network, vehicles, logistics request, renewable generations, and depots.

1) *Aerial Transportation Network*: We define the transportation network on an directed graph $\mathscr{G} = \langle V, E \rangle$. Each vertex $v \in V$ stands for a point of interest, which could be a task request, a recharging facility or the depot. Each edge $e \in E$ represents the path segment linking every pair of vertexes. Vertexes are connected with the corresponding logistics components and each edge is associated with a series of way-points, including the take-off, landing locations and predefined mid-way route points.

2) *Autonomous Aerial Vehicles*: The UAVs executing in the system are expected to be heterogeneous, denoted as $a \in \mathscr{A}$. Each of them is specified with an independent battery capacity B_a and a State of Charge (SoC) C_a. UAVs are normally sent out from the depot with fully charged batteries and are expected to get recharged before the battery uses up. Moreover, it is assumed that the logistics capacity of UAVs is 1, i.e., it only processes one request at one time.

3) *Logistics Requests*: Customer requests are denoted as $q \in \mathscr{Q}$, which consist of pickup tasks $q_p \in \mathscr{Q}_p$ and delivery tasks $q_d \in \mathscr{Q}_d$. Each pair of pickup and delivery tasks is bound together and constrained by precedence relationship. They share same load weight w_q, task value r_q, severity index s_q, appearance time a_q, and execution deadline d_q.

4) *Renewable Generation*: Renewable facilities $g \in \mathscr{G}$ are deployed in GALS to extend the range of service and promote environmental friendliness. Batteries equipped with UAVs get recharged or

swapped in these locations. If recharged, the recharging rate is denoted as ρ_a. Additionally, we consider at most that Ω_g UAVs are allowed to be charged at point g at the same time.

5) *Depots*: Depots $d \in \mathscr{D}$ refer to those places where UAVs take off and are recycled at the end of the route. In GALS, the depots also serve to charge UAVs.

3.3 Route Optimization in GALS

To make sure the GALS system functions well, an online mission planner is set to get logistics requests allocated among the UAV fleet and also to plan an optimal route for each team member in terms of certain objectives. It is characterised as a dynamic and stochastic green vehicle routing problem (SD-GVRP), with optimal solution indicating maximizing a predefined objective function while satisfying the battery budget and mission requirements.

3.3.1 Objective function

In logistics systems, optimizing the delivery flow indicates maximizing the rewards collected from requests, while minimizing the loss if failed to arrive at one task within its time window. To quantify rewards and losses, we denote the delivery request q using the following four parameters:

- delivery reward r_q: the reward obtained by finishing request q;
- severity index s_q: the penalty for failing to execute request q;
- appearance time a_q: the time when request q arrives;
- execution deadline d_q: the latest finish time for request q.

Also, for the ease of formulation, the status of UAV a at vertex j is described by

- visiting time t_j^a: the time when UAV a arrives at j;
- duration of stay dur_j^a: the duration of stay at j for a ;
- state of charge (SoC) c_j^a: the SoC of UAV a when it arrives at j.

For GALS, we consider practical objectives from four aspects: 1) maximizing the sum of collected rewards; 2) minimizing the losses caused by missing delivery requests; 3) minimizing the sum of delivery delay time; 4) minimizing the overall energy consumption.

Regarding a route $\mathscr{R}^a = [v_1, v_2, \cdots, v_n]$ for UAV a and using above notations, the objective function is formally given by

$$\text{maximize } f(\mathscr{R}^a) = w_r \sum_{j=1}^{n} r(v_j) - w_d \sum_{j=1}^{n} dla(v_j) - w_e \sum_{j=1}^{n-1} e(v_j, v_{j+1}) \tag{1}$$

where w_r, w_d, and w_e represents weights for each components, $e(v_j, v_{j+1})$ indicates the energy cost flying from v_j to v_{j+1}, the reward function $r(v_j)$ and the time-delay function $dla(v_j)$ are defined as

$$r(v_j) = \begin{cases} 0 & \text{if } v_j \notin \mathscr{Q}_d \\ r_j & \text{if } v_j \in \mathscr{Q}_d \text{ and } t_j^a \le d_j \\ -s_j & \text{if } v_j \in \mathscr{Q}_d \text{ and } t_j^a > d_j \end{cases} \tag{2}$$

$$dla(v_j) = \begin{cases} 0 & \text{if } v_j \notin \mathscr{Q}_p \\ t_j^a - a_j & \text{if } v_j \in \mathscr{Q}_p \end{cases} \tag{3}$$

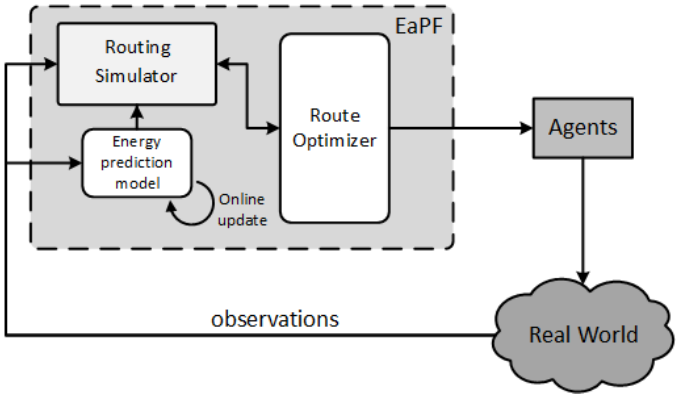

Fig. 1 Energy-aware Planning Framework.

3.3.2 Constraints

1) Route and precedence constraints

For logistics systems, the precedence relationship between pickup and delivery processes should be satisfied, which could be guaranteed by ensuring $t_{q_p} < t_{q_d}$ if q_p and q_d are for the same request, or by defining the connectivity of the transportation network. Additionally, every route is expected to start and terminate at the depot, i.e. $v_0, v_n \in D$, which could be the same one or not.

2) Energy constraint

Distinct from ground vehicles, small commercial UAVs always come with harsh restrictions on battery and payload capabilities. Thus we impose a capacity constraint that the UAV only processes one request at one time, and for flight safety, the SoC of UAVs must be guaranteed positive at any point of the route, which means $c_j^a \geq 0, \forall j \in \mathscr{R}, \forall a \in \mathscr{A}$.

4 Energy-aware Planning Framework

To handle the above route optimization for the GALS, we propose an Energy-aware Planning Framework (EaPF), depicted in Fig.1. The main goal of EaPF is to produce robust routes and manage energy risks in the face of poorly known uncertainties. In EaPF, an event-driven routing simulator serves as the simulated environment, where we plug in an energy consumption model to yield stochastic state transitions. To reduce the discrepancies between the planning model and real-world dynamics, the energy prediction model is derived from realistic data and online refined using empirical observations. Also, the energy prediction model is used to quantify the energy risk, assessing the reliability of the route. By means of filtering the reliable solution candidates using stochastic simulators, a route optimizer is expected to produce robust routes via Simulation-based Optimization.

In this paper, we leverage a heuristic method as the route optimizer in EaPF. The heuristic algorithm performs in a 'sim-heuristic' way via incorporating the routing simulator. Other planning methods, e.g., deep reinforcement learning (DRL), could alternatively be implemented in EaPF. In this case, the simulator would also serve as the training environment. Observations are extracted from the simulated environment and passed into the DRL network. Then the agent randomly chooses the next node following the output probability from the policy network while states in the event-driven simulator get updated accordingly. During the process, rewards are collected and used to train the networks via DRL algorithms.

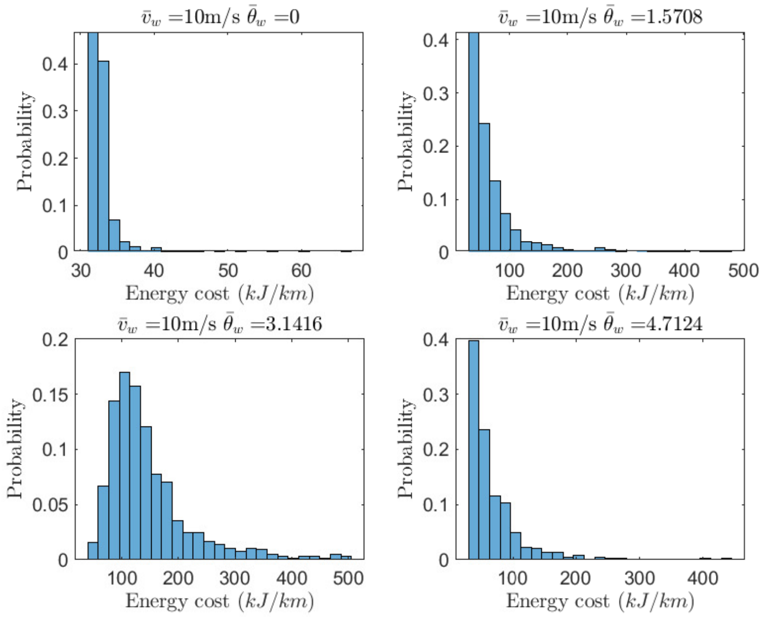

Fig. 2 Histograms of drone energy cost under winds. Four cases are simulated in presence of winds from different directions: $0deg$ **(tailwind),** $90deg$**,** $180deg$ **(headwind),** $270deg$**. The wind velocity** v_w **and direction** θ_w **follow:** $v_w \sim N(10m/s, 1.5m/s)$**,** $\theta_w \sim N(\bar{\theta_w}, 30deg)$

5 Energy Modelling via Mixture Density Network (MDN)

5.1 Energy Prediction Model

Researchers have adopted regression techniques to predict the energy cost, the goal of which is always to find a single-value prediction with maximum likelihood or a Gaussian conditional distribution $p(t|x)$. However, under varying wind conditions, the UAV energy cost for a given path is distributed across a large range and is shown to be non-Gaussian (Fig.2).

In this study, we use a Mixture Density Network (MDN) to estimate the energy cost for a more general formalism and more precise prediction. The MDN approximates conditional density functions by mixing several component densities with mixing coefficients. Both densities and coefficients are flexible functions of the input vector x. Here components are chosen as Gaussian Mixtures (GM):

$$p(t|x) = \sum_{k=1}^{K} \pi_k(x) \mathcal{N}(t|\mu_k(x), \sigma_k^2(x)) \tag{4}$$

Since we mainly focus on the effect of wind on UAV energy consumption, the input vector is set as $x = [\mu(v_w), \sigma^2(v_w), \mu(\theta_w), \sigma^2(\theta_w)]$, containing the mean value and std value of both wind speed and direction.

The parameters in the above Gaussian mixtures are governed by the outputs of a conventional neural network that takes x as its input. If there are K components in the mixture model, the network will have 3*K outputs, representing mixing coefficients $\pi_k(x)$, means $\mu_k(x)$ and variances $\sigma_k^2(x)$ respectively. Notably, since the sum of mixing coefficients must be 1, an extra SoftMax layer is imposed on coefficient outputs, and an exponential operation is added to the variance outputs to satisfy the positive variances.

5.1.1 Parameter Update

Ahead of the mission, the MDN network forms the prior distribution of energy cost $p(t|\theta_i)$, which is parameterized by θ_i and trained using historical and simulated data. Once practical flights launch, the network is then modified on board with latest empirical results (x_i, y_i), producing a posterior distribution $p(t|\theta_i, x_i, y_i)$.

In terms of the parameter update, the trainable parameters θ of this density network, including weights and biases in the neural network, is set by maximizing the likelihood, or equivalently by minimizing a loss function defined as the negative logarithm of the likelihood. The loss function takes the form

$$L(\theta) = -\sum_{n=1}^{N} \ln \left\{ \sum_{k=1}^{K} \pi_k(x_n, \theta) N\Big(t_n | \mu_k(x_n, \theta), \sigma_k^2(x_n, \theta)\Big) \right\} \tag{5}$$

To minimize the loss function, we calculate the derivatives of the loss $L(\theta)$ with respect to parameters θ and update the network via the standard back-propagation procedure. Unlike the fixed energy prediction model, the MDN is updated continuously, beneficial to long-term performance in real-world applications.

5.1.2 Model Test

The MDN is set as a simple dense network, consisting of an input layer, an output layer, and two full-connected layers with 256 neurons each. We apply the Adam optimizer with a learning rate of 0.001. The component number of the Gaussian mixture model is set as $K = 5$. The dataset used for training comprises randomly generated wind parameters and the energy cost calculated via the deep energy model from [8], which is based on an energy usage dataset collected from 100-hour real-world flights [9]. Here, we adopt the application case 1 in [8] as the training environment, where the test site is surrounded by ten buildings ranging in heights from 5 to 35 meters and the wind field is modelled using computational fluid dynamics (CFD). The learning cures with regard to the likelihood of the training dataset and validation dataset are shown in Fig. 3.

One instance is extracted for further illustration, where the inlet wind condition is sampled as $[\mu(v_w), \sigma^2(v_w), \mu(\theta_w), \sigma^2(\theta_w)] = [-169.27°, 47.7°, 9.06m/s, 1.32m/s]$. A series of 1000 Monte Carlo simulations are also conducted using the above-mentioned deep energy model for comparison. The probability density function via well-trained MDN and a histogram of numerical Monte Carlo results are both depicted in Fig. 4. As seen from the figure, the MDN model achieves a very good approximation to the histogram of Monte Carlo simulations.

5.2 Energy Risk Assessment

5.2.1 Risk Metric

For UAVs' route planning, it is vital to make sure that the vehicle returns to a recharging point before its battery uses up. Considering uncertainties in the operation environment, using the mean value or the expected value to constrain the path might result in a relatively poor plan. Thus we employ a chance-constrained strategy that guarantees the energy constraint within a certain risk threshold. Specifically, a route candidate $r = [v_1, v_2 \ldots, v_k]$ is denoted as valid only if it satisfies the following change-constrained metric:

$$\mathscr{R}^* = \Big\{ r : p_\theta\Big(C(r) < B_r\Big) \geq 1 - \varepsilon) \Big\} \tag{6}$$

where $p_\theta(\bullet)$ denotes the probability defined by uncertainty parameters θ, and B_r is UAV's remaining budget. If the agent chooses a route from $\mathscr{R}^*$, it can be guaranteed to finish the route within the allowable risk threshold specified by ε.

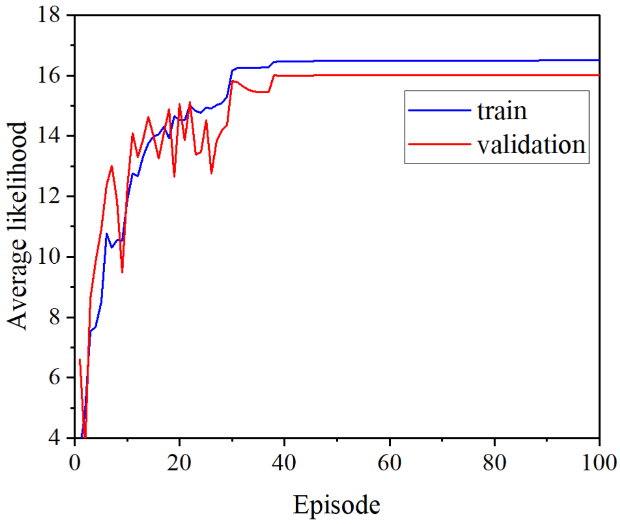

Fig. 3 Likelihood curves of the MDN for energy cost prediction

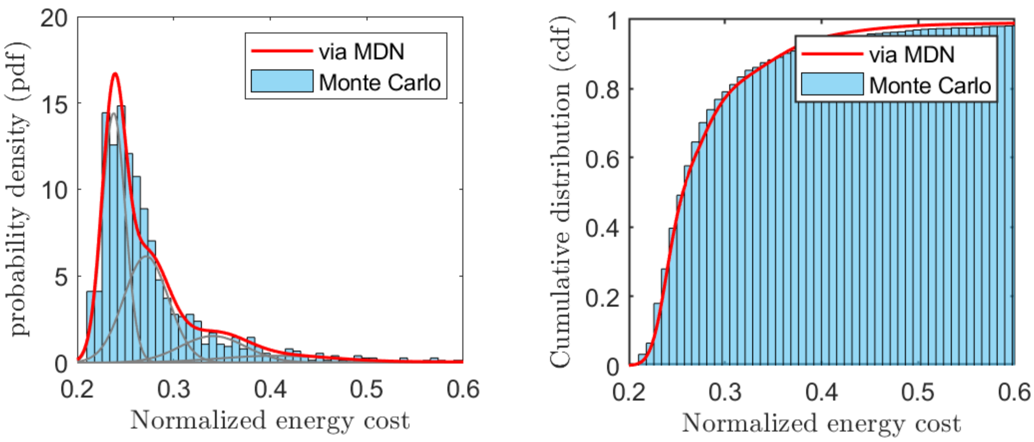

Fig. 4 An example illustration for energy prediction via MDN. Results obtained by MDN and Monte Carlo simulations are plotted in the form of probability density functions and cumulative distribution respectively.

5.2.2 *Joint Energy Distribution*

For flexibility and computational efficiency, energy cost functions are applied on each network edges rather than route solutions consisting of connected edges. Under the assumption of independent energy costs, the probability density of energy cost for flying over a mission route is formed by integrating each probability density function with regard to its path components. As a simple illustration, suppose there is a route from waypoint A to way point C passing through waypoint B, the probability function of the route is then given by

$$e_{ac} = e_{ab} + e_{bc} \sim \sum_i \sum_j \pi_i \phi_j \mathcal{N}(\mu_i + \mu_j, \sigma_i^2 + \sigma_j^2) \tag{7}$$

when energy costs for both path segments follow mixture Gaussian distributions:

$$e_{ab} \sim \sum_{i=1}^{n} \pi_i \mathcal{N}(\mu_i, \sigma_i^2), \quad e_{bc} \sim \sum_{j=1}^{m} \phi_j \mathcal{N}(\mu_j, \sigma_j^2) \tag{8}$$

5.2.3 Mixture Weight Pruning

The component number of the energy prediction model increases exponentially when integrating GMs of all path segments. To retain the number of components within a reasonable range, we resort to mixture reduction techniques, which attempt to reduce the number of components in a mixture through emerging or pruning. Here we primarily employ the emerging operation, which is inspired from [19].

The emerging operation replaces a pair of components in a mixture density with a single component of the same type. Here we adopt the KL divergence as the criterion for selecting the emerging pair to minimize the divergence between the emerged and original components. Formally, the KL divergence criterion $B(i,j)$ is calculated as

$$B(i,j) = w_i D_{KL}\Big(p(x;\mu_i,\sigma_i^2)||p(x;\mu_{ij},\sigma_{ij}^2)\Big) \tag{9}$$

$$+ w_j D_{KL}\Big(p(x;\mu_j,\sigma_j^2)||p(x;\mu_{ij},\sigma_{ij}^2)\Big) \tag{10}$$

where merged parameters are got by moment matching:

$$\mu_{ij} = \frac{w_i\mu_i + w_j\mu_j}{w_i + w_j} \tag{11}$$

$$\sigma_{ij}^2 = \frac{(\sigma_i^2+\mu_i^2)w_i + (\sigma_j^2+\mu_j^2)w_j}{w_i+w_j} - \mu_{ij}^2 \tag{12}$$

Merging operations are implemented in a greedy manner. The pair of components with the minimum global KL divergence is recursively selected to be emerged until reaching the target component number.

5.3 Objective Function in the Form of Expectation

Uncertain energy costs of drone couriers caused by airflow leads to stochastic route optimization for GALSs. As defined in Section3.3, the objective component regarding total energy cost is directly governed by the GM model of drone energy consumption, while the reward function and time delay penalty is decided by arriving times, which implicitly depend on energy cost through charging time if ground speed is assumed to be constant. Accordingly, the expected objective function with respect to the GM model of energy cost is formulated as

$$\text{maximize } E(f(\mathscr{R}^a) \mid \theta) = w_r \sum_{j=1}^{n} E(r(v_j) \mid \theta) - w_d \sum_{j=1}^{n} E(dla(v_j) \mid \theta) - w_e \sum_{j=1}^{n-1} E(e(v_j, v_{j+1}) \mid \theta) \tag{13}$$

where the expectation components are as follows

$$E(r(v_j) \mid \theta) = r_j p(t_j^a \leq d_j \mid \theta) - s_j p(t_j^a > d_j \mid \theta) \tag{14}$$

$$E(dla(v_j) \mid \theta) = \mu(t_j^a \mid \theta) - a_j \tag{15}$$

$$E(e(v_j, v_{j+1}) \mid \theta) = \mu(e(v_j, v_{j+1}) \mid \theta) \tag{16}$$

where $\mu(\bullet \mid \theta)$ indicates the expected value of the GM model with its parameters θ.

6 Event-driven Simulator for GALS

A Routing Simulator is developed based on the scheme of discrete event simulation (DES) to emulate the delivery process in real-world GALSs. In the GALS, most time increments produce no change to the system, implying the fixed time advance simulation is clearly time-wasting. Therefore, we resort to DES in which the simulation clock could skip to the next event time. The Routing Simulator implements the standard components of a DES, as presented in [20], and is then extended regarding specific logistics applications.

1) *System State*: The collection of state variables of GALS system entities (i.e., transportation network, vehicles, requests, renewable facilities and depots), which is used to describe the system at a particular time and forms the input data for route optimization.

2) *Simulation Clock*: A variable giving the simulated time.

3) *Initialization Routine*: Initialize the system entities model at time zero.

4) *Timing Routine*: Determine the next event from the event list and then advance the simulation clock to the time.

5) *Event Routine*: Update the system state when an event occurs, and gather the reward value from the state update.

6) *Library Routine*: A set of subprograms used to generate random observations from probabilism models. For GALS, we employ modular models to get stochastic energy updates and reward functions.

7) *Main program*: A subprogram that invokes the timing routine to determine the next event and then transfers control to the corresponding event routine to update the system state appropriately.

Based on the components listed above and the workflow presented in [20], we integrate the Simulator with a route optimization module, as depicted in Fig. 5. The optimization module takes the system state as input data and generates optimized routes via an optimizer. These routes compose the main behaviours of the system and always return its first element in the sequence. The optimization module is first invoked after the system is initialized, and might also be triggered by specific events or after a re-planning time interval. For example, when a breakdown event from a vehicle or a new customer request occurs, the optimization module might be re-invoked, redefining all the tour events in the event list. This trigger mechanism enables the implementation of the dynamic rolling horizon planning scheme and could be flexibly designed in terms of the trade-off between performance and computational expense.

7 Heuristic Optimizer

A typical and efficient heuristic method, the neighbourhood search (NS) algorithm, is employed in this paper as the route optimizer. It is implemented on two basic local-search operators: intra-route relocation and inter-route relocation, after which a recharging insertion operation is appended to finish the route consideration of the energy constraint.

Intra-Route relocation: The order of two pairs of pickup and delivery nodes are swapped for the same vehicle.

Inter-Route relocation: A pair of pickup and delivery nodes are removed from a vehicle and inserted into another vehicle.

Recharging insertion: Every time after applying either of the above operations, the utility of the route is calculated with consideration of the required recharging. Herein applies a quite simple recharging policy: vehicles whose battery is insufficient to cover both the next assigned request and the flight to recharge stations go to the closest station to get fully recharged.

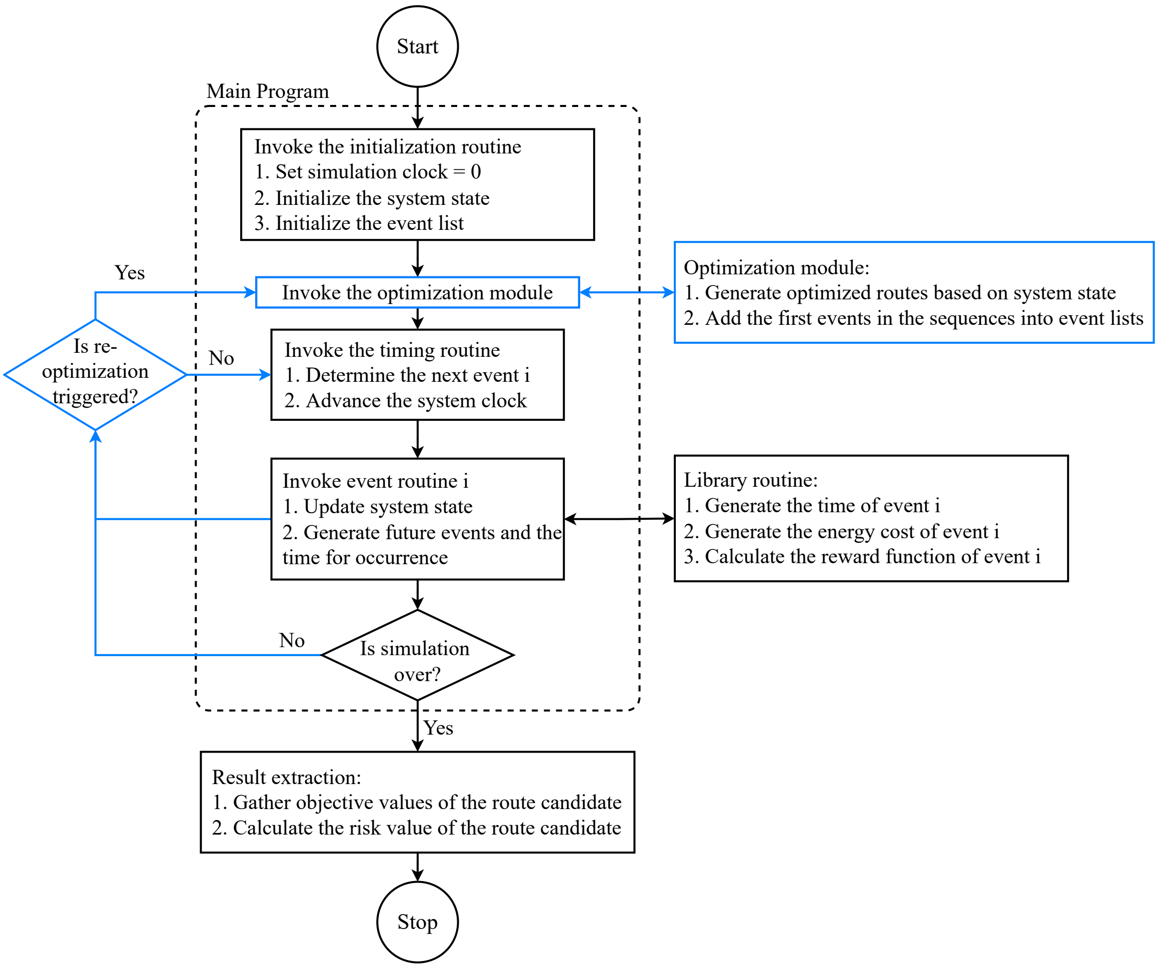

Fig. 5 The workflow of Discrete Event Simulation integrated with optimization.

8 Case Study: Drone Medical Delivery

The effectiveness of the EaPF is verified through simulations. We consider a medical delivery mission as the demonstration environment. The mission is initialized with a heterogeneous task list consisting of 20 delivery requests, each of which features different delivery properties in view of realistic medical applications, as listed in Tab. 1. These requests are expected to be accomplished by a fleet of 3 vehicles, with 5 recharging stations deployed in the mission area to provide recharging services. All vehicles are sent out from one same depot and finally recycled there. Stations, the depot and all customer requests, including pickup and delivery tasks, are randomly located in a 10km×10km mission area. Each UAV's battery capacity is considered to be the same, 0.3kWh. The commanded airspeed is set as 15m/s.

Table 1 Properties of medical delivery requests.

Delivery Product	Value	Severity	Time requirement
Blood sample	4	1	20 min
Test result (critical)	3	1	30 min
Medical supply (critical)	2	1	35 min
Test result (non-critical)	1	1	40 min
Medical supply (non-critical)	1	1	60 min
Medical equipment	1	1	2 h

To demonstrate the capability of EaPF in terms of settling stochastic energy consumption, we optimize routes using the original NS algorithm, the NS algorithm with 20% battery safety margin (NS20%) and the NS algorithm integrated within the EaPF framework (NS-EaPF) respectively. The battery safety margin means a specific percentage reduction to the available battery capacity considered in planning with the aim to prevent battery exhaustion. The weights in objective functions are set as $w_r = 1.0, w_d = e^{-4}, w_e = 1.0$.

Tab. 2 compares the planning performance of the above three approaches in terms of risk regulation and quality of routes. Results are averaged using 1000 random scenario instances. Among them, one particular instance is extracted in Fig. 6 for further illustration. In Fig. 6, the route sequences planned by NS and NS-EaPF are depicted, along with the cumulative probability curves of finishing time for each delivery request or the energy consumption on recharging stations and the depot.

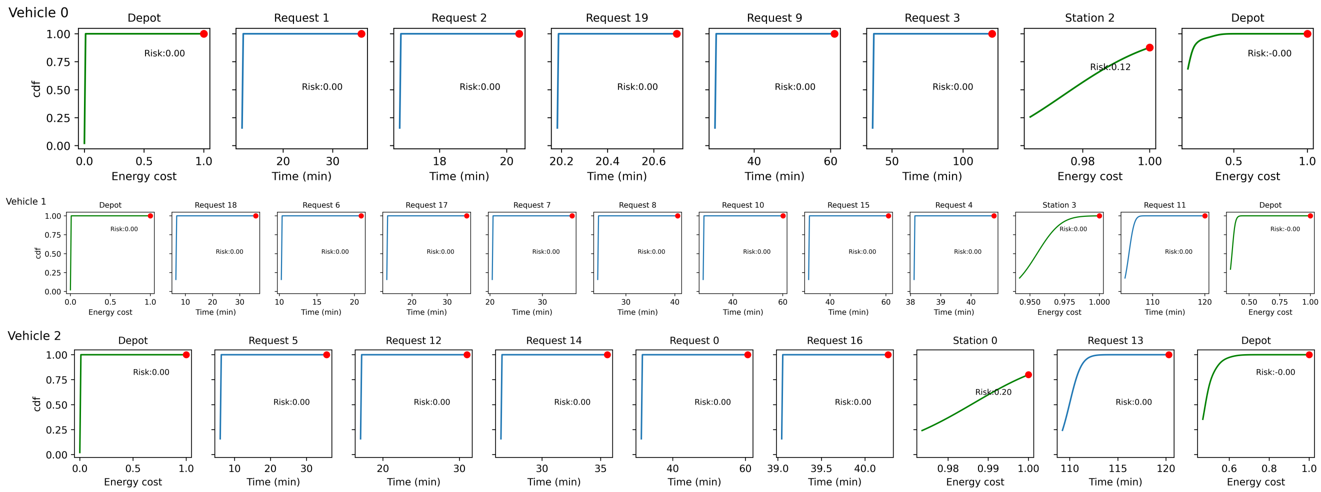

(a) Routes via NS

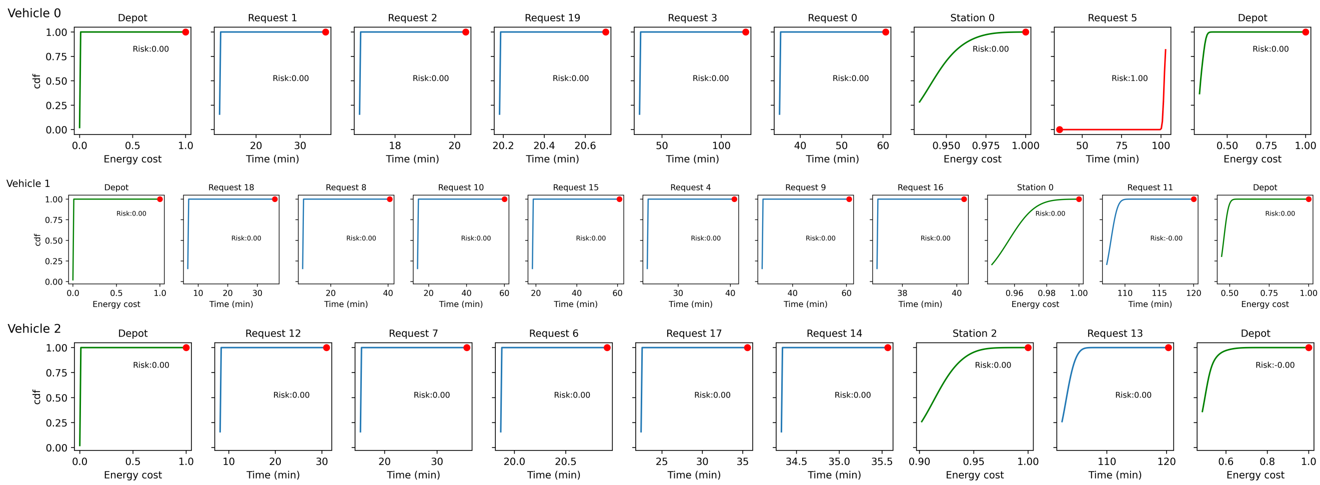

(b) Route via NS-EaPF

Fig. 6 Routes planned by NS and NS-EaPF. The executing order is from left to right. For customer request nodes, the cdf curve with respect to finishing time is plotted while the time consumption curve is presented for other nodes. Red dots indicate deadlines for requests or energy limitations for vehicles

As shown in Tab. 2, the approach combined with the EaPF yields a more reliable performance in terms of the maximum energy risks in comparison with original NS algorithm. Additionally, compared to NS20%, the NS-EaPF achieves higher objective values within a risk threshold of 1.0%. Setting a certain margin for battery usage is the strategy that is mostly used in industrial applications, which, however, comes with the question that how to define the margin value. With the UAV operating in varying wind fields, the fixed safety margin might either become risky or conservative.

Table 2 Performance comparison

Method	Task reward/penalty	Time delay	Energy cost	Overall objective	max energy risk
NS	27.60	780.9 min	4.77	18.14	14.37%
NS20%	20.80	894.6 min	4.85	10.58	0.02%
NS-EaPF	24.80	841.2 min	4.76	15.01	0.39%

Fig. 6 further illustrates the risk regulation function of NS-EaPF using one specific mission instance. In this case, vanilla NS tries to execute all delivery tasks while taking the risk of overusing the battery, e.g., the energy risk arrives at 0.20 when Vehicle2 is flying to Station 0 after finishing Task 16. In contrast, NS-EaPF follows a risk threshold, set as 1.0%, guaranteeing the safety of executing vehicles even at the cost of violating the time window constraint for a request.

9 Conclusion

This paper investigated the energy-risk aware planning problem for unmanned aerial logistic systems. Constrained by limited battery capacities, UAVs tend to visit recharging facilities frequently, featuring a relatively fluctuating state of charge, which makes its energy risk management worth studying. For that reason, we proposed an Energy-aware Planning Framework (EaPF), which integrates an energy prediction model, a routing simulator and a route optimizer. The energy prediction model adopts mixture density networks, facilitating insight into the distribution of energy costs without the need for sampling. Based on it, the energy risk is also analytically assessed. Then an event-driven simulating workflow is presented, which accommodates the simulation with the uncertainty model and a heuristic route optimizer. Finally, we evaluated the risk management ability of the EaPF via a medical delivery case stud. For comparison, an independent optimizer, an optimizer setting a given battery safety margin, and an optimizer integrated with the EaPF are respectively implemented. Results have shown that the approach incorporated in EaPF produces more reliable solutions where the risk is regulated within a certain threshold, and the quality of routes is also improved compared to the safety strategy using battery margins.

Dynamic events, such as the breakdown of UAVs and the arrival of new requests, are happening in the logistic systems from time to time. Our proposed discrete event simulator has the potential to simulate these scenarios. However, this kind of case is not within the scope of this paper. Also, the adopted heuristic approach is not efficient enough to respond to the dynamic event on time. Thus, for future work, we are planning to extend the implementation of EaPF to dynamic logistic systems and employ a more advanced route optimizer.

References

[1] Tugrul U Singireddy, Shiva Ram Reddy and Daim. Technology roadmap: Drone delivery–amazon prime air. In *Infrastructure and Technology Management*, pages 387—-412. Springer, 2018.

[2] Jack W. Langelaan, Sven Schmitz, Jose Palacios, and Ralph D. Lorenz. Energetics of rotary-wing exploration of Titan. *IEEE Aerospace Conference Proceedings*, 2017. DOI: 10.1109/AERO.2017.7943650.

[3] Thomas Kirschstein. Comparison of energy demands of drone-based and ground-based parcel delivery services. *Transportation Research Part D: Transport and Environment*, 78(December 2019):102209, 2020. DOI: 10.1016/j.trd.2019.102209.

[4] Zhilong Liu, Raja Sengupta, and Alex Kurzhanskiy. A power consumption model for multi-rotor small unmanned aircraft systems. *2017 International Conference on Unmanned Aircraft Systems, ICUAS 2017*, pages 310–315, 2017. DOI: 10.1109/ICUAS.2017.7991310.

[5] Joshuah K. Stolaroff, Constantine Samaras, Emma R. O'Neill, Alia Lubers, Alexandra S. Mitchell, and Daniel Ceperley. Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery. *Nature Communications*, 9(1):1–13, 2018. DOI: 10.1038/s41467-017-02411-5.

[6] Kevin Dorling, Jordan Heinrichs, Geoffrey G. Messier, and Sebastian Magierowski. Vehicle Routing Problems for Drone Delivery. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1):70–85, 2017. DOI: 10.1109/TSMC.2016.2582745.

[7] Alex S. Prasetia, Rong Jong Wai, Yi Lun Wen, and Yu Kai Wang. Mission-based energy consumption prediction of multirotor uav. *IEEE Access*, 7:33055–33063, 2019. DOI: 10.1109/ACCESS.2019.2903644.

[8] Arnav Choudhry, Brady Moon, Jay Patrikar, Constantine Samaras, and Sebastian Scherer. CVaR-based Flight Energy Risk Assessment for Multirotor UAVs using a Deep Energy Model. 2021.

[9] Thiago A. Rodrigues, Jay Patrikar, Arnav Choudhry, Jacob Feldgoise, Vaibhav Arcot, Aradhana Gahlaut, Sophia Lau, Brady Moon, Bastian Wagner, H. Scott Matthews, Sebastian Scherer, and Constantine Samaras. In-flight positional and energy use data set of a DJI Matrice 100 quadcopter for small package delivery. *Scientific Data*, 8(1):6–13, 2021. DOI: 10.1038/s41597-021-00930-x.

[10] Hoong Chuin Lau, William Yeoh, Pradeep Varakantham, Duc Thien Nguyen, and Huaxing Chen. Dynamic stochastic orienteering problems for risk-aware applications. *Uncertainty in Artificial Intelligence - Proceedings of the 28th Conference, UAI 2012*, pages 448–458, 2012.

[11] Ty Nguyen and Tsz Chiu Au. Extending the range of delivery drones by exploratory learning of energy models. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 3:1658–1660, 2017.

[12] Angel A. Juan, Javier Faulin, Scott E. Grasman, Markus Rabe, and Gonçalo Figueira. A review of simheuristics: Extending metaheuristics to deal with stochastic combinatorial optimization problems. *Operations Research Perspectives*, 2:62–72, 2015. DOI: 10.1016/j.orp.2015.03.001.

[13] D. E. Mazzuco, A. M. Carreirão Danielli, D. L. Oliveira, P. P.P. Santos, M. M. Pereira, L. C. Coelho, and E. M. Frazzon. A concept for simulation-based optimization in Vehicle Routing Problems. *IFAC-PapersOnLine*, 51(11):1720–1725, 2018. DOI: 10.1016/j.ifacol.2018.08.208.

[14] Thomas Mayer, Tobias Uhlig, and Oliver Rose. An open-source discrete event simulator for rich vehicle routing problems. *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, pages 1305–1310, 2016. DOI: 10.1109/ITSC.2016.7795725.

[15] Edward Farhan, Mohammed and Gohre, Brett and Junprung. Reinforcement learning in anylogic simulation models: a guiding example using pathmind. *2020 Winter Simulation Conference (WSC)*, 4(5):3212—-3223, 2020. DOI: 10.1038/s41562-020-0884-z.

[16] Markus Rabe, Felix Dross, and Alexander Wuttke. Combining a Discrete-event Simulation Model of a Logistics Network with Deep Reinforcement Learning. pages 765–774, 2017.

[17] Kunal Menda, Yi Chun Chen, Justin Grana, James W. Bono, Brendan D. Tracey, Mykel J. Kochenderfer, and David Wolpert. Deep Reinforcement Learning for Event-Driven Multi-Agent Decision Processes. *IEEE Transactions on Intelligent Transportation Systems*, 20(4):1259–1268, 2019. DOI: 10.1109/TITS.2018.2848264.

[18] James J.Q. Yu and Albert Y.S. Lam. Autonomous Vehicle Logistic System: Joint Routing and Charging Strategy. *IEEE Transactions on Intelligent Transportation Systems*, 19(7):2175–2187, 2018. DOI: 10.1109/TITS.2017.2766682.

[19] Andrew R. Runnalls. Kullback-Leibler approach to Gaussian mixture reduction. *IEEE Transactions on Aerospace and Electronic Systems*, 43(3):989–999, 2007. DOI: 10.1109/TAES.2007.4383588.

[20] Mark J. Schervish, Averill M. Law, and W. David Kelton. *Simulation Modeling and Analysis.*, volume 78. 1983. DOI: 10.2307/2288169.