

7-10-2021

Learning-to-Dispatch: Reinforcement Learning Based Flight Planning under Emergency

Kai Zhang
Embry-Riddle Aeronautical University, zhangk3@my.erau.edu

Yupeng Yang
Embry-Riddle Aeronautical University, yang5@my.erau.edu

Chengtao Xu
Embry-Riddle Aeronautical University, xuc3@my.erau.edu

Dahai Liu
Embry-Riddle Aeronautical University, liu89b@erau.edu

Houbing Song
Embry-Riddle Aeronautical University, h.song@ieee.org

Follow this and additional works at: <https://commons.erau.edu/publication>



Part of the [Aviation Commons](#), and the [Systems and Communications Commons](#)

Scholarly Commons Citation

Zhang, K., Yang, Y., Xu, C., Liu, D., & Song, H. (2021). Learning-to-Dispatch: Reinforcement Learning Based Flight Planning under Emergency. *IEEE Internet of Things Journal*, (). Retrieved from <https://commons.erau.edu/publication/1770>

This Article is brought to you for free and open access by Scholarly Commons. It has been accepted for inclusion in Publications by an authorized administrator of Scholarly Commons. For more information, please contact commons@erau.edu.

Learning-to-Dispatch: Reinforcement Learning Based Flight Planning under Emergency

Kai Zhang¹, Yupeng Yang¹, Chengtao Xu¹, Dahai Liu² and Houbing Song¹

Abstract—The effectiveness of resource allocation under emergencies especially hurricane disasters is crucial. However, most researchers focus on emergency resource allocation in a ground transportation system. In this paper, we propose Learning-to-Dispatch (L2D), a reinforcement learning (RL) based air route dispatching system, that aims to add additional flights for hurricane evacuation while minimizing the airspace’s complexity and air traffic controller’s workload. Given a bipartite graph with weights that are learned from the historical flight data using RL in consideration of short- and long-term gains, we formulate the flight dispatch as an online maximum weight matching problem. Different from the conventional order dispatch problem, there is no actual or estimated index that can evaluate how the additional evacuation flights influence the air traffic complexity. Then we propose a multivariate reward function in the learning phase and compare it with other univariate reward designs to show its superior performance. The experiments using the real-world dataset for Hurricane Irma demonstrate the efficacy and efficiency of our proposed schema.

Index Terms—Evacuation, reinforcement learning, air traffic management, graph theory.

I. INTRODUCTION

Natural disasters such as floods, earthquakes, and wildfires occur around the world almost every year and cause loss of life or damage property [1]. Hurricane is the most common flooding disaster in the Southeastern United States and it has killed or injured many populations. Therefore, how to design an effective evacuation plan before a hurricane to reduce casualties becomes a significant problem for emergency planning and management [2], [3].

According to the Hurricane Irma Local Report, approximately 6+ million South Floridians joined in evacuating. Although most of them began paying very close attention to Irma’s forecast up to a week and left early, they still encountered significant traffic jams upstate. To deal with this issue, recent studies [4]–[6] focus on vehicle routing or route choice problem during hurricane evacuation considering both dynamic evacuation demand [7] and hurricane characteristics. In summary, they estimated evacuee behavior and assessed road accessibility using statistical models, then applied dynamic traffic assignment approaches [8] to solve for the paths and

* This work was supported by the Center for Advanced Transportation Mobility (CATM), USDOT Grant #69A3551747125.

¹Kai Zhang, Yupeng Yang, Chengtao Xu and Houbing Song are with the Security and Optimization for Networked Globe Laboratory (SONG Lab), Embry-Riddle Aeronautical University, Daytona Beach, FL 32114 USA. {zhangk3, yangy5, xuc3}@my.erau.edu, h.song@ieee.org

²Dahai Liu is with the College of Aviation, Embry-Riddle Aeronautical University, Daytona Beach, FL 32114 USA. dahai.liu@erau.edu

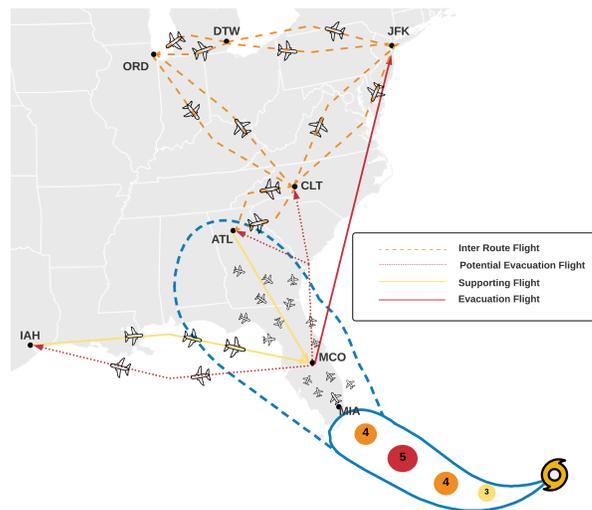


Fig. 1: Flight Dispatch under Emergency.

travel times of each driver. Indeed, the practical application of these researches is in doubt because people do not usually follow the recommended route, especially in emergencies. The indeterminate driving behaviors bring unknowable disturbance to the system. Consequently, taking a flight to evacuate before the hurricane is a better option since air traffic control (ATC) is a centralized system in which “drivers” and paths are managed by the air transportation authorities. However, to the best of our knowledge, there has been no literature thus far that tries to address the flight dispatch during a hurricane evacuation. On the other hand, recent advances in Internet of Things and computing power made large-scale traffic data be collected and analyzed in the modern intelligent transportation system (ITS) [9]–[13]. As the transportation management and control becomes more data driven [14]–[16], machine learning has drawn a lot of academic interest in the application of disaster evacuation planning [17]–[20]. However, previous work purely focused on classification tasks, such as driver’s route choice [17], [19], contraflow activation [18] and human’s sentiments [20], while they naturally ignore the potential application of data-driven evacuation planning via machine learning.

In the paper, we propose Learning-to-Dispatch (L2D), a reinforcement learning (RL) based flight dispatch system that aims to maximize the air traffic capacity, and meanwhile minimize the potential negative impact of additional flights on the air route network. To accomplish this, we first gen-

erate synthetic flight transactions and model the dynamics of the air route network during hurricane evacuation based on the proposed reward function which balances the trade-off between evacuation efficiency and airspace capacity. The flight planning is then formulated as a matching problem, where each individual decision of matching an aircraft to an evacuation flight is based on the instant reward for the aircraft serving this flight and the impact of this decision in the future.

The contribution of this work is summarized as follows:

- To our knowledge, we firstly propose an effective flight dispatch system for supporting disaster evacuation. The system considers both instance requirements of evacuation efficiency and the expected future air traffic complexity.
- We present an elaborated reward function. The experimental results prove that it results in a more balanced dispatch plan in terms of flight length, flight elapsed time, and air traffic controller’s workload, which is likely to play a vital role in the future optimization of evacuation flight dispatching research.

The remainder of this paper is structured as follows. Section II formulates the flight dispatch problem. The L2D system design is detailed in Section III, including synthetic data generation and model construction. Experiments are then described in Section IV, followed by the conclusion of the paper in Section V.

II. PROBLEM FORMULATION

Figure. 1 shows the basic idea behind the L2D system. More specifically, flights are categorized by task type: (a) inter route is the scheduled flight which serves regular trips; (b) evacuation route is the extra flight or charter flight for hurricane evacuation; (c) supporting route can be considered a kind of evacuation route but it is not "direct" one whose departure is in the evacuation zones. Given the estimated arrival time of hurricane T_h , the flight dispatching aims to complete the urgent egress of people away from a city at least T hours before a hurricane approaches. To achieve this goal via a coordinated and optimized way, we should take two principle problems into account: (a) *How to measure the long-term increase of air traffic complexity caused by the evacuation flight?*, and (b) *How to allocate available resources on basis of future complexity estimation?*

To solve these issues, we model flight dispatch as a sequential decision problem, say fully observable Markov decision process (MDP), in which there is an embedded decision node of matching aircraft-airway at each stage. The key terms of MDP we build in the paper are discussed in the following paragraphs and illustrated in Figure 2.

A. MDP Definition

An MDP is defined by a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$, where $\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma$ are the set of states, set of actions, transition probability function, reward function and a discount rate, respectively. The purpose of the agent is to maximize the long-run reward $G_t = \sum_{k=0}^T \gamma^k R_{t+k+1}$ it receives. There are



Fig. 2: MDP formulation in the L2D. An agent represents an aircraft/pilot. The action in flight dispatch represents assigning the aircraft/pilot to serve a particular evacuation flight.

two strategies to model an agent: (a) To model the dispatch system as an agent; and (b) To model the individual aircraft preparing for serving evacuation as an agent. For the first strategy, it is hard to define components in an MDP in an appropriate manner. As an example, we have to cope with the large combinatorial action spaces. Therefore, we adopt the second strategy. However, it still exists a major problem that the state spaces and action spaces are changing for each agent at different stages. Here we use a generalized setting. Specifically, we regard each agent as the same so that each agent performs the same policy π . Other components in an MDP can be defined as follows:

- **State** \mathcal{S} is determined by a spatial-temporal vector $\langle g, t \rangle$ where g is the airport index where the aircraft is located and t is the time index.
- **Action** \mathcal{A} is defined as a set $\{0, 1\}$ where 0 and 1 represent that the aircraft is assigned to a flight or not in a time slot, respectively.
- **State Transition** \mathcal{P} depends on the estimated future information such as flight delay and air flow. We use a naive way to estimate these information for simplicity that we assume the relevant variables are distributed normally with the calculated means and variances based on historical records.
- **Reward** \mathcal{R} is defined as an index with hybrid factors which we think can reflect the influence of an evacuation flight to the corresponding airports. Details are seen in § II-B, the definition of ATCI impact factor (ATCI).
- **Discount rate** γ determines how important future rewards to the current state. We set $\gamma = 0.9$ in the paper.

B. Dynamics of the Airspace Complexity

We use the ATCI to measure the airspace complexity which is defined as follows:

$$\text{ATCI} = \frac{100}{\lambda \times w + \mu \times et + \nu \times d} \quad (1)$$

where w is an indicator for air traffic controller’s workload. et denotes the elapsed time of a flight, d represents the distance, while λ, μ, ν are hyperparameters. The details of how to calculate indicators and select hyperparameters are described in § III-A.

Recall the MDP definition, an agent performs a policy π in an environment (the air route network in our problem setting) and tries to maximize its gain. To get the optimal policy, one path is to learn the value function $V_\pi(s) = \mathbb{E}_\pi[G_t|S_t = s]$ that estimate how good it is for an agent does a given action in a given state. Logically, it is clear that $V_\pi(s)$ measures the dynamics of airspace complexity since $V_\pi(s) \sim \text{ATCI}$ and ATCI is related to factors of airspace complexity. In other words, the learned value function captures the spatio-temporal patterns of the air route network that can further used to dispatch evacuation flights.

C. Flight Dispatching under Emergency

Now we assume there are total m available aircraft and n airports at a time slot t that are able to assist people to evacuate. The goal of flight dispatching is to find the best match between potential aircraft and potential evacuation flights that are defined mathematically as follows:

- **Potential Aircraft** is denoted by a set of quadruple $\langle I_a, S_a, A_a, C_a \rangle$, where I_a, S_a, A_a, C_a are the index of aircraft, current state, availability index and expected cumulative cost respectively.
- **Potential Evacuation Flights** is represented by a set of quintuple $\langle I_o, S_D, A_o, C_o, S_A \rangle$, where I_o, S_D, A_o, C_o, S_A are the index of flight, departure state, availability index, expected ATC impact, and expected arrival state respectively. Note that in the evacuation task, the airway could be a single itinerary from city C to the arrival city or multi-way itinerary. For example, an aircraft takes off in the city C_D , and flies to C to support the evacuation and then goes to the other city C_A .

The best match means that at each time step, potential airways are paired with the unoccupied aircraft resulting in the maximum expected dispatch gain or minimum expected air traffic complexity. Formally, this task's objective function is written as follows:

$$\begin{aligned} & \operatorname{argmax}_{a_{ij}} \sum_{i=0}^m \sum_{j=0}^n W(i, j) a_{ij} \\ & \text{s.t.} \sum_{i=0}^m a_{ij} = 1, \quad j = 1, 2, 3 \dots, n \\ & \sum_{j=0}^n a_{ij} = 1, \quad i = 1, 2, 3 \dots, m \end{aligned} \quad (2)$$

where $a_{ij} = 1$ if airway j is assigned to aircraft i while $a_{ij} = 0$ if airway j is not assigned to aircraft i . $W(i, j)$ is a function that could point out the long-term return or cost if an airway is dispatched.

The Eq. 2 can be further defined as a bipartite graph matching problem with weighted edges as shown in Figure. 3. The weights here are learned by RL. To put it another way, the weights come from the value function. We will explain how to obtain the value function in § III-B.

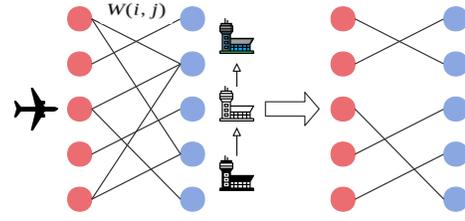


Fig. 3: An aircraft-airway maximum-weight matching example, in which i is the aircraft index while j is the airway index. The weights of each potential connection are obtained by the estimated value functions.

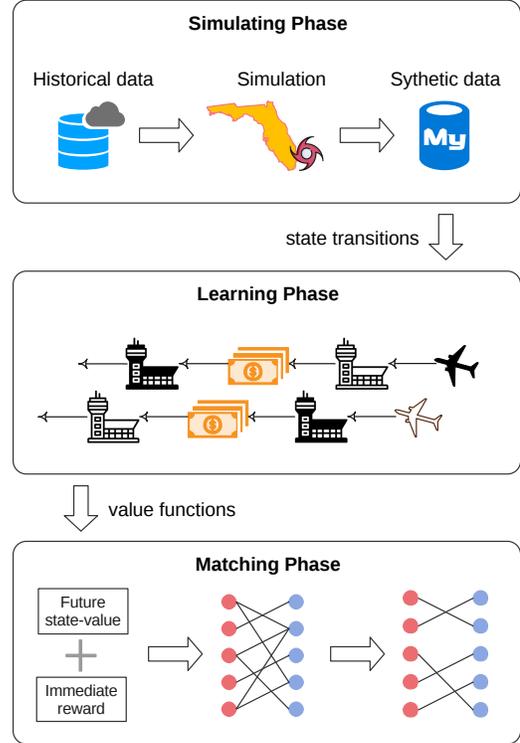


Fig. 4: Overview of L2D system.

III. SYSTEM DESIGN

L2D system consists of three modules: offline simulating, offline learning, and online matching, as illustrated in Figure. 4. Although it is computationally expensive for learning the value function, it could be completed before disasters approach. Thus, offline process is not time-sensitive in practice.

A. Generating Simulated Evacuation Flight Transactions

The flight transactions in the historical database record only a series of one-way itineraries such as from Miami to New York City. However, in the context of disaster evacuation, an evacuation airway could be a multi-way itinerary. Thus we generate simulated evacuation flight transactions based on the statistical properties from the historical real-world data, which will be utilized to learn the value functions later. The progress of producing synthetic transactions is described in

Algorithm 1. In each episode, we randomly select departure and arrival states including airport and time from the historical flight dataset, and treat the availability index as either 0 (not available) or 1 (available).

At a time t , workload $w^k = air^k/atc^k$, time et^k , and distance d^k could measure a one-way flight as well as a multi-way flight. Let's see an example to make the process clearer. There is a transaction from k_1 , supporting k_2 to transport people to k_3 , therefore the w_t^k should be the accumulated workload when an aircraft takes off or lands in these three airports at the corresponding time step, which is calculated as $w_t^k = w_{t_1}^{k_1} + w_{t_2}^{k_2} + w_{t_3}^{k_3}$ where t_1 represent the time of departure and t_2, t_3 represent the estimated time of arrival.

Algorithm 2 is the hyperparameter selection strategy for computing ACTI. The intuition behind the hyperparameter selection algorithm is to let every variable be negatively correlated with the ATCI as much as possible so that the reward function could not be heavily biased toward either factor.

B. Building Weighted Aircraft-Airway Graph

We adopt the advantage function trick in [21] to reduce the computational complexity by removing connections between aircraft and the "idle" action. The advantage function measures what is the advantage of selecting a certain action in a certain state compared to expected reward of all possible actions in that state. Mathematically, the advantage function is defined as:

$$A_\pi(i, j) = Q_\pi(i, j) - V_\pi(s_i) \quad (3)$$

where $Q_\pi(i, j)$ is the action-value function of aircraft i performing action of serving an evacuation transaction j or "do nothing". The value of this function shows the expected return taking action a in state s under policy π .

According to Eq. 2, we denote the case of $i = 0$ or $j = 0$ as the situation that the airway is not assigned to any aircraft. Let s_i and s'_i denote the current state of the aircraft i , and the final state when an evacuation flight transaction is completed, the state-value function is thus obtained with dynamic programming based on Temporal-Difference (TD) method [22] as described in Algorithm 3. Therefore, we can further rewrite Eq. 3 in the following:

$$A_\pi(i, j) = \gamma^{\Delta t_j} V(s'_{ij}) + R_j - V(s_i) \quad (4)$$

Now edge weight of each potential evacuation flight can be calculated using Eq. 4 to build the weighted bipartite graph. If $Q_\pi(i, j) \leq V_\pi(s_i)$, we let the weight between i and j be zero to reduce the potential pairs, then to reduce the computational complexity.

C. Matching Maximum-Weighted Connections

The aforementioned bipartite graph contains two sets of nodes – aircraft and airway, and the edge between them has a weight of $A_\pi(i, j)$. We adopt the widely used Hungarian algorithm to finds a maximum matching [23].

Algorithm 1 Evacuation Flight Transaction Productor

Input: Historical flight data \mathbf{F} and human-factor data \mathbf{H} ; pre-defined number of transactions \mathbf{N} . The airports k and time t are extracted randomly from the historical datasets in each iteration.

Output: Synthetic airway set $\{\langle I_o, S_D, A_o, C_t^k, S_A \rangle\}$

- 1: **for** $i = 0$ to $i = \mathbf{N}$ **do**
 - 2: Generate a potential airway tuple $\langle I_o, S_D, A_o, S_A \rangle_i$ but without the ATCI.
 - 3: $atc_t^k \leftarrow$ number of air traffic controllers \times the corresponding training time
 - 4: $air_t^k \leftarrow$ the average of delay \times the number of flights
 - 5: $et_t^k \leftarrow$ sample from the historical elapsed time distribution complying with Gaussian
 - 6: $d_t^k \leftarrow$ flight length
 - 7: $C_t^k = 100/[\lambda \times (air_t^k/atc_t^k) + \mu \times et_t^k + \nu \times d_t^k]$
 - 8: Take the C_t^k as the ATCI in the airway tuple to construct $\langle I_o, S_D, A_o, C_t^k, S_A \rangle_i$.
 - 9: **end for**
-

Algorithm 2 Hyperparameter Selection

Input: A pre-defined hyperparameter candidate set $\{\mathcal{H}_i\}$ where i is the index of each candidate, $i \in [0, N_{hyper}]$.

Output: An optimal hyperparameter tuple $\langle \lambda^*, \mu^*, \nu^* \rangle$.

Initialize ρ_{hyper} and $\mathcal{H}^* = \langle \lambda^*, \mu^*, \nu^* \rangle$ of all zeros.

- 1: **for** $i = 0$ to $i = N_{hyper}$ **do**
 - 2: $\rho_i^w, \rho_i^d, \rho_i^t \leftarrow$ calculate the Spearman rank-order correlation coefficient of ATCI and workload, flight length, elapsed time in \mathcal{H}_i , respectively.
 - 3: **if** $\rho_i \leq 0.4$ for each coefficient **then**
 - 4: $\rho_i^{neg} = |\rho_i^w + \rho_i^d + \rho_i^t|$
 - 5: **if** $\rho_i^{neg} > \rho_{hyper}$ **then**
 - 6: $\rho_{hyper} \leftarrow \rho_i^{neg}$, $\mathcal{H}^* \leftarrow \mathcal{H}_i$
 - 7: **end if**
 - 8: **end if**
 - 9: **end for**
-

Algorithm 3 State-value Lookup Generator

Input: Collect state transitions $\{\langle s_i, a_i, r_i, s'_i \rangle\}$ from synthetic evacuation transactions. State is represented as a spatio-temporal tuple: $s_i = \langle g_i, t_i \rangle$, $s'_i = \langle g'_i, t'_i \rangle$. T is the time of last state in the finite-horizon MDP.

Output: Value function $V(s)$ for all states.

Initialize $V(s)$ and $N(s)$ of all zeros

- 1: **for** $t = T - 1$ to $t = 0$ **do**
 - 2: $\{\langle s_i, a_i, r_i, s'_i \rangle\}^t \leftarrow$ A subset of state transitions in which $t_i = t$
 - 3: **for** each instance in the subset **do**
 - 4: $N(s_i) \leftarrow N(s_i) + 1$
 - 5: TD error: $\delta_i \leftarrow \gamma^{\Delta t(a_i)} V(s'_i) + R(a_i) - V(s_i)$
 - 6: $V(s_i) \leftarrow V(s_i) + \frac{1}{N(s_i)} \delta_i$
 - 7: **end for**
 - 8: **end for**
-

TABLE I: Hyperparameters and Correlation Coefficients

Candidate $\langle \lambda^*, \mu^*, \nu^* \rangle$	ρ^t	ρ^d	ρ^w
(1,0.01,0.2)	-0.77	-0.73	-0.44
(1,0.01,0.1)	-0.89	-0.84	-0.2
(1,0.01,0.5)	-0.45	-0.5	-0.63
(1,0.005,0.2)	-0.66	-0.6	-0.5
(1,0.006,0.2)	-0.65	-0.65	-0.5
(1.5,0.006,0.2)	-0.72	-0.73	-0.42
(1.5,0.006,0.4)	-0.47	-0.5	-0.66

In the scenario of online evacuation flight dispatching, L2D collects all available aircraft and potential airways at every time step to build aircraft-airway pairs with weights that are estimated using Eq. 4, and afterward the matching is solved by Hungarian algorithm. Remaining resources will go to the next assignment until the iteration meets the termination.

IV. PERFORMANCE EVALUATION

Detailed setup and results of our experiments are illustrated in this section, which provide a more complete understanding of the efficacy and efficiency of our proposed schema.

A. Implementation Details

Data source: We collect flight and human-factor data for 2017 September from the Airline On-Time Performance Database and 123ATC.com¹.

Evacuation Plans: We assume that residents in Miami (MIA) would like to be evacuated by flight 6 hours before the arrival of the hurricane T_h , and there are 5 supporting and destination airports – ATL, BNA, CLT, DFW, and JFK. The time step in the L2D system is set as 1 hour. 200 evacuation flights are generated at each time step, and the evacuation plan must be completed before T_h with a fixed number of evacuation-oriented aircraft N , which are evenly assigned to the relevant airports. We also suppose that aircraft can transport people immediately without waiting for overhaul, and passenger boarding and landing.

Reward function settings: Table I shows some pre-defined hyperparameter candidates and their Spearman rank-order correlation coefficients in terms of ATCI. As described in Algorithm 2, we select $\lambda = 1, \mu = 0.01, \nu = 0.2$ in our reward function design. The correlations of indexes for the optimal hyperparameter tuple are depicted in Figure 5. We can see that w, et, d are all roughly negatively correlated with the ATCI, which tallies with the reward function design intention.

Evaluation metrics: Our pre-defined ATC impact factor is the multiplicative inverse for workload, flight length and elapsed time, thus the larger it is, the lower incremental air traffic complexity is. We also consider two more intuitive evaluation metrics: "How many evacuation flights can be completed within 6 hours?" and "How many evacuation flights each airport has to coordinate with?"

B. Hurricane Irma Example

The eye of Irma is close to Southern Florida in the morning on September 10, 2021. To coordinate evacuation flights in a safer manner, we set 6 AM on that day as the evacuation end time. N is set as 30, 60, 120 respectively. All other experimental settings have been described in the implementation details.

Figure. 6 shows comparison results of the four methods for the number of completed flights within 6 hours starting at 12 AM on September 10, 2021. Note that the transit flights via MIA are not counted and there are a total of 5 departure flights. It is not astonishing that the time-based and distance-based methods tend to dispatch evacuation flights to ATL and CLT since they are the closest airports to MIA compared to the other 4 airports, while the elapsed time is usually positively correlated with flight length. On the contrary, very few evacuation flights will take off or land from ATL and CLT using the workload-based method. Such a result reconciles with the real conditions since they are two of the world's busiest airports as well as the nearest airports to MIA, evacuation flights could increase airspace complexity in a short time. Besides, airports in the experimental settings are almost large hub/core airports in the US. Therefore, it is easy to explain that why the workload-based method produces far fewer airway-aircraft pairs than others, the same result is also shown in Figure. 7. It is worth pointing out that a more balanced flight allocation plan is obtained with our designed reward function. The extra workload caused by evacuation is nearly uniformly shared by all airports.

The comparison results for different airway-aircraft ratios are illustrated in Figure. 7. Here the "200" in the airway-aircraft ratio indicates that 200 potential airways are generated at each time step, then they can be paired with available aircraft. As the ratio increases, there is a convergent phenomenon in the evaluation metrics depending on average elapsed time and flight length. In other words, all strategies tend to achieve similar performance because of sufficient aircraft supply. Otherwise, when the supply is not enough such as $N = 30$ or $N = 60$, the workload-based method results in more flight time and flight distance for each evacuation flight compared to the other three methods.

Assume that an evacuation aircraft can transport 400 people at a time, approximately 28,000 people can evacuate towards other cities within 6 hours based on L2D even if only 30 aircraft serves for evacuation. Therefore, Figure. 6 and Figure. 7 jointly state that the ATCI-based method can produce good and unbiased evacuation flight plans.

V. CONCLUSION

In this paper, we propose a new flight dispatch system L2D under emergency that aims to optimize the air route network's long-term efficiency, as well as satisfying instant evacuation demands. The flight dispatch is modeled as an MDP, where the value of the aircraft-airway pair is obtained by the flight's utility and the future expected ATCI value learned from synthetic data based on historical records. Matching between multiple

¹<https://123atc.com/>

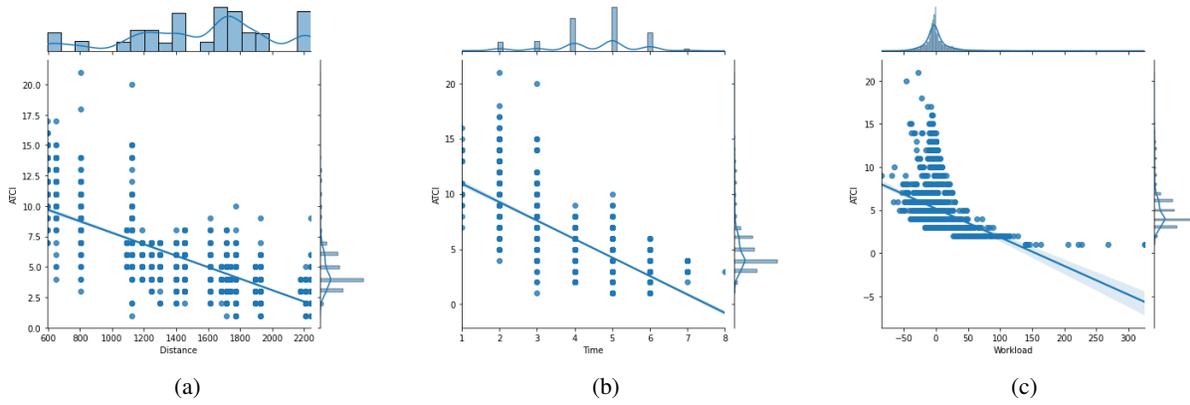


Fig. 5: The distribution and relationships amongst multiple indexes.

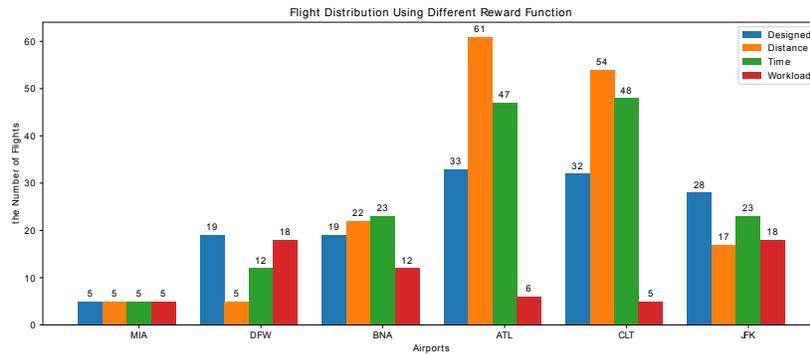


Fig. 6: Flight distribution with different reward functions. The transit flights via MIA are not counted while $N = 30$.

aircraft and potential evacuation airways is then solved by the Hungarian algorithm. The real-world case study based on Hurricane Irma reveals the effectiveness and efficiency of the proposed L2D and ATCI-based reward function compared to the baseline univariate methods. Although it is an early work on flight dispatch under emergency, the results show its potential to be deployed in the real-world disaster management system in the future.

Our proposed L2D is still in the early stage. For future work, we are interested in developing a large-scale and real-time flight dispatch schema based on passenger demand and weather fluctuation, which could be applied for nationwide evacuation planning. Meanwhile, it is beneficial to investigate the approach which can better capture the dynamics of the air route network during disaster evacuation.

REFERENCES

- [1] A. B. Smith and R. W. Katz, "Us billion-dollar weather and climate disasters: Data sources, trends, accuracy and biases," *Natural hazards*, vol. 67, no. 2, pp. 387–410, 2013.
- [2] Y. Zhu, K. Ozbay, K. Xie, and H. Yang, "Modeling of incident-induced capacity loss for hurricane evacuation simulation," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 613–618.
- [3] N. K. Tri, S. J. Wright, J. M. Betts, and H. L. Vu, "Studying the impact of public transport on disaster evacuation," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–6.
- [4] W. Yi, L. Nozick, R. Davidson, B. Blanton, and B. Colle, "Optimization of the issuance of evacuation orders under evolving hurricane conditions," *Transportation Research Part B: Methodological*, vol. 95, pp. 285–304, 2017.
- [5] Y.-J. Zhu, Y. Hu, and J. M. Collins, "Estimating road network accessibility during a hurricane evacuation: A case study of hurricane irma in florida," *Transportation research part D: transport and environment*, vol. 83, p. 102334, 2020.
- [6] M. A. Dulebenets, J. Pasha, O. F. Abioye, M. Kavousi, E. E. Ozguven, R. Moses, W. R. Boot, and T. Sando, "Exact and heuristic solution algorithms for efficient emergency evacuation in areas with vulnerable populations," *International journal of disaster risk reduction*, vol. 39, p. 101114, 2019.
- [7] W. Yin, P. Murray-Tuite, S. V. Ukkusuri, and H. Gladwin, "An agent-based modeling system for travel demand simulation for hurricane evacuation," *Transportation research part C: emerging technologies*, vol. 42, pp. 44–59, 2014.
- [8] Y. Lv, X. Zhang, W. Kang, and Y. Duan, "Managing emergency traffic evacuation with a partially random destination allocation strategy: A computational-experiment-based optimization approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 4, pp. 2182–2191, 2015.
- [9] H. Song, R. Srinivasan, T. Sookoor, and S. Jeschke, *Smart cities: foundations, principles, and applications*. John Wiley & Sons, 2017.
- [10] Y. Lv, Y. Duan, W. Kang, Z. Li, and F.-Y. Wang, "Traffic flow prediction with big data: a deep learning approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 865–873, 2014.
- [11] Y. Sun, H. Song, A. J. Jara, and R. Bie, "Internet of things and big data analytics for smart and connected communities," *IEEE access*, vol. 4, pp. 766–773, 2016.

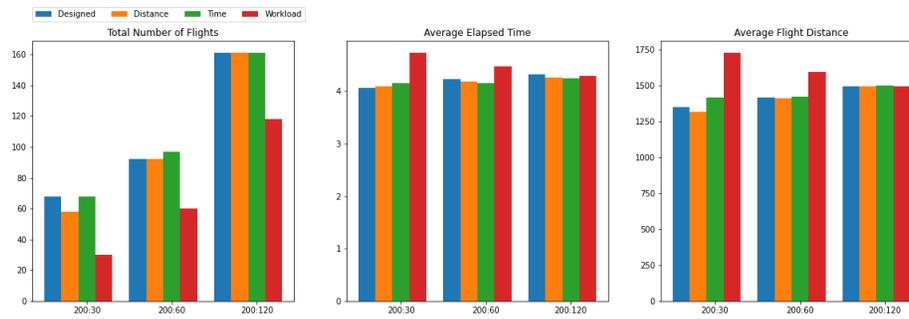


Fig. 7: Comparison of distance-based method, time-based method, workload-based method and the ATCI-based method in three metrics on the Hurricane Irma example. X-axis stands for the ratio of potential airway-aircraft at each time step.

- [12] Y. Liu, X. Weng, J. Wan, X. Yue, H. Song, and A. V. Vasilakos, "Exploring data validity in transportation systems for smart cities," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 26–33, 2017.
- [13] G. Dartmann, H. Song, and A. Schmeink, *Big data analytics for cyber-physical systems: machine learning for the internet of things*. Elsevier, 2019.
- [14] L. Zhu, F. R. Yu, Y. Wang, B. Ning, and T. Tang, "Big data analytics in intelligent transportation systems: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 1, pp. 383–398, 2018.
- [15] Y. Jiang, S. Niu, K. Zhang, B. Chen, C. Xu, D. Liu, and H. Song, "Spatial-temporal graph data mining for iot-enabled air mobility prediction," *IEEE Internet of Things Journal*, 2021.
- [16] K. Zhang, Y. Jiang, D. Liu, and H. Song, "Spatio-temporal data mining for aviation delay prediction," in *2020 IEEE 39th International Performance Computing and Communications Conference (IPCCC)*. IEEE, 2020, pp. 1–7.
- [17] X. Zhao, R. Lovreglio, and D. Nilsson, "Modelling and interpreting pre-evacuation decision-making using machine learning," *Automation in Construction*, vol. 113, p. 103140, 2020.
- [18] J. W. Burris, R. Shrestha, B. Gautam, and B. Bista, "Machine learning for the activation of contraflows during hurricane evacuation," in *2015 IEEE Global Humanitarian Technology Conference (GHTC)*. IEEE, 2015, pp. 254–258.
- [19] X. Song, Q. Zhang, Y. Sekimoto, T. Horanont, S. Ueyama, and R. Shibasaki, "Modeling and probabilistic reasoning of population evacuation during large-scale disaster," in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013, pp. 1231–1239.
- [20] D. Reynard and M. Shirgaokar, "Harnessing the power of machine learning: Can twitter data be useful in guiding resource allocation decisions during a natural disaster?" *Transportation research part D: Transport and environment*, vol. 77, pp. 449–463, 2019.
- [21] Z. Xu, Z. Li, Q. Guan, D. Zhang, Q. Li, J. Nan, C. Liu, W. Bian, and J. Ye, "Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 905–913.
- [22] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [23] D. Bruff, "The assignment problem and the hungarian method," *Notes for Math*, vol. 20, no. 29-47, p. 5, 2005.