

1 [TITLE PAGE]

2
3 **Multi-Objective Reinforcement Learning for Autonomous Drone Navigation in Urban**
4 **Areas with Wind Zones**

5
6 Jiahao Wu¹, Yang Ye², Jing Du, Ph.D.^{3*}

7 ¹Ph.D. Student, Informatics, Cobots and Intelligent Construction (ICIC) Lab, Department of
8 Civil and Coastal Engineering, University of Florida, Gainesville, FL 32611; e-mail:
9 jiahaowu@ufl.edu

10 ² Ph.D. Candidate, Informatics, Cobots and Intelligent Construction (ICIC) Lab, Department of
11 Civil and Coastal Engineering, University of Florida, Gainesville, FL 32611; e-mail:
12 ye.yang@ufl.edu

13 ³Professor, Informatics, Cobots and Intelligent Construction (ICIC) Lab, Department of Civil and
14 Coastal Engineering, Univ. of Florida, Gainesville, FL 32611 (corresponding author); Email:
15 eric.du@essie.ufl.edu

16

17

18 **Reinforcement Learning for Autonomous Drone Multi-Objective Navigation in Urban**

19 **Areas with Wind Zones**

20 **ABSTRACT**

21 Drones can be used for tasks such as data collection and logistics in civil engineering. Current
22 research on autonomous drones mainly focuses on planning a safe path and avoiding obstacles in
23 a static environment. However, navigating a drone in complex environments like urban areas
24 involves many dynamic constraints, such as building layout, winds, and signal coverage, which
25 are interdependent. The wind factor is the most important among these environmental factors,
26 which may cause a drone to lose control or even crash. This paper presents a multi-objective
27 navigation reinforcement learning algorithm (MONRL) for the drone to navigate and avoid
28 obstacles in an unknown environment when dynamic wind zones present, with only imagery data
29 about the building layouts. Based on a deep reinforcement learning and memory architecture, the
30 drone develops policies to prioritize navigation decisions, optimizing the path while minimizing
31 negative impact of winds with only sparse sensor data, in our case, camera inputs. By leveraging
32 the advantages of the proposed method in estimating the environmental factors from previous
33 trials, no aerodynamic force sensors are needed for the drone to develop effective strategies to
34 navigate to target while counteracting to milder winds and dodging away from stronger winds.
35 The proposed method was tested in a virtual environment, and a real model of New York City.
36 The method is expected to contribute to new automation algorithms for urban aerial logistics and
37 future automated civil infrastructure inspection.

38 **KEYWORDS:** Autonomous drones; Reinforcement Learning; Policy network; Wind simulation

40 **1. INTRODUCTION**

41 Unmanned Aerial Vehicles (UAVs), or drones, have gained increasing attention and popularity
42 in many industrial applications due to their flexibility, versatility, and efficiency in disaster
43 response [54], search and rescue [42], and environmental monitoring [15]. Most UAVs are
44 piloted by a human operator from a distance, while recent advances in control algorithms and
45 sensor systems have started to support autonomous drones in complex tasks and workplaces

46 [24]. Autonomous drones can plan out their path and achieve the flying goals automatically
47 based on dynamic environmental conditions from the onboard or remote sensors, enabling them
48 to reach distant areas or access difficult or dangerous locations that would be challenging or
49 impossible for humans to access [51]. In addition, autonomous drones can free human pilots
50 from repetitive tasks such as logistics delivery, helping the growth of industries despite the
51 shortage of human pilots [59]. Especially, it is expected that the scaled adoption of autonomous
52 drones in the logistics industry can help reduce traffic congestion and carbon emissions by
53 delivering goods through the air instead of on the ground [53]. These drones can also be
54 equipped with a variety of sensors and cameras, such as thermal imaging and LiDAR, to collect
55 high-quality data and imagery on the fly, supporting future tasks such as aerial mapping,
56 surveying, and monitoring [22] [65] [21].

57 The key to the safe and effective implementation of autonomous drones in logistics is
58 path planning, i.e., determining the optimal route for a drone to follow from its starting point to
59 its destination while satisfying various requirements, such as avoiding obstacles and minimizing
60 travel time [34]. Most existing path planning methods for the drone's navigation system mainly
61 focus on identifying the shortest route and avoiding obstacles [39]. While in real-world
62 applications, dynamic environment conditions (e.g., weather conditions) could have unexpected
63 impacts on the navigation controls of autonomous drones. For example, in urban areas
64 surrounded by tall buildings, the changing aerodynamic conditions can significantly impact the
65 desired outcomes of drone navigation. Another example is that unexpected turbulences
66 occurring between two close buildings may influence the stabilization of a drone [12]. On a
67 larger scale, the layout of buildings in a region, such as a metropolitan district, can affect the
68 wind zone distribution in a complicated way, making a pre-determined drone path plan less
69 optimal [11]. Dynamic winds happening anywhere in a large region, including strong gusts of
70 wind, can cause the drone to lose control or even crash. Even milder winds can affect the
71 maneuverability of the drone, causing it to drift away from its intended path. Lastly, headwinds
72 can impact the drone's battery life and flight time, as the drone requires using more power to
73 compensate for the wind resistance. All of these dynamic environmental conditions make the
74 traditional path planning of autonomous drones less effective in achieving goals in a safe way.

75 This paper proposes a multi-objective navigation reinforcement learning (MONRL)
76 algorithm for the drone’s path planning while safely across the normal wind zones and dodging
77 the strong wind zones. A strong wind zone refers to an area where the wind speed exceeds the
78 safe allowance of a particular drone, which could cause clashes, while the normal wind zones
79 would cause drafting problems and affect the ability of the drone to maintain the desired path.
80 The goal is for an autonomous drone to develop a successful strategy to adapt to aerodynamic
81 conditions while accomplishing navigation tasks with a learning approach. To simulate a more
82 realistic scenario, in our investigation, we consider only an onboard camera as the sole sensor for
83 autonomous drones to sense the environment and learn from the errors. The rationale is that
84 there is a latent relationship between the layout of buildings and the distribution of wind zones,
85 and we hope that the proposed MONRL can still learn effective strategies based on only sparse
86 information about the environment (i.e., visual information). During the MONRL training, we
87 implemented curriculum learning to train the model and used a 3D bounding box algorithm to
88 detect the obstacles based on the picture captured by the camera. Curriculum learning is a
89 machine learning technique that involves training a model on increasingly complex or difficult
90 examples over time. The idea is to gradually increase the difficulty of the training data so that the
91 model can learn more efficiently and effectively. Besides, we develop a conversion algorithm to
92 quantify the wind zone distribution around tall buildings based on computational fluid dynamics
93 (CFD) results. After the training, the drone should be able to sense multiple environmental and
94 physical conditions of the near field, and develop dynamic policies for prioritizing the
95 navigational decision to optimize the path while minimizing the negative environmental impact.
96 The remainder of this paper introduces the point of departure of our study, the proposed method,
97 case studies of the MONRL training, and a validation case based in Manhattan, NYC.

98

99 **2. RELATED WORKS**

100 **2.1. Autonomous Drones Navigation**

101 Autonomous drones have gained significant attention due to their ability to operate without
102 direct human intervention [24]. The versatility of autonomous drones has resulted in a wide
103 range of applications across industries, including agriculture, disaster management, infrastructure
104 inspection, and wildlife monitoring [50] [4] [18] [38]. Autonomous drones employ various self-

105 navigation methods, such as GPS-based, vision-based, and sensor-based navigation for path
106 planning and flying controls. GPS-based drone navigation is considered the most common and
107 mature method, which relies on GPS data for navigation through pre-defined waypoints [35]. In
108 most cases, a global map that provides a high-level representation of the navigation mission
109 space is available, allowing the drone to plan its trajectory and avoid obstacles more efficiently
110 [57]. Even for partially unknown and dynamic 3D environments, access to GPS data can still
111 enable hybrid navigation methods to facilitate self-governing drone operations, such as rapid
112 reflexive reactions to newly detected obstacles [20].

113 However, many drone missions are performed in GPS-denied environments characterized
114 by tall buildings or dense plants, limiting the use of GPS-based navigation [64]. As a result,
115 vision-based or sensor-based techniques are explored. Vision-based navigation utilizes computer
116 vision algorithms and onboard cameras to perceive and analyze the environment, enabling
117 drones to avoid obstacles and follow paths [5]. Examples of vision-based navigation systems
118 include Simultaneous Localization and Mapping (SLAM) [66] and Visual-Inertial Odometry
119 (VIO) [61]. Sensor-based navigation techniques leverage sensors like LiDAR, radar, and
120 ultrasonic to measure distances between the drone and surrounding objects, which is useful in
121 environments with limited or impaired visual navigation [13]. These techniques often employ
122 sensor fusion, combining data from multiple sources to enhance the accuracy and robustness of
123 the navigation system [36]. Recently, the use of machine learning has greatly improved the
124 capabilities of autonomous drones in complex environments [37]. For example, deep learning
125 techniques, such as Convolutional Neural Networks (CNNs) [6] and Recurrent Neural Networks
126 (RNNs), have been applied to enhance vision-based navigation by enabling more accurate object
127 recognition, segmentation, and tracking [71]. Despite these advances, challenges remain for
128 autonomous drones' navigation in complex and dynamic environments. Obstacle detection and
129 avoidance in areas such as urban settings or dense forests is critical, but developing effective
130 algorithms remains challenging due to the variety of objects and their proximity [40]. The
131 following section unfolds the details of challenges related to drone navigation in dynamic
132 environments and introduces the existing efforts in addressing the self-navigation in dynamic
133 environments.

134 **2.2. Self-Navigation in Dynamic Environments**

135 One of the primary challenges in autonomous drone navigation is operating in dynamic
136 environments. The challenge roots from the difficulty of environmental sensing to incorporate a
137 large amount of dynamic information and the corresponding information processing and
138 computational demands. On the one hand, dynamic environments with moving objects (e.g.,
139 vehicles and pedestrians) and adverse environmental conditions (e.g., poor lighting conditions)
140 require advanced sensing methods to enable real-time adaptation and accurate movement
141 prediction [41]. On the other hand, with the increasing input of sensor data, robust data
142 processing, and control algorithms are needed to convert the external information about the
143 environment into safe and effective drone navigation and path planning decisions. As for the
144 sensing techniques for autonomous drones, a popular solution is the use of LiDAR-based
145 sensors. Bolognini and Fagiano [13] developed a real-time algorithm for LiDAR-based drone
146 control, enabling navigation in unknown environments. El-Sheimy and Li [19] proposed an
147 innovative algorithm to reduce computational cost and energy consumption for indoor LiDAR
148 target detection. Aldao, González-de Santos and González-Jorge [2] evaluated LiDAR's ability
149 to detect dynamic obstacles' location and speed with positive results. Patrik, Utama, Gunawan,
150 Chowanda, Suroso, Shofiyanti and Budiharto [48] used course-over-ground information and an
151 IMU sensor for outdoor drone navigation without obstacles, relying on GNSS data. In addition,
152 camera-based methods provide a more detailed visual representation of the environment,
153 allowing drones to respond to a broader range of obstacles. Chakravarty, Kelchtermans, Roussel,
154 Wellens, Tuytelaars and Van Eycken [16] used a single front-facing camera and trained a CNN
155 for depth perception, guiding the drone to avoid obstacles. Sani and Karimian [55] proposed a
156 complete solution for indoor drone navigation using two cameras, a Kalman filter, and an inertial
157 sensor. García Carrillo, Dzul López, Lozano and Pégard [25] placed a stereoscopic camera on a
158 drone and trained a CNN to process depth maps, detecting and locating drones up to eight meters
159 away with a 10% error rate. Jung, Hwang, Shin and Shim [32] introduced a new CNN to
160 estimate the locations of key objects in drone navigation, applying a line-of-sight guidance
161 algorithm for drone direction, validated on cost-efficient hardware. Baskar and Gorodetsky [9]
162 proposed a new dataset of wind simulation to validate the fixed-wing UAV route choice when
163 considering the energy efficiency. Nathanael, Wang and Low [47] used steady-state RANS

164 simulation generated by Open FOAM to simulate the wind field in Singapore and then
165 researched its impact on the flight of the Unmanned Aerial Systems (UAS).

166 Literature has also been investigating advanced algorithms capable of processing
167 additional sensor information and functioning effectively in diverse conditions [68]. Sun, Li,
168 Wang and Zhao [62] proposed the use of the Rapidly-exploring Random Tree (RRT) algorithm
169 for path planning in complex environments. RRT has become a popular choice for UAV path
170 planning due to its ability to quickly explore search spaces and generate feasible paths. Later,
171 Wen, Su, Ma, Zhao and Zhang [69] presented an enhanced version of RRT, called Safe-RRT,
172 which guarantees safe and optimal paths for UAVs in cluttered environments. Optimization-
173 based methods have also gained traction for addressing path planning problems. Albert, Leira
174 and Imsland [1] applied mixed-integer linear programming (MILP) to generate collision-free
175 paths for multiple UAVs simultaneously. Their method demonstrated the ability to handle
176 complex scenarios with multiple UAVs and obstacles efficiently. Similarly, Huang, Zhou, Ran,
177 Wang, Chen and Deng [28] investigated the use of particle swarm optimization (PSO) for drone
178 path planning, showing its effectiveness in generating optimal paths in real-time. Recent
179 developments in machine learning have led to new approaches for UAV path planning and
180 navigation. Fan, Chu, Zhang, Song and Li [23] explored the use of imitation learning to transfer
181 expert knowledge for drone navigation, leveraging human demonstration data for improved
182 performance. In addition to path planning algorithms, sensor fusion, and localization play a
183 crucial role in autonomous drone navigation in a dynamic environment. Cadena, Carlone,
184 Carrillo, Latif, Scaramuzza, Neira, Reid and Leonard [14] provided a comprehensive review of
185 simultaneous localization and mapping (SLAM) techniques for drones, emphasizing the
186 importance of integrating vision, inertial, and other sensors for accurate localization.
187 Furthermore, Baca, Hert, Loianno, Saska and Kumar [8] demonstrated the effectiveness of
188 vision-based navigation for UAVs, enabling them to operate in GPS-denied environments.
189 Collision avoidance and obstacle detection are essential components of safe drone navigation.
190 Imrane, Melingui, Mvogo Ahanda, Biya Motto and Merzouki [29] utilized a bio-inspired
191 approach based on the artificial potential field (APF) method for obstacle avoidance, showing its
192 effectiveness in real-world scenarios. Meanwhile, Andersson, Wzorek, Rudol and Doherty [3]
193 developed a predictive collision avoidance algorithm using Bayesian methods, which offered
194 improved performance of drone navigation in dynamic environments.

195 What is worth noting, many of the advanced control algorithms and machine learning
196 techniques for autonomous drone self-navigation require excessive data from the sensing
197 systems about the dynamic environment [63]. We consider a more realistic scenario where not
198 all relevant sensor data is available or very limited, while the primary source of data is still from
199 imagery sensors [40]. For example, for an algorithm that can control a drone through heavy wind
200 zones (aerodynamic optimization), we consider the situation when wind sensors or similar
201 pressure sensors are not accessible, but only onboard RGB cameras are available. The following
202 section introduces the potential of reinforcement learning in tackling issues with only sparse
203 data.

204 **2.3. Reinforcement Learning for Self-Navigation with Limited Data**

205 Reinforcement learning (RL) is a branch of machine learning methods that involves training an
206 agent to make decisions based on interactions with an environment [33]. The agent learns to
207 maximize a numerical reward signal that it receives from the environment in response to its
208 actions. This trial-and-error process allows the agent to discover the optimal policy for achieving
209 its goal [43]. Compared to other machine learning approaches, RL does not need a training set,
210 and the agent will adjust the action based on continuous feedback to maximize the final reward.
211 One key advantage of RL is its ability to learn from experience and adapt to changing
212 environments, which is key to autonomous drone navigation in dynamic environments. Pham et
213 al. used a PID+Q-learning algorithm as the training strategy to train the drone to navigate in an
214 unknown environment without any obstacles [49]. They created a simulation environment for the
215 drone to learn, then tested the navigation system in the real world. The method's validation and
216 simulation yielded comparable outcomes, indicating that the drone could adeptly learn to
217 navigate through its surroundings without relying on any pre-established model. Hodge et al.
218 built a 2D training environment in Unity and tried safely training the drone to arrive at the
219 destination while avoiding all the obstacles [27]. They used LiDAR sensors to obtain the
220 obstacle information near the drone and then combined RL with the Proximal Policy
221 Optimization algorithm to train the drone navigation system. Muñoz et al. used the DDQN to
222 train a drone in the 3D simulation environment and proposed a new neural network named JNN
223 to learn obstacle information based on depth cameras [45]. Shin et al. created a virtual
224 environment for drone training and applied the U-Net as the reinforcement learning network

225 during the training [60]. To test the model’s validity, they transferred it to the real drone, then
226 tested it in a reconfigurable maze that was similar to the virtual maze. They used only one RGB
227 camera as the model’s input, which still performed well in the real-world test. Xue and
228 Gonsalves suggested a novel deep reinforcement learning approach that used Variational
229 Autoencoder (VAE) for image data preprocessing [70]. This approach facilitated fast and
230 effective learning for UAVs to avoid obstacles without depending on sensors other than the
231 depth camera. The advantage of this method over other visual obstacle avoidance algorithms was
232 that it allowed for obstacle avoidance in a continuous action space and eliminated the need for
233 complex image recognition networks during RL training. Tong et al. proposed an enhanced Deep
234 Reinforcement Learning (DRL) approach to address the inherent limitations of UAV navigation
235 in dynamic environments [63]. The proposed method involved a distributed DRL framework that
236 broke down the navigation task into two simpler sub-tasks, each tackled by a Long Short-Term
237 Memory (LSTM) based DRL network using a portion of the interactive data. Additionally, a
238 clipped DRL loss function is introduced to seamlessly integrate the two sub-solutions into a
239 single comprehensive solution for UAV navigation.

240 In summary, RL methods offer significant advantages for autonomous drones in self-
241 navigation, particularly when certain sensor data is missing. These advantages stem from RL’s
242 ability to learn from experience, adapt to dynamic environments, exhibit inherent robustness,
243 generalize knowledge across environments, incrementally learn, and scale with deep learning
244 techniques [46]. Through trial-and-error interactions with the environment, RL enables drones to
245 develop navigation policies that account for missing data. RL’s adaptability ensures effective
246 navigation in dynamic and unpredictable scenarios, even with limited sensor data. The inherent
247 robustness of RL algorithms allows drones to prioritize actions leading to successful navigation,
248 compensating for missing information. Additionally, RL’s generalization capabilities enable
249 drones to apply learned navigation policies to similar environments without extensive retraining
250 [17]. Incremental learning allows RL-based drones to continually refine their navigation policies
251 in response to changing or missing sensor data, while the integration of deep learning techniques,
252 such as Deep Q-Networks (DQNs) and Proximal Policy Optimization (PPO), ensures scalability
253 and the ability to learn complex navigation policies in various environments [7]. Overall, RL
254 provides a powerful solution for autonomous drones to overcome challenges posed by
255 incomplete or missing sensor data and achieve successful self-navigation.

256

257 **2.4. Multi-Objective Reinforcement Learning for Self-Navigation**

258 In recent years, Multi-Objective Reinforcement Learning (MORL) has emerged as a pivotal
259 approach to tackle complex decision-making problems involving multiple conflicting objectives.
260 MORL techniques offer a framework to optimize agents' policies in environments where
261 multiple objectives must be simultaneously considered and balanced. At its core, MORL seeks to
262 find a balance between different objectives by exploring the trade-off space. Algorithms within
263 the MORL framework aim to generate policies representing diverse solutions, each offering a
264 distinct compromise between objectives. This set of policies is known as the Pareto frontier, and
265 it empowers decision-makers to choose the most suitable policy according to their priorities.
266 Ramezani Dooraki and Lee [52] built a multi-objective reinforcement learning network to train
267 the drone navigation system in a dynamic environment, which includes static and moving
268 obstacles. Shantia, et al. [58] proposed a two-stage training methodology for drone navigation
269 based on visual input and proved that the performance for the multi-objective training network
270 was better than the single-objective. MORL addresses scenarios where a single optimal solution
271 may not suffice due to competing goals or trade-offs. This is particularly relevant in our context,
272 where our problem involves optimizing drone navigation while minimizing wind impact and
273 maximizing safety. MORL provides a robust methodology to navigate these challenges and
274 derive policies that lead to a Pareto-optimal set of solutions, allowing decision-makers to make
275 informed choices based on their preferences. By incorporating MORL into our study, we
276 leverage its capabilities to tackle the intricate trade-offs inherent in our multi-objective drone
277 navigation problem. MORL equips our approach with the flexibility to handle conflicting
278 objectives, ensuring that our solution is not only technically robust but also practically adaptable
279 to real-world scenarios.

280 In summary, Multi-Objective Reinforcement Learning offers a well-established framework to
281 address multi-dimensional optimization challenges, making it an ideal choice for our problem of
282 optimizing drone navigation while considering wind fields in urban environments. The adoption
283 of MORL empowers us to explore the trade-off space and present a set of solutions that cater to a
284 spectrum of decision-makers preferences.

285 **3. METHODOLOGY**

286 3.1. Overview

287 This paper proposes a MONRL method for drones to autonomously navigate through dynamic
288 wind zones while reaching the given target and avoiding obstacles only based only on local
289 information obtained from the RGB cameras and GPS equipment. Unlike previous studies that
290 assume that global information about the environment is known, or relevant sensors are available
291 for obtaining all necessary information, our approach assumes that only GPS equipment and
292 limited imagery sensors, such as RGB cameras, are accessible. The dynamic and complex
293 interactions between the built environment and the wind zones would be captured by the system
294 while learning.

295 The framework showcased in **Fig. 1** provides a clear outline of the sequential steps that
296 constitute the proposed Multi-Objective Navigation Reinforcement Learning (MONRL)
297 technique, which is designed to enhance the navigation capabilities of drones. Here's an in-depth
298 breakdown of the process:

299 The first step involves the sensing mechanisms employed by the drone agent. Equipped with
300 two core sensing devices, namely the RGB camera and GPS equipment, the drone gathers crucial
301 information from its environment. The RGB camera can offer a three-dimensional view of the
302 surroundings by applying bounding box algorithms. This facilitates the drone's understanding of
303 the landscape's geometry, helping it gauge the proximity of objects. The GPS equipment, on the
304 other hand, not only provides accurate positional data for navigation but also aids in determining
305 wind force.

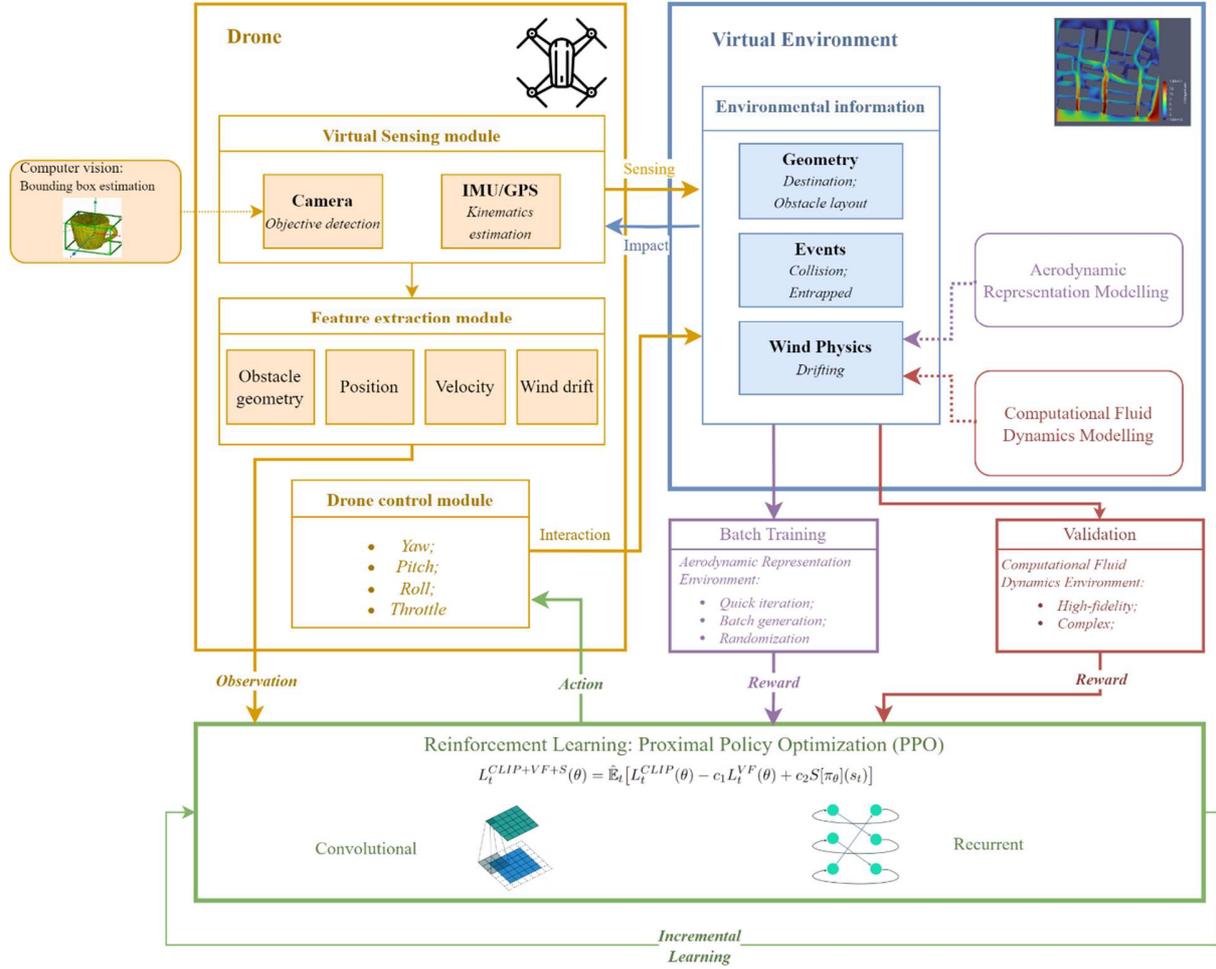
306 In the second step, the drone utilizes the information it collects to estimate wind force
307 through drifting information. When subjected to wind forces, the drone deviates from its
308 intended trajectory. By comparing the anticipated position (based on movement commands) with
309 the actual position (from GPS data), the drone calculates the extent of drifting deviation. This
310 deviation information is then utilized to gauge the wind force acting upon the drone. Analyzing
311 the degree and direction of drift provides valuable insights into the magnitude and orientation of
312 the wind.

313 The third step involves the integration and pre-processing of data. The data acquired from the
314 RGB camera and the estimated wind force from GPS are merged and subjected to preparatory

315 procedures. This ensures that the data is formatted in a manner conducive to its utilization within
316 the subsequent reinforcement learning network. Potential processes here include data
317 normalization and flattening, which streamline the input data.

318 In the fourth step, the integrated data is fed into a complex reinforcement learning network.
319 This network is characterized by its deep architecture, encompassing different layers designed to
320 handle diverse information. Convolutional layers process the depth information from the camera,
321 while dense layers manage drifting and wind force data. The core objective of this network is to
322 learn optimal navigation strategies for the drone across various environments, taking into
323 account obstacles derived from geometry information as well as environmental factors such as
324 wind.

325 The ultimate output of this network is a set of actions or decisions that guide the drone's real-
326 time movement. These actions are carefully formulated based on the drone's assimilation of data
327 and its understanding of the environment. In essence, the MONRL framework enables the drone
328 to navigate intricate landscapes, making informed decisions that optimize its movements by
329 intelligently adapting to the prevailing conditions.



330

331

Fig. 1 Framework of proposed MONRL method

332

333

334

335

336

337

338

339

340

341

342

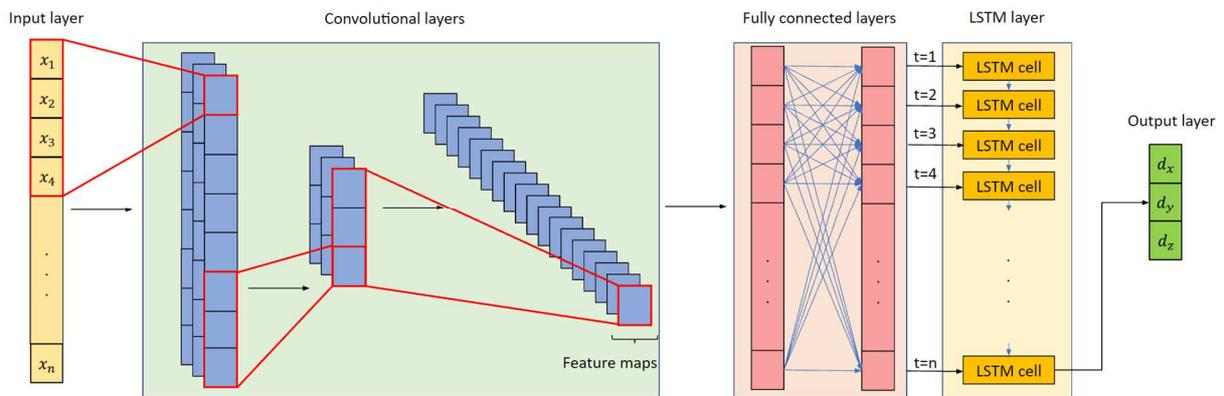
Among these inputs, the external inputs to the drone agent are solely processed camera data. We use a 3D bounding box algorithm to convert the imagery data about the built environments (such as buildings) captured by the front camera to 3D vectors [44]. The training can be divided into two phases. The first phase is to train the drone in the simulation environment with no wind zones by using the Proximal Policy Optimization (PPO) algorithm and the LSTM structure [26]. The PPO algorithm is a RL algorithm that offers advantages such as stability, efficiency, sample efficiency, and robustness, making it a powerful and reliable algorithm for learning good policies with relatively few samples [56]. In addition, LSTM is used to remember information from previous time steps in a sequence. This recurrent model can capture temporal dependencies in the input data and help the drone make more informed decisions. The experiment was performed

343 in multiple phases, gradually adding more environmental factors to the simulation. This is to test
344 how incorporating environmental factors could alter the original decisions made by the drone AI.
345 Specifically, the purpose of the first phase was to validate the feasibility of the RL architecture
346 for identifying the shortest path from the starting point to the target point without considering
347 other conditions, such as environmental factors. It also supports a simulated navigation system
348 that grants the drone an obstacle avoidance ability in windless conditions. In the second phase,
349 after the drone’s basic navigation and obstacle avoidance system was built up (i.e., phase one),
350 the wind information was added to examine how MONRL could facilitate navigation decision-
351 making in an environment with more conditions, in our study, dynamic winds. Therefore, we
352 added a simulation of the wind dynamics to the MONRL training during the second phase, i.e.,
353 adding aerodynamic forces in the simulated environments to replace the windless condition in
354 the prior phase, especially when the drone approaches buildings. As will be shown later in the
355 results, we found that incorporating environmental factors indeed changes the decision of the AI.
356 Such a comparison is only possible with a multi-phase study. Because MONRL in our study is a
357 multi-objective training, i.e., reaching the target, avoiding obstacles, and counteract the negative
358 effects of dynamic wind zones on the gesture and path stability of the drone, a multivariate
359 reward function is utilized for training the drone to learn to dodge the strong wind zones, react to
360 milder wind zones, avoid visible obstacles, in addition to finding a new route to the target when
361 needed. After the two training phases, real-world data is used to validate the training results. We
362 use OpenFOAM [30], an open-source computational fluid dynamics (CFD) software, to create
363 the CFD results for a metropolitan district in New York City. Once the CFD result is generated,
364 the data is transferred to Unity and used as the wind zone input in the test environment.

365 **3.2. Deep Reinforcement Learning Network**

366 The core of the proposed MONRL method is a deep reinforcement learning (DRL) network that
367 enables the robust learning of the drone to navigate in a dynamic environment. We designed a
368 DRL network with seven layers, as shown in **Fig. 2**. The first layer is the input layer with 87
369 inputs. Among the 87 inputs, 72 inputs are assigned to store the local geometric information. In
370 our case, it refers to the bounding box results from the drone-carried imagery sensor to represent
371 obstacles (mainly buildings) in the visible proximity of the drone. For each obstacle, the 3D
372 bounding box algorithm can return the 8 space positions of all its vertex, and each space position

373 contains three values, which represent the XYZ position data, respectively. In this way, we can
 374 use a short array to represent the key geometric features of a building (i.e., the contour of a
 375 building). The other 12 inputs are used to record all the information related to the drone agent.
 376 The information includes the distance between the drone and the target, the drone's current speed,
 377 the drone's position, and wind data from a given aerodynamic map. Each of the inputs is a vector
 378 that contains three values. To teach the drone to avoid the stronger wind zones, we allocated
 379 three inputs to enable the drone to memorize the position of the last strong wind zone that can be
 380 used in the reward function.



381

382 **Fig. 2** Deep reinforcement learning network architecture

383

384 Moving forward from the input layer, there are two convolutional layers. Since the vertex
 385 information of a bounding box is interdependent, using convolutional layers helps extract the key
 386 features. Then there are 2 hidden layers, each with 256 nodes. These layers allow the network to
 387 capture more complex relationships between the input features. A single LSTM layer is used
 388 after the hidden layers. The advantage of this recurrent layer is that it can effectively capture
 389 long-term dependencies and temporal patterns in sequential data, which can let the drone avoid
 390 an infinite loop. The final layer is the outcome layer with three outputs. The output signifies the
 391 intended direction for the drone's movement along the x, y, and z axes. This vector comprises
 392 three values spanning from 0 to 1. Given its adjustable magnitude, the drone's behavior can be
 393 altered by modifying this magnitude. The ultimate velocity will be determined by the product of
 394 the output and the maximum speed of the simulated drone during the simulation.

395 In the simulation, the wind force would directly affect the drone’s acceleration rate.
 396 Therefore, the agent must consider the movement decision based on the wind information. In our
 397 experiment, we assumed that the drone could fully follow the commands from the RL algorithm.
 398 In addition, our model did not consider the modeling of sensor noises other than the visual
 399 sensors, i.e., the 3D bounding box recognition based on the virtual visual sensors. We admit that
 400 this is a simplified assumption as in real world, drone controls could be affected by a variety of
 401 factors and thus, the control could be off. Nonetheless, the scope of this study is to examine that
 402 if the incorporation of multiple objectives at the same time, especially the aerodynamic forces, in
 403 the drone navigation decisions, could yield paths that are different than a direct consideration of
 404 shortest path. As a result, we focused on a more controlled experiment in the simulator while
 405 ignoring the real-world control challenges. It is recognized as one of the limitations in the
 406 discussion.

407 In general, its decision network can be defined as Eq (1):

$$408 \quad D_t = F(b_t, d_t, v_t, w_t, p_t, D_{t-1}) \quad \text{Eq (1)}$$

409 Where D_t is the speed decision for this step, b_t represents the obstacle data (relative
 410 position, dimensions) from the 3D bounding box algorithm, d_t represents the relative distance in
 411 all directions between the agent and the target, v_t represents the current speed of the drone, w_t is
 412 the wind data based on the drifting of the drone, p_t refers to the local position of the drone,
 413 which is used to control the drone fly in the designed scope, D_{t-1} is the movement decision for
 414 the last step, which is the velocity.

415 3.3. Reward Function

416 According to the proposed MONRL, the goal of the agent is to learn a policy that maximizes its
 417 cumulative reward with a series of sub-objectives over time by taking actions in an environment
 418 and receiving feedback in the form of rewards. The reward function design is based on the
 419 project goal, previous experience, and the drone’s performance during the experiment. The
 420 reward system that we designed in this experiment is as Eq (2).

$$421 \quad R = r_{distance} + r_{target} + r_{wind} + r_{collision} + r_{stuck} + r_{timeout} + r_{time} \quad \text{Eq (2)}$$

422 Where R represents the total reward that the drone can receive. It is a linear combination
 423 of a set of award terms: $r_{distance}$ represents the reward depending on how close the drone moves
 424 towards the target, r_{target} is a binary reward if the drone safely arrives at the target, r_{wind} refers
 425 to the reward if the drone successfully avoids the strong wind zone, $r_{collision}$ represents the
 426 cumulative penalty when the drone has a collision with any obstacle, r_{stuck} is a penalty when the
 427 drone is stuck at an infinite loop in navigation, r_{time} is a penalty term based on the total time
 428 spent on the navigation task, and $r_{timeout}$ is a binary term to represent penalty when the drone
 429 uses out all the time. To simplify, we did not assign any weight to the sub-objectives. It deserves
 430 future investigations into how changing the priorities of each of the sub-objectives would affect
 431 the overall performance of the MONRL in drone navigation. The followings provide more
 432 detailed information about each of the sub-award terms.

433 The first reward term encourages the drone to move towards the target. It is a continuous
 434 reward that can guide the drone at every step in one epoch. In our experiment design, the
 435 positions of the drone and target change every epoch. In order to normalize the distance reward
 436 for every epoch, we design the distance reward function as Eq (3).

$$437 \quad r_{distance} = \sum_1^T \frac{dis_t - dis_{t-1}}{dis_0} \times 100 \quad \text{Eq (3)}$$

438 Where $r_{distance}$ represents the rewards the drone can receive, dis_t is the distance
 439 between the drone and the target at this step, dis_{t-1} is the distance between the drone and the
 440 target at the last step, dis_0 is a fixed value that represents the distance between the drone and the
 441 target at the start of the current epoch.

442 The second reward term provides a positive reward for reaching the target quickly and
 443 safely while avoiding stronger wind zones and obstacles, as shown in Eq (4). This can be a
 444 comprehensive approach to incentivize the drone to learn to navigate effectively in the wind
 445 zone.

$$446 \quad r_{target} = \begin{cases} +10 & \text{if arrives target} \\ 0 & \text{otherwise} \end{cases} \quad \text{Eq (4)}$$

447 The third reward term encourages the drone to find alternative routes to avoid the
 448 stronger wind zones. For every epoch, if the drone encounters a strong wind zone and avoids it, it
 449 will receive a positive reward, as indicated in Eq (5).

$$450 \quad r_{wind} = \begin{cases} +10 & \text{if avoids strong wind zone} \\ 0 & \text{otherwise} \end{cases} \quad \text{Eq (5)}$$

451 The fourth reward function is a negative term that penalizes the drone once there is a
 452 collision. In the real world, collisions with obstacles can be disastrous for the drone. Besides, the
 453 epoch will stop immediately if a collision occurs, thus the drone will not receive any further
 454 rewards and will receive a very low total reward. By providing a significant penalty for
 455 collisions, the drone will learn to avoid them (see Eq (6)).

$$456 \quad r_{collision} = \begin{cases} -10 & \text{if collision occurs} \\ 0 & \text{otherwise} \end{cases} \quad \text{Eq (6)}$$

457 The fifth reward term is a penalty for the drone flying around a strong wind zone, as
 458 shown in Eq (7). It is a negative reward for every step the drone spends in a strong wind zone.
 459 This will encourage the drone to minimize the time it spends in those areas and encourage it to
 460 find a route around them. We set this function as a linear function. The closer the drone is to a
 461 wind zone, the larger the penalty it will have. This will encourage the drone to move away from
 462 the strong wind zones.

$$463 \quad r_{stuck} = \begin{cases} -0.1 \times \frac{dis_{wind}}{20} \text{ per step} & \text{if in strong wind zone} \\ 0 & \text{otherwise} \end{cases} \quad \text{Eq (7)}$$

464 Where r_{stuck} represents the penalty the drone will receive, dis_{wind} is the distance
 465 between the drone and the stored strong wind position.

466 The sixth reward term is a penalty if the drone uses all the time but does not arrive at the
 467 target, as shown in Eq (8). Battery consumption is vital to drones as they rely on onboard power
 468 supplies to power their propulsion systems, sensors, cameras, and other components. If the drone
 469 spends too much energy on a one-trip flight, it may lose power and be unable to complete its
 470 tasks or return to the start location. Since it does not complete the tasks, a negative penalty is
 471 given to the drone.

$$r_{timeout} = \begin{cases} -10 & \text{if step reaches maximum} \\ 0 & \text{otherwise} \end{cases} \quad \text{Eq (8)}$$

Finally, the last reward function is the time penalty function for the drone. As shown in Eq (9), the penalty increases as the time spent on the mission increase. The purpose of this penalty is to let the drone choose the fastest way to the target.

$$r_{time} = 0.001 \text{ per step} \quad \text{Eq (9)}$$

3.4. 3D Bounding Box Algorithm

It should be noted the advancements in stereoscopic vision systems offer extended range and improved computational efficiency. The compactness and enhanced capabilities of systems like the Intel RealSense indeed make them suitable for a variety of applications, including construction surveying. The compactness and enhanced capabilities of systems like the Intel RealSense indeed make them suitable for a variety of applications. In our study, we primarily relied on RGB information for drone navigation not only due to its technical feasibility but also its cost viability. RGB cameras, widely used in various applications, are generally more cost-effective compared to depth cameras. While advanced depth cameras, including models like the RealSense, offer superior capabilities, they often come at a higher cost. This factor is particularly crucial in the selection of sensory equipment for drones, especially for large-scale or budget-sensitive projects. Therefore, we use RGB camera combined with 3D bounding box to do the obstacle detection tasks. A 3D bounding box is a rectangular prism that surrounds a 3D object, such as a building. The bounding box is defined by its three dimensions - length, width, and height - as well as its position and orientation in space. 3D bounding boxes are commonly used in computer vision and robotics applications to represent the spatial extent of objects in a scene [44]. In general, for a 3D box projected on the image, we can always wrap it with a 2D box. Mousavian et al. proposed a deep learning model to predict the 3D bounding box position based on the 2D bounding box results [44]. They built up constraints between the 2D and 3D bounding box images, called point-to-side correspondence constraint equations. One of the equations is as follows.

498

$$x_{min} = \begin{pmatrix} \frac{d_x}{2} \\ K[R T]^{-1} \frac{d_y}{2} \\ \frac{d_z}{2} \\ 1 \end{pmatrix} \quad \text{Eq (10)}$$

499

500

501

502

Where x_{min} represents the minimum x position of the 2D bounding box, K represents the camera intrinsic matrix, R represents the frame rotation, T represents the frame translation, d_x is the length of the 3D bounding box, d_y is the width of the 3D bounding box, d_z is the height of the 3D bounding box.

503

504

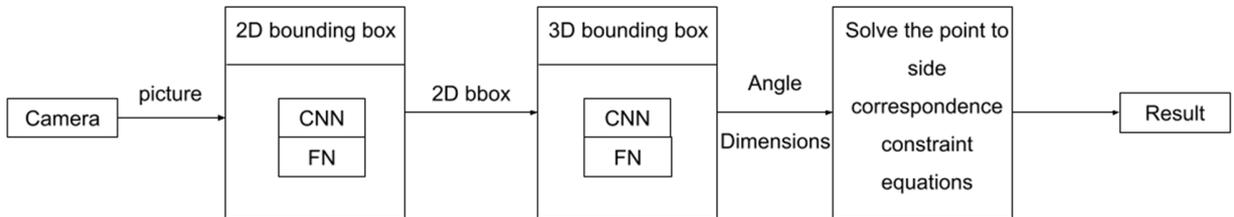
505

506

507

508

Because there are four vertex points in a 2D bounding box, which are x_{min} , x_{max} , y_{min} , y_{max} , a total of four equations of the above form can be written. This is the set of four constraint equations for the 2D bounding box and the 3D bounding box. And there are nine parameters to be solved in this set of equations (six affine transformation parameters and three bounding box dimensions). In our work, we built the 3D bounding box algorithm based on this model. The details of the algorithm are shown in **Fig. 3**:



509

510

Fig. 3 Bounding box algorithm

511

512

513

514

515

516

517

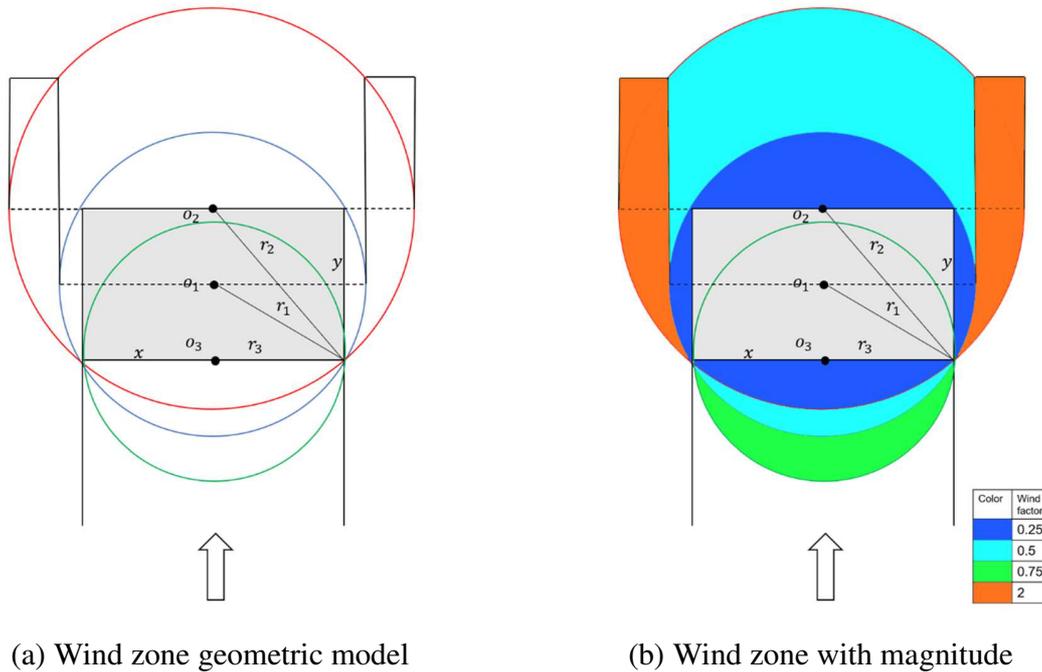
518

The camera on the drone will provide a picture of all the obstacles in front of the drone. In our design, we limit the number of obstacles used in the model to simulate the limited data storage and processing capacities of most drones and stress test the performance of the proposed MONRL method given only sparse data. Specifically, if the number of obstacles captured by the sensor is larger than three, we only consider the three closest obstacles. Once a drone determines any obstacle, we will first apply the 2D bounding box algorithm to compress the raw data. A 2D detector is used to find the bounding box for the obstacles. Then, its results are used as the input

519 of the 3D bounding box network. This network can convert the convolutional features to the
520 angle and dimensions of the obstacles, which are parameters in the four constraint equations.
521 Finally, since we have already known six of the nine parameters, it is trivial for us to obtain the
522 rest three parameters from the four constraint equations, which is the relative position of the 3D
523 bounding box. By combining the dimensions and relative position of the obstacles, we can
524 calculate all eight vertex positions of the obstacles and transfer them to the reinforcement
525 learning model.

526 **3.5. Aerodynamic Representation Modeling**

527 We also design a framework for representing the aerodynamic characteristics of an environment.
528 Although CFD can be used to simulate aerodynamic forces, we found it to be inefficient for
529 meeting the needs of a large amount of simulation for the DRL algorithm. DRL usually requires
530 thousands of simulated scenarios with changing conditions. The cost of performing CFD for each
531 of them is significant. In contrast, we found that there were modellable aerodynamic features
532 around a given structure if the initial conditions of the wind and the geometric features of the
533 structure were known. Our proposed representation model can generate a wide range of wind
534 conditions, including random gusts and turbulences. This allows the drone to be trained in a
535 variety of changing conditions, making it more adaptable to different weather scenarios. For
536 example, we can quickly magnify the wind factor in our work to create a strong wind zone
537 between two close buildings. This will make the training environment more suitable for our goal.
538 In addition, representation model simulation typically requires less computational time than CFD,
539 which can be important for training large-scale drone fleets. This allows for faster DRL training
540 and more efficient use of computing resources. During the training period, we change the
541 obstacle layout every epoch.



(a) Wind zone geometric model

(b) Wind zone with magnitude

542 **Fig. 4** Representation modeling for simplifying aerodynamic features around building

543

544 **Fig. 4** shows the details of the proposed aerodynamic representation modeling method
 545 based on the key features extracted from actual CFD results [31]. Different colors represent
 546 different wind factors. This modeling method aims to simulate the wind distribution around a
 547 single building in the wind zone. Assume that we have a single obstacle in front of the wind
 548 direction in the x-y plane. We set the obstacle's length as x and the obstacle's width as y. Then
 549 we need to determine circles with O_1 , O_2 and O_3 as the center and r_1 , r_2 and r_3 as the radius
 550 respectively, where O_1 is the center of the obstacle, O_2 is the middle point of the upper line, and
 551 O_3 is the middle point of the under line. Eq (11) – Eq (13) indicates the quantification of the
 552 geometric models of a wind zone.

553
$$r_1 = 0.5 \times \sqrt{x^2 + y^2} \quad \text{Eq (11)}$$

554
$$r_2 = \sqrt{0.25 \times x^2 + y^2} \quad \text{Eq (12)}$$

555
$$r_3 = 0.5 \times x \quad \text{Eq (13)}$$

556 After completing these three circles, we can also identify two tangents of the blue circle
557 and find their intersection with the red circle. Finally, we divide the whole area around the
558 obstacle into three regions. The first region is in front of the obstacle (front region), which
559 contains three subregions. In the front region, the wind force is reduced due to obstructions.
560 Based on the CFD results, we set all the wind directions in this area toward the closest corner of
561 the obstacle. The second region is the sides of the obstacle (side region), where the valley effect
562 may happen, i.e., the wind speed will increase since air cannot pile up in a large amount in a
563 small space [12]. According to CFD results, the wind magnitude in this region is related to the
564 distance between the wind vector and the obstacle. When the wind is closer to the obstacle in this
565 region, its magnitude will be slower. Otherwise, the flow of the current will increase. As a result,
566 the valley effect can occur if two buildings are very close, causing disturbances to the drone.
567 Besides, the direction of the wind is parallel with the direction of the tangent of the red arc. The
568 last region is behind the obstacle (back region). Similar to the front region, the obstructions cause
569 a decrease in the strength of the wind. In summary, the aerodynamic force on the drone can be
570 written as Eq (14) and Eq (15).

$$571 \quad F_{drone} = F_{wind} \times f_{wind} \quad \text{Eq (14)}$$

$$572 \quad a_{drone} = \frac{F_{drone}}{w_{drone}} \quad \text{Eq (15)}$$

573 Where F_{drone} represents the aerodynamic force exerted on the drone, F_{wind} is the origin
574 aerodynamic force, f_{wind} is the wind factor, which depends on the position of the drone relative
575 to the obstacle, a_{drone} is the passive acceleration of the drone, w_{drone} is the weight of the drone.
576 It is important to note that this aerodynamic representation modeling framework is a simplified
577 approximation of the precise CFD results. It may not function well in applications where
578 accurate and precise high-fidelity CFD results are needed. However, the primary purpose of this
579 research is to train the MONRL model to guide the drone to avoid stronger wind zones and
580 counteract the milder winds. And the main function of this representation framework is to
581 provide high-level, low-resolution aerodynamic information for training only. In our test case,
582 we found that such a simplified representation of the actual aerodynamic features worked well in
583 yielding effective MONRL training.

584 3.6. Proximal Policy Optimization (PPO)

585 In our architecture, we use Proximal Policy Optimization (PPO) to improve the stability and
586 sample efficiency of policy gradient methods by constraining the size of the policy update in
587 each iteration. PPO is an algorithm used in DRL that has gained significant attention in recent
588 years due to its ability to achieve state-of-the-art performance on a wide range of DRL problems
589 [56]. PPO is a model-free, on-policy DRL algorithm that learns directly from experience using a
590 neural network to represent the policy function. The PPO algorithm aims to optimize a policy
591 function $\pi(a|s)$, which represents the probability of taking an action a in state s . The goal is to
592 maximize the expected cumulative reward from time step 0 to T , which is defined as:

$$593 \quad J(\theta) = E[\sum_{t=0}^T \gamma^t r_t] \quad \text{Eq (16)}$$

594 where θ represents the parameters of the policy function, 0^T represents the initial time
595 step in the problem, T is the time horizon, γ is the discount factor, and r_t is the reward received
596 at time step t . It employs a clipped surrogate objective function to update the policy, which helps
597 to ensure stability and prevent catastrophic updates. This function is given by:

$$598 \quad L(\theta) = E \left[\min \left(\frac{r_t(\theta)}{r_{t_{old}}(\theta)} \times A_t, \text{clip} \left(\frac{r_t(\theta)}{r_{t_{old}}(\theta)}, 1 - \varepsilon, 1 + \varepsilon \right) \times A_t \right) \right] \quad \text{Eq (17)}$$

599 where $r_t(\theta)$ is the probability of taking the action that was taken at time step t under the
600 updated policy, $r_{t_{old}}(\theta)$ is the probability of taking that same action under the old policy, A_t is
601 the advantage function, and ε is a hyperparameter that controls the size of the clip. The
602 advantage function is defined as:

$$603 \quad A_t = \sum_{k=t}^T \gamma^k (r_k + 1 + \gamma V(s_{k+1}) - V(s_k)) \quad \text{Eq (18)}$$

604 where $V(s)$ is the value function, which estimates the expected cumulative reward from
605 state s onwards, k is a time index that refers to the current time step in the advantage function
606 calculation. By using stochastic gradient descent, the PPO algorithm updates the policy and
607 value function parameters, with the clipped surrogate objective as the loss function. The value
608 function is updated using the mean squared error loss between the estimated value and the actual
609 reward. Overall, the PPO algorithm is designed to balance exploration and exploitation while
610 ensuring stability and avoiding catastrophic updates.

611 3.7. LSTM for Capturing Temporal Dependencies

612 In our pilot test, we found that the drone would be easily stuck in the local optimum without
613 memory from previous epochs. This is probably due to the complexity of the environment and
614 the high dimensionality in the DRL input. As a result, we use Long Short-Term Memory (LSTM)
615 to help the drone remember the previous action to avoid getting stuck in a dead-end cycle. LSTM
616 is a type of recurrent neural network (RNN) that is designed to overcome the vanishing gradient
617 problem, which is a common issue with traditional RNNs [26]. It uses a specialized architecture
618 that includes memory cells, input gates, output gates, and forget gates. At each time step t , the
619 network takes an input vector x_t and produces an output vector h_t . The network also maintains
620 an internal state vector c_t , which represents the memory of the network at time step t . The key to
621 LSTM's ability to model long-term dependencies is its memory cell, which is denoted by the
622 variable c_t . The memory cell is updated at each time step using the following equations:

$$623 \quad f_t = \sigma(W_f [h_{t-1}, x_t] + b_f) \quad \text{Eq (19)}$$

$$624 \quad i_t = \sigma(W_i [h_{t-1}, x_t] + b_i) \quad \text{Eq (20)}$$

$$625 \quad o_t = \sigma(W_o [h_{t-1}, x_t] + b_o) \quad \text{Eq (21)}$$

$$626 \quad g_t = \tanh(W_c [h_{t-1}, x_t] + b_c) \quad \text{Eq (22)}$$

627 where σ denotes the sigmoid function, \tanh denotes the hyperbolic tangent function, W_f ,
628 W_i , W_o , and W_c are weight matrices, b_f , b_i , b_o , and b_c are bias vectors, and $[h_{t-1}, x_t]$ denotes
629 the concatenation of the previous output h_{t-1} and the current input x_t . These equations define
630 four gates that control the flow of information into and out of the memory cell. The forget gate f_t
631 determines how much of the previous memory cell value to retain. The input gate i_t determines
632 how much new information to add to the memory cell. The cell gate g_t computes a candidate
633 memory cell value based on the current input and the previous output. The input gate i_t also
634 determines how much of the candidate value from g_t to add to the memory cell and the output
635 gate o_t controls how much of the current memory cell value to output, which is given by:

$$636 \quad c_t = f_t * c_{t-1} + i_t * g_t \quad \text{Eq (23)}$$

$$637 \quad h_t = o_t * \tanh(c_t) \quad \text{Eq (24)}$$

638 where * denotes element-wise multiplication. These gates control the flow of information
639 into and out of the memory cells, allowing the network to remember or forget information from
640 previous time steps selectively. This allows LSTMs to learn long-term dependencies and to
641 model sequential data effectively.

642

643 **4. EXPERIMENT**

644 We used a simulator to train drones to navigate and avoid obstacles that offers several
645 advantages over traditional physical training methods. The simulator allows for the creation of
646 complex and varied training scenarios that may be difficult or impractical to replicate in the real
647 world. This enables the drone to encounter a wide range of obstacles and challenges, allowing
648 for more comprehensive and effective training. The simulated environments also provide a safe
649 and cost-effective way to train drones. In traditional physical training methods, crashes or other
650 accidents can be costly and time-consuming to repair, and may even pose a risk to human
651 operators. In contrast, simulated crashes or other incidents in a virtual environment pose no risk
652 to humans or property, allowing for more extensive and realistic training without the associated
653 costs and risks. Finally, the simulated environments allow for more efficient and iterative
654 training processes. Training algorithms and parameters can be quickly adjusted and refined, and
655 the drone's performance can be evaluated and analyzed in real-time. This enables faster and
656 more effective learning, leading to more skilled and capable drones in a shorter amount of time.
657 We chose Unity 3D as the training platform for the drone, given its capability in creating a scaled
658 simulation environment with analytical functions for reading and generating external data from
659 the experiment [67].

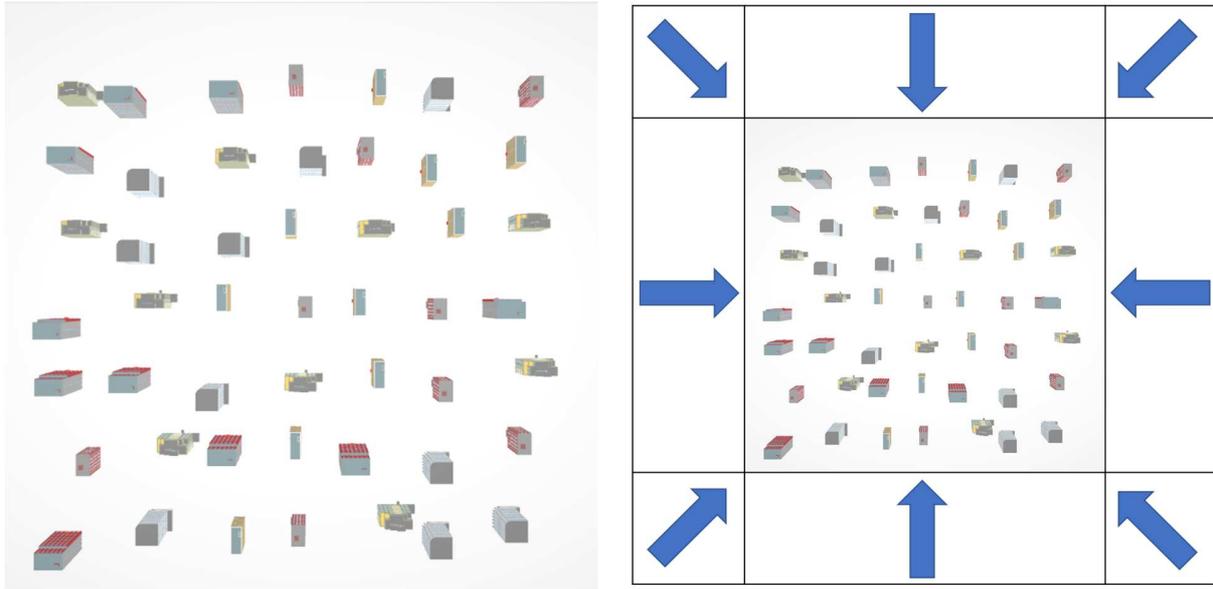
660 **4.1. Simulated Environments**

661 The simulated site represented a 3000 × 3000 feet district in an urban area. We randomly placed
662 49 tall buildings in the simulated environment, with the width and length ranging from 130 to
663 260 feet and height ranging from 200 to 400 feet. Besides, we did not set any physical boundary
664 for the simulated site. Thus, the wind could pass through the site. For each epoch, the drone was
665 randomly spawned at a location on one side of the site, while the target point was randomly
666 placed on the other side. The drone must safely cross the whole site to reach the target. We

667 allowed 3D navigation for the drone. As such, the drone could adjust its altitude. However, in
668 order to prevent the drone from always flying over the buildings, we set a height limitation of
669 160 feet for the drone. This ensured that the drone learned strategies to fly through the buildings
670 and dynamic wind zones instead of flying over them. As a result, the total dimension of the main
671 scenario is a $3000 \times 3000 \times 160$ feet cube.

672 As for the wind simulation, we set the initial wind speed at 11 mph. Using the
673 aerodynamic representation framework as discussed earlier, the wind speed during the training
674 was set to change from 3 to 45 mph, which covered the most common real-life situations.
675 Another important parameter in the wind simulation is the initial wind direction. Given that we
676 focused on the drone's navigation on the X-Y plane, we simulated eight initial wind directions
677 shown in **Fig. 5**. For each epoch, the drone agent would experience one of the eight wind
678 directions or windless conditions randomly, adding variability to the simulated environment.
679 Depending on the flying direction of the drone and the initial wind direction, the wind should
680 have exerted different effects on the drone. For example, if the wind direction is parallel with the
681 drone's flight direction, the drone's speed would increase. A faster speed can be challenging for
682 automatic navigation control, as there would be a short window for collision avoidance. Another
683 challenging situation is the side wind, i.e., the wind direction and the drone's flight direction are
684 perpendicular. In this case, it is nontrivial for the control algorithm to maintain the desired path
685 or for stabilization controls. The goal of MONRL is to teach the drone to recognize the hidden
686 patterns of the effect of the wind and spontaneously generate optimal solutions to offset its
687 impact.

688 Regarding the simulation of GPS functionality in our experiment, we utilized Unity's
689 local coordinate system to relay position information to the drone. The main objective of the
690 GPS equipment in this experiment was to provide the local position information of both the
691 drone and its target. It's important to note that our experimental design assumed the drone was
692 equipped to access only local positional information. Consequently, there was no requirement for
693 global positioning data in this context. This approach allowed us to focus on the drone's ability
694 to navigate and interact effectively within its immediate environment, based on the local
695 information provided.



(a) Scenario design

(b) Wind direction design

Fig. 5 Example of simulated environments and drone

696

697

698 **4.2. Drone Simulation**

699 To closely emulate the realistic behaviors of a drone in our simulation, it is important to ensure
 700 that all parameters utilized in the drone simulation accurately reflect those in the real world. In
 701 this simulation, our selection is the DJI Mavic 2 Pro drone—an esteemed creation by DJI, a
 702 renowned pioneer in unmanned aerial vehicles (drones) and camera stabilization systems. This
 703 drone boasts sophisticated GPS technology that guarantees meticulous positioning and
 704 navigation. This pivotal attribute augments flight stability while enabling intelligent flight modes
 705 that simplify the capture of intricate shots. Drawing from its official specifications, the drone
 706 attains a top speed of 45 mph and achieves a maximum flight altitude of approximately 1640 feet.
 707 However, given the considerations pertaining to urban environments and air traffic control,
 708 prudent restrictions are placed on the drone’s height and velocity. Considering the foregoing, the
 709 drone’s specifications for the simulation have been set as follows:

- 710 • Weight: 2 lbs
- 711 • Maximum speed: 30 mph
- 712 • Maximum altitude: 160 feet

713

714 **4.3. Curriculum Learning**

715 We utilized curriculum learning to help the drone learn complex missions, such as detouring the
716 strong wind zone. Curriculum learning is a popular and effective technique in machine learning
717 that aims to train models by gradually increasing the complexity of the tasks they need to
718 perform [10]. The basic idea behind curriculum learning is to present the model with a sequence
719 of tasks that start off relatively easy and become progressively more difficult over time. By
720 following this curriculum, the model can learn more efficiently and effectively while avoiding
721 getting stuck in sub-optimal solutions. There are various metrics available to decide the
722 appropriate number of epochs required for training each lesson, which may include loss, entropy,
723 or mean final reward. This approach is inspired by the way humans learn, starting with simple
724 concepts and gradually building upon them to tackle more complex ideas. In our work, there is a
725 multi-objective task for the drone, which includes navigation, dodging the obstacles, controlling
726 the drone in the wind zone, and detouring when there is a strong wind zone. It is hard for the
727 drone to learn all these things in a single training environment. Therefore, we applied the
728 curriculum learning strategy in our work and divided the whole training process into four phases
729 as follows.

730 In the first phase, the primary purpose was to train the drone to learn to navigate and
731 dodge all obstacles. For the environment setting, we only placed building obstacles on the
732 scenario with no wind simulation. For every epoch, the positions of the drone and target were
733 randomly generated. Besides, all the obstacles were also randomly distributed in the scenario.
734 The purpose of this setting was that we did not want the drone to remember the obstacles'
735 location. During the training, if a collision was encountered or the drone flew out of the scenario,
736 a negative reward was given to the drone, stopping the epoch. In addition, if the step reached the
737 maximum number, the drone would be assigned a timeout penalty. A limit of 15,000 steps was
738 used for every epoch. Once the success rate reached a high number, we moved to the next phase.

739 In the second phase, since the drone was able to navigate and dodge the obstacles in a no-
740 wind environment after phase 1, a wind zone was added to the training environment. The main
741 goal of this phase was to train the agent to learn to control the drone in a wind zone environment.
742 Except for the wind simulation, all other settings were the same as phase 1. For every epoch, the

743 drone agent would choose one direction as the wind direction, and then this direction would be
744 used to calculate the simulation wind zone around each building.

745 After the first two phases of training, the drone was able to fly safely in an environment
746 with a random wind zone. However, in some cases, the drone could still be caught in an infinite
747 loop and be stuck at a specific position due to the strong wind zone’s existence and the
748 complexity of the layout. In order to allow the drone to learn to detour, a more dynamic training
749 environment was required. In the third phase, we changed the wind direction choices to make
750 sure all the initial wind direction was opposite to the drone’s flight direction. In addition, we
751 developed a new training strategy for the training in this phase. Unlike the previous approach,
752 where an epoch stopped if the step reached the maximum in previous phases, we retained all the
753 environmental factors and the drone status as in the last epoch. This strategy, combined with a
754 specific penalty function r_{stuck} , aimed to enable the drone to overcome its inclination to local
755 traps and find a solution for an infinite loop situation.

756 The final phase combined all the previous training features to improve the model’s
757 performance. One of the advantages of curriculum learning is that the system does not discard
758 previously learned instances. Instead, each lesson or training criterion generates a unique set of
759 weights that builds upon the previous ones during the training process. However, the sub-training
760 criterion can add too much weight to the network. In that case, the whole model will be
761 overfitted for a specific task and unsuitable for the multi-objective model. During this phase, the
762 model would be retrained to seek better performance in the comprehensive environment. In
763 addition, a variable learning rate was used to overcome issues with local optima and saddle
764 points, which helped escape from the trap.

765 In the reinforcement learning segment of our research, we utilized the ML-Agents toolkit as
766 our platform. ML-Agents, which stands for Machine Learning Agents, is a toolkit provided by
767 Unity. Its primary purpose is to transform games and simulations into conducive environments
768 for training machine learning algorithms. This bridges the gap between Unity’s game
769 development environment and the world of machine learning, thereby paving the way for the
770 development of smarter and more interactive applications. **Table. 1** presents the settings we used
771 for reinforcement learning within ML-Agents.

Parameter	Value
Max Steps	1,000,000,000
Time Horizon	1,000
Summary Frequency	100,000
Keep Checkpoints	10
Hyperparameter	
Batch Size	2,048
Buffer Size	20,480
Learning Rate	0.0003
Beta	0.01
Epsilon	0.2
Lambda	0.95
Num Epoch	3
Learning Rate Sch	Linear
Beta Schedule	Constant
Epsilon Schedule	Linear
Network Setting	
Normalize	True
Hidden Units	256
Number of Layers	2
Visualization Type	Simple
Memory Size	256
Sequence Length	64
Reward Signal	
Extrinsic Gamma	0.99
Extrinsic Strength	1.0

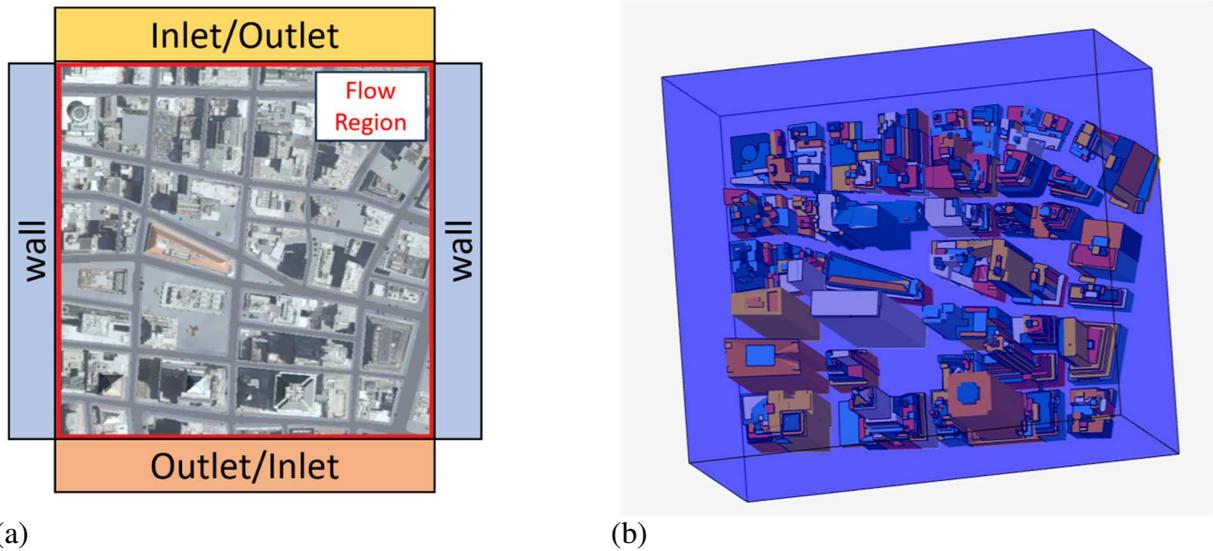
772

Table. 1 Details of ML-Agents setting

773

4.4. Computational Fluid Dynamics Setting

774 Our study generated the Computational Fluid Dynamics (CFD) data to evaluate our algorithm's
 775 performance rigorously. We set up the CFD simulations using the OpenFOAM software package
 776 to provide a brief overview. We defined the simulation domain, boundary conditions, and
 777 relevant physical parameters to the model, as shown in **Fig. 6**.



778 **Fig. 6** Details of CFD setting

Aspect/Parameter	Details/Choice
Material in Flow	Air
Air Density	1.196 kg/m ³
Wind Situations	Four different situations with two wind directions
Wind Speed (Inlet)	5 m/s and 10 m/s
Air Pressure (Outlet)	0.015 kPa and 0.0625 kPa
3D Domain for CFD	As shown in Fig. 6(b)
OpenFOAM Solver	simpleFoam
Meshing Generation	BlockMesh, snappyHexMesh
Grid Convergence Study	Yes

779 **Table. 2** Details of OpenFOAM setting

780 **Fig. 6(a)** shows the domain of the fluid and the boundary conditions of the CFD setting
 781 environment. The material in the flow region is air, whose density is assigned to 1.196 kg/m³. In

782 the validation case, we assume four different wind situations with two different wind directions.
 783 For each wind direction, we assign the wind speed at the inlet as 5 m/s and 10 m/s, and the air
 784 pressure at the outlet is 0.015 kPa and 0.0625 kPa, respectively. **Fig. 6(b)** shows the 3D domain
 785 for the CFD calculation. **Table. 2** lists all the settings in OpenFOAM.

786 For the meshing part in OpenFOAM, we began with the blockMesh utility to generate a
 787 structured base hexahedral mesh for our computational domain. This provided a coarse mesh that
 788 roughly captured the geometrical boundaries of our domain. The snappyHexMesh utility was
 789 then used to refine the mesh and conform it to our complex geometry. After mesh generation, we
 790 checked the quality of the mesh using Open Foam’s checkMesh utility. This ensured that the
 791 mesh met the quality criteria, such as non-orthogonality, skewness, and aspect ratio. Based on
 792 the mesh quality metrics, iterative refinements were made until the mesh met the desired quality
 793 standards for our simulations. **Table. 3** shows the mesh quality metrics of our model.

Metric	Min	Average	Max
Skewness	0	0.090	2.836
nonOrthogonality	0	10.918	70.305
Aspect Ratio	1	1.229	3.592
Volume Ratio	1	1.346	17.222

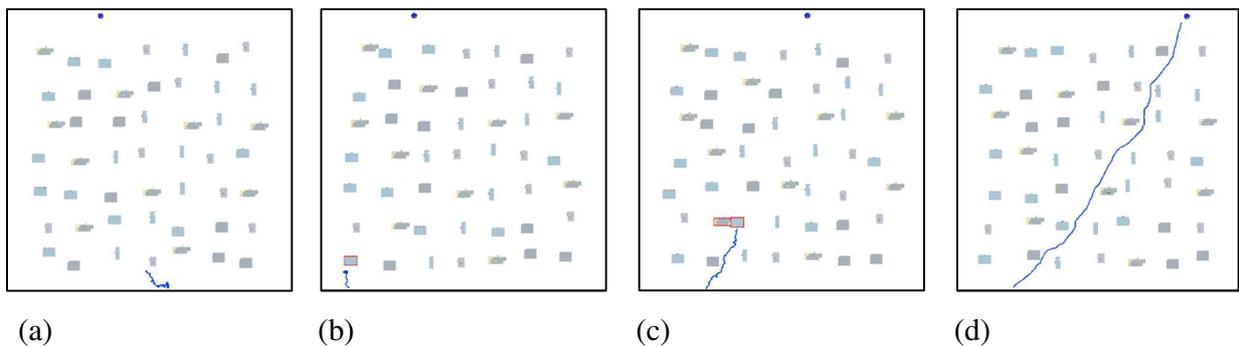
794 **Table. 3** Mesh quality metrics

795 5. RESULTS

796 5.1. Simulated Cases

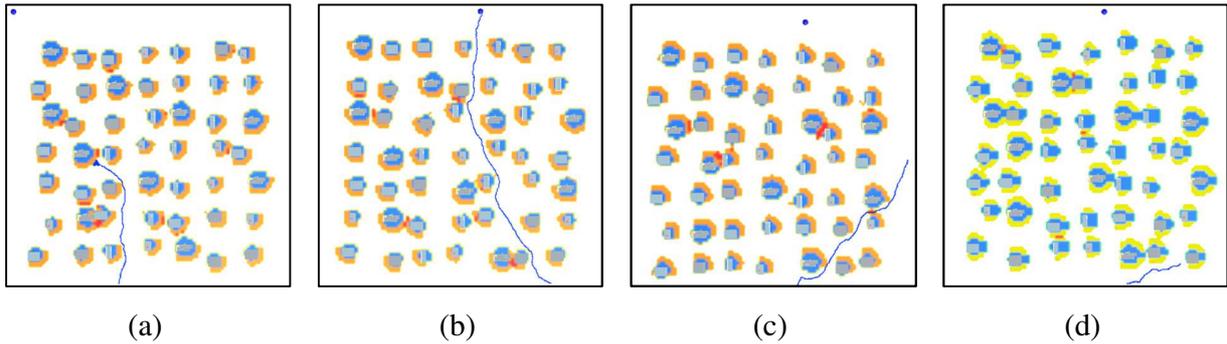
797 We performed a series of tests to examine the performance of the proposed MONRL method in
 798 piloting a drone to find the final target while avoiding and dodging against dynamic wind zones.
 799 **Fig. 7** shows the process of the drone agent learning how to navigate and avoid collision with the
 800 obstacles in the training environment in phase 1. The blue sphere in **Fig.7** refers to the target, the
 801 red rectangle represents the bounding box result, and the buildings are represented as the cube.
 802 To clearly display the drone’s route, we use a blue line to describe the drone’s trajectory in space.
 803 **Fig. 7(a)** and **Fig. 7(b)** show the behavior of the drone at the earlier phase of the training. Since
 804 the agent did not know its task and the meaning of the inputs, it randomly explored the
 805 environment and often collided with the obstacles or the geometries. **Fig. 7(c)** shows that the

806 drone already learned its task and started moving toward the target based on previous rewards.
 807 Besides, it could control its flight in the training scene's scope and does not collide with the
 808 geometries. However, the agent still hit buildings. The reason is that the agent had not
 809 understood the meaning of the input, which represented the locations and dimensions of the
 810 building obstacles. **Fig. 7(d)** shows the drone could safely arrive at the target while avoiding all
 811 the obstacles in its way, and its success rate was 98% after 100,000 epochs. Furthermore, the
 812 drone learned how to choose the shortest path based on the local information from the start point
 813 to the target.



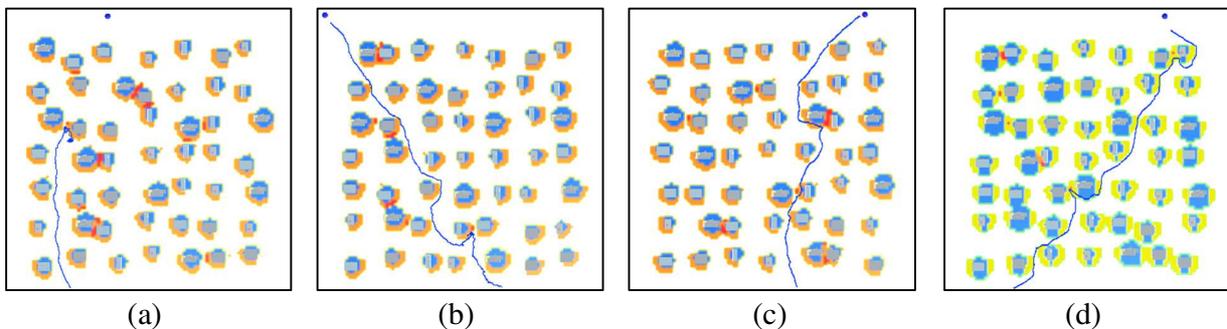
814 **Fig. 7** Examples of training results in phase 1

815
 816 **Fig. 8** shows the process of the drone agent learning how to adapt to the wind force and
 817 change its decision network in phase 2. As we mentioned, we added simulation of wind zones
 818 into the training scenario in this phase, which may have significantly affected the drone's flight
 819 route. To visualize the strength and the kind of wind, we created a heat map on the ground of the
 820 scene. The blank background means that there is no initial wind. Depending on the strength of
 821 the wind, the wind zones around the buildings were assigned three different colors. The red zone
 822 referred to the strongest wind zones, whose speed was larger than the drone's max speed (30
 823 mph). The orange zone showed the magnitude of the wind that was larger than 50% of the
 824 drone's max speed but less than the maximum. Then the weaker wind zones were shown in blue.



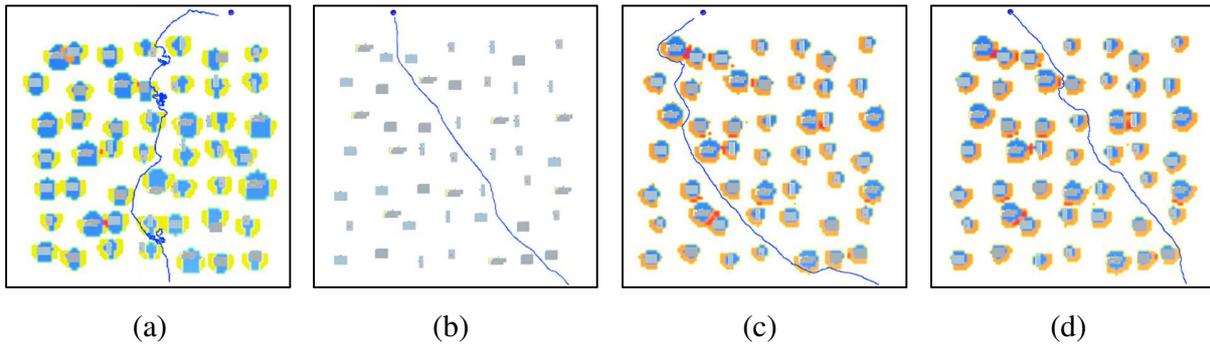
825 **Fig. 8** Examples of training results in phase 2

826 **Fig. 8(a)** shows an example of the poor performance of the drone in the earlier phase of
 827 the training scenario with a wind zone. Unlike the previous high success rate in a windless
 828 environment, the drone failed in most cases during the initial part of the training because it still
 829 generated the same strategies as the wind force did not exist. Then the wind forced the drone to
 830 deviate from its original route until it collided with the buildings. **Fig. 8(b)** shows that the drone
 831 starts recognizing the existence of the wind and changing its action decision. **Fig. 8(c)** shows that
 832 the drone could safely reach the target while dodging all the obstacles and crossing the wind
 833 zones with a success rate of 88% after 40,000 epochs. However, about 10% of the total cases
 834 were timeout cases in training phase 2. **Fig. 8(d)** shows the behavior of drones in timeout cases.
 835 As we can see, the drone was trapped at the boundary of the red wind zone, which indicated that
 836 the magnitude of the wind in this area was larger than the drone's maximum speed. In other
 837 words, if the drone aimed to enter this area and the wind direction was opposite to the drone's
 838 moving direction, the drone would be pushed back. And once the drone returned to its previous
 839 position, it would still decide to go forward because all the environmental information was the
 840 same, making an infinite loop of local trap.



841 **Fig. 9** Examples of training results in phase 3

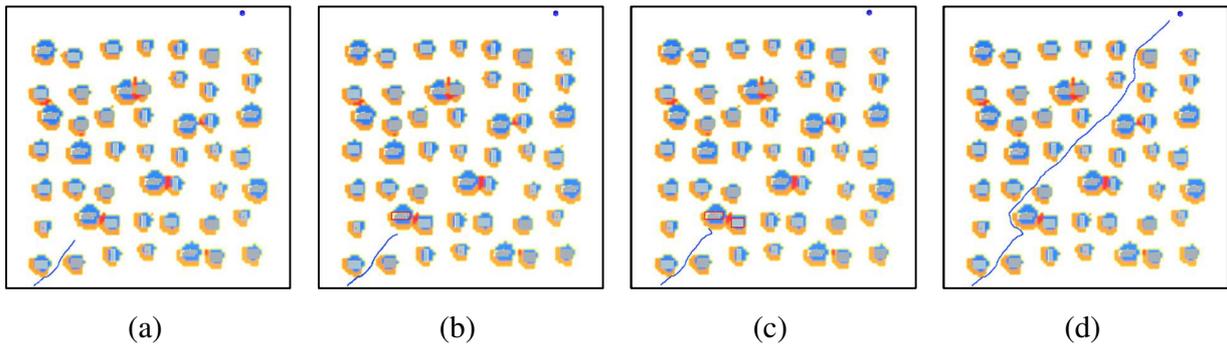
842 **Fig. 9** shows the procedure of the drone agent learning how to detour and avoid the
843 strong wind zone, which was red in the figure. In training phase 3, to increase the probability of
844 timeout cases, we retrained the drone in the same environment once there was a timeout case.
845 Attributed to the LSTM network, the drone could memorize its previous actions, making it
846 possible to bypass the strong wind area. **Fig. 9(a)** shows that the drone started to find another
847 route to leave away from the strong wind zone but failed in the early stage of training. A high
848 penalty for staying around the strong wind zone forced the drone to fly towards different
849 directions to detour, but the drone would return to the area when it left the range of the penalty
850 area. **Fig. 9(b)** shows the drone could find the alternate route after spending sufficient time
851 around the strong wind zone. This means the drone successfully recognized the existence of a
852 strong wind zone (even without any aerodynamic force sensor) and could change its decision
853 based on previous decisions. **Fig. 9(c)** and **Fig. 9(d)** show that after enough epochs of training,
854 the drone could turn quickly when it entered into a strong wind zone and chose to fly away to
855 bypass it. The success rate in this training environment was 98% after 20,000 epochs.



856 **Fig. 10** Examples of training results in phase 4

857
858 **Fig. 10** shows the procedure of training the drone agent in the comprehensive
859 environment in training phase 4. **Fig. 10(a)** shows the poor behaviors of the drone in the early
860 phase with all possible wind directions at the beginning of the training. The model seemed
861 overfitted, and the agent had forgotten how to control the drone in a downwind environment. **Fig.**
862 **10(b)-(d)** show that the drone agent could successfully bypass the strong wind zone and safely
863 arrive at the target on the same map with varying wind directions after 20,000 epochs. In **Fig.**
864 **10(b)**, since there was no wind zone in the environment, the drone chose the shortest path to the

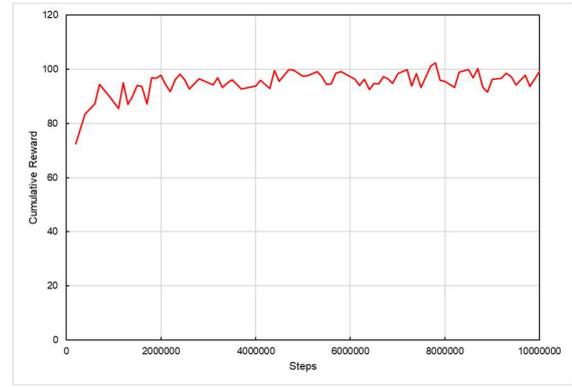
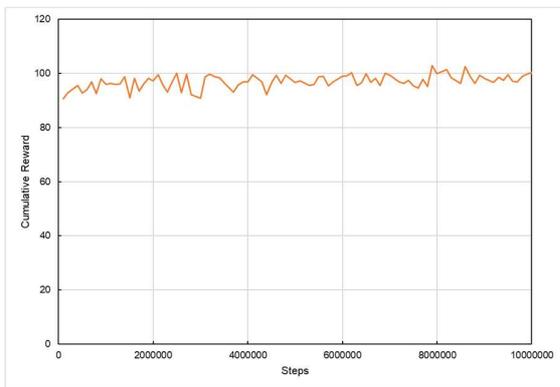
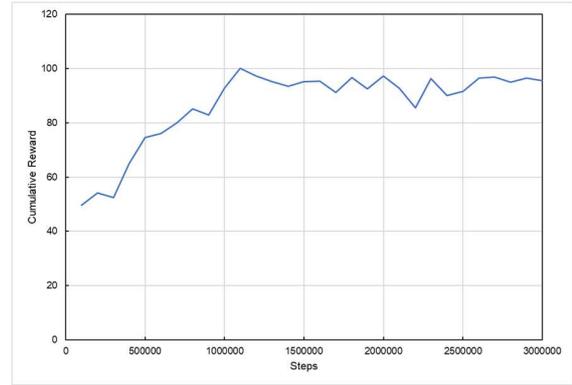
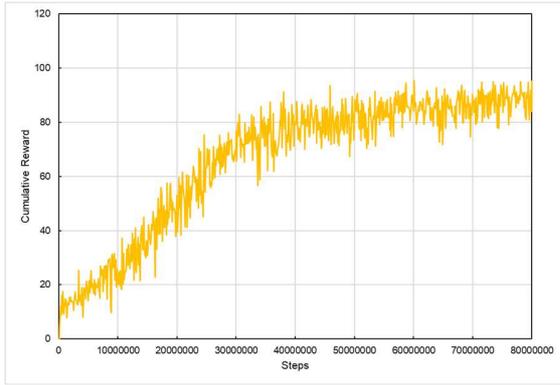
865 target, which proved that the agent may have learned the optimal global route based on the local
866 information. In **Fig. 10(c)**, the map's initial wind direction was from the top right to the bottom
867 left. In this case, the drone chose to fly in the partial wind direction to save energy and increase
868 the speed. At the same time, we found that the drone always kept a certain distance from the
869 nearby buildings to reduce the wind impact. Similarly, in **Fig. 10(d)**, the drone's flight path was
870 more to the right compared to the windless situation. After the training, the final success rate for
871 the model was about 93% after 20,000 epochs.



872 **Fig. 11** Process of drone learning to detour strong winds based only on imagery data

873

874 **Fig. 11** shows the whole process of a drone bypassing a strong wind zone without
875 entering it. **Fig. 11(a)** shows that the drone was flying toward the obstacle, and its camera
876 captured the building on the left of the strong wind zone. At this moment, the drone did not
877 change its direction. **Fig. 11(b)** shows that the drone continued flying in this direction, and its
878 camera captured both buildings next to the strong wind zone, shown in the red box. **Fig. 11(c)**
879 shows that the drone immediately changed its flight direction, bypassed the left building, and
880 continued moving toward the target. **Fig. 11(d)** shows the drone's whole flight route in this case.
881 This drone's behavior was very dynamic because the drone could only obtain the wind
882 information for the imagery data (such as the relative locations of two adjacent buildings) and its
883 own positions. Since the drone was not equipped with any aerodynamic force sensors, there was
884 no relative information about the aerodynamic forces in the inputs of the model. However, the
885 drone agent could still decide to detour before entering the strong wind zone. This proves that the
886 drone could establish some knowledge about the relationship between the wind and the building
887 layout information, such as the distance between the two obstacles and their dimensions.



(a)

(b)

(c)

(d)

888

889

890

891

892

893

894

895

896

897

898

899

900

901

Fig. 12 Learning curve of all phases

As shown in **Fig. 12**, we collect the data from every epoch and record the mean reward every 100,000 steps. For training phase 1, we can observe from **Fig. 12(a)** that the agents received greater rewards as the number of iterations increased. The reward curve raised rapidly at the first period of the training, then increased slower when the steps reached 40M, and eventually, it became stable when the steps reached 80M. **Fig. 12(b)** shows the mean rewards change in phase 2. At first, the rewards were very low. After 1M training, the rewards became stable, which meant the drone could fly safely in an environment with winds presenting. As shown in **Fig. 12(c)**, there was little improvement in the mean reward in phase 3 because the agent had already learned how to control the drone with winds presenting. The goal for the agent in phase 3 was to learn how to quickly detour in a more dynamic situation, which only reduced

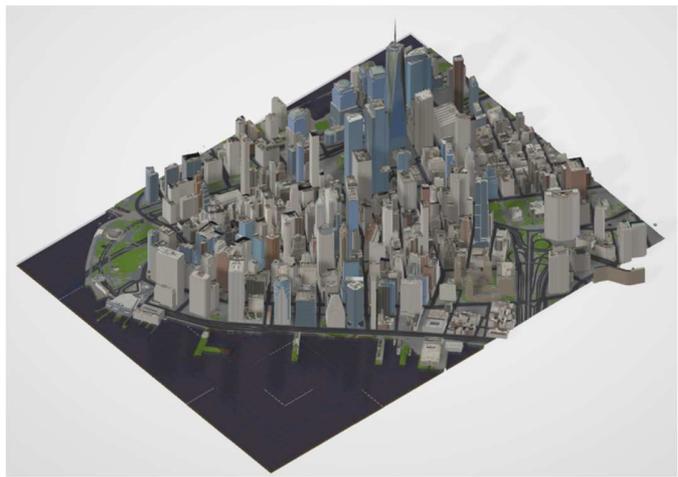
902 the time penalty on the rewards. **Fig. 12(d)** shows the overfitted phenomenon during the training.
903 The mean rewards were lower the expected at the beginning of the training, but the agent quickly
904 reviewed the previous lessons and took 2M epochs to recover its high-performance strategies.

905 **5.2. Validation Case**

906 To further validate the efficacy of the proposed MONRL method in guiding a drone navigating
907 in a more realistic urban setup with precise aerodynamic data, we chose the Manhattan district in
908 New York City as the validation case, as shown in **Fig. 13**. Manhattan is known for its tall
909 buildings, particularly in its central business district. The island has a high concentration of
910 skyscrapers, including iconic structures such as the Empire State Building, the Chrysler Building,
911 and the One World Trade Center. These buildings not only define the city’s skyline but also
912 impact its local environment by creating unique wind patterns, shading effects, and other
913 microclimatic conditions. The main wind direction for Manhattan is generally from the
914 northwest, as the city is located on the east coast of the United States and is influenced by
915 prevailing westerly winds. Therefore, we choose two opposite wind directions, northwest, and
916 southeast, as the main direction of the wind to test the algorithm. **Fig. 13 (c)** and **(d)** show the
917 first-person view from the drone in the virtual simulator.



(a) Selected area in CFD case

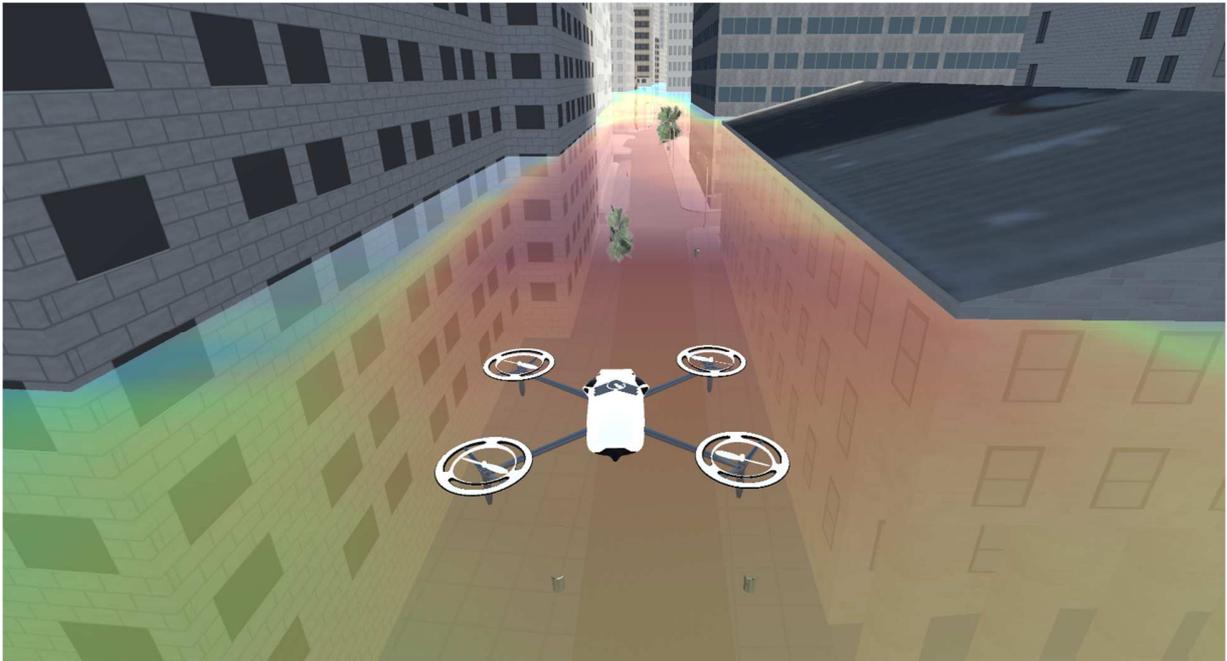


(b) Manhattan district 3D model

918

919

920



921

(c) Drone is flying along with airflow direction

922

(d) Drone is approaching target (blue ball)

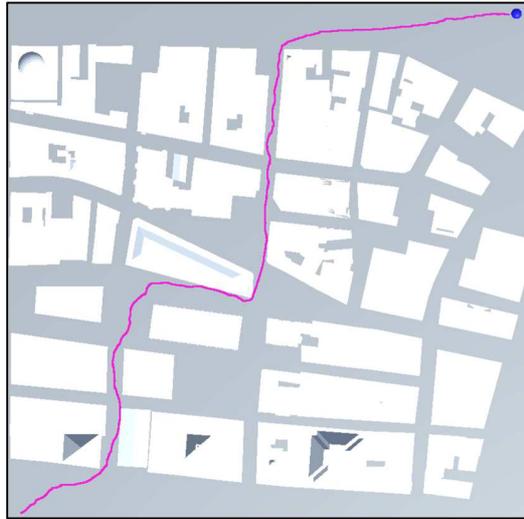
923

Fig. 13 Details of Manhattan district model

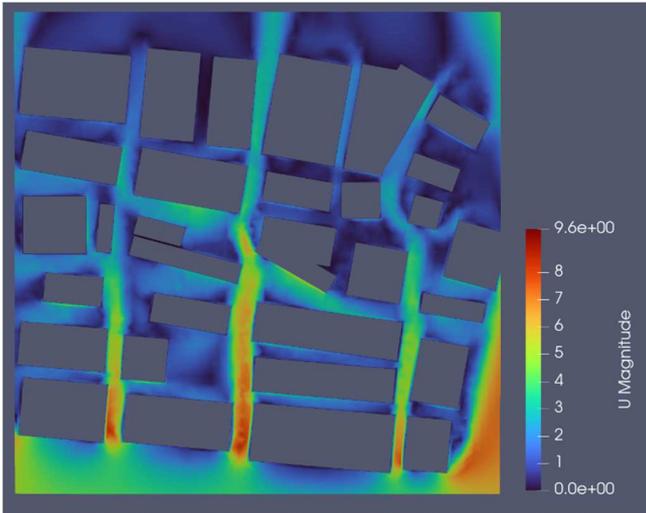
924 There are two main differences between the training environment and the validation test
925 case. The first difference is the layout of the environment. In the training environment, the layout
926 of the buildings is generated and distributed randomly by the Unity program. However, in the
927 validation case, we use this Manhattan district in New York City, a real city model, to test its
928 behavior. The second difference is that we use “Aerodynamic Representation Modeling,”
929 introduced in section 3.5, to simulate the wind map during the training. However, we use Open
930 FOAM to generate the CFD results for the Manhattan model in the validation case, which has a
931 higher authenticity of the wind flow data.

932 **Fig. 14** illustrates the varying trade-off behaviors of the drone within the New York City
933 model, where two conflicting objectives are at play. These objectives involve selecting the most
934 efficient route to the target while circumventing areas of high wind intensity. The results
935 effectively demonstrate the capability of our proposed algorithm to train the drone in adapting to
936 wind dynamics and proficiently navigating through unfamiliar terrain. In **Fig. 14 (a)**, the drone’s
937 behavior is showcased under calm wind conditions. In this context, the drone intelligently opted
938 for the shortest path between the starting point and the destination. **Fig. 14 (c)** and **(e)** present
939 distinct route choices made by the drone in response to varying wind strengths. Specifically,
940 when the wind direction shifted from southeast to northwest, favorably aiding the drone’s flight,
941 the choice to embark on a longer, direct journey emerged as preferred, as highlighted in **Fig. 14**
942 **(e)**. On the other hand, **Fig. 14 (g)** and **(i)** exhibit alternative route selections when the wind
943 forces opposed the drone’s movement. In these scenarios, the presence of potent wind zones,
944 indicated by the red segments in **Fig. 14 (g)** and **(i)**, impeded the drone’s ability to follow the
945 shortest path observed in **Fig. 14 (a)**. To mitigate the impact of adverse wind forces, the drone
946 strategically elected routes that traversed through milder wind zones, relying solely on imagery
947 data and previously memorized positions from earlier epochs. This deviation from conventional
948 paths underscores the drone’s calculated trade-off decision, opting to cover more distance in
949 order to minimize the detrimental influence of the environmental variable—wind force. A demo
950 can be found here: <https://youtu.be/ycMFNs-YCRU>.

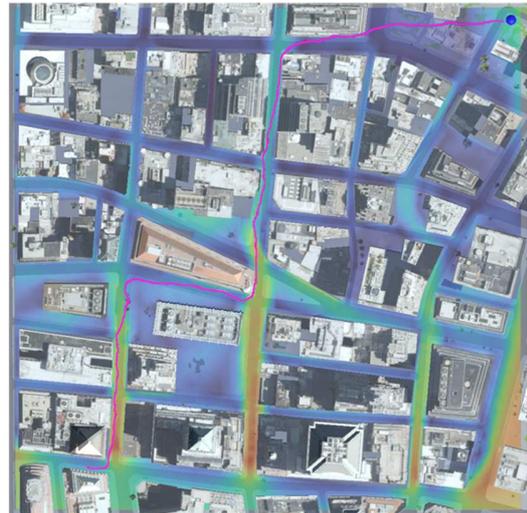
951



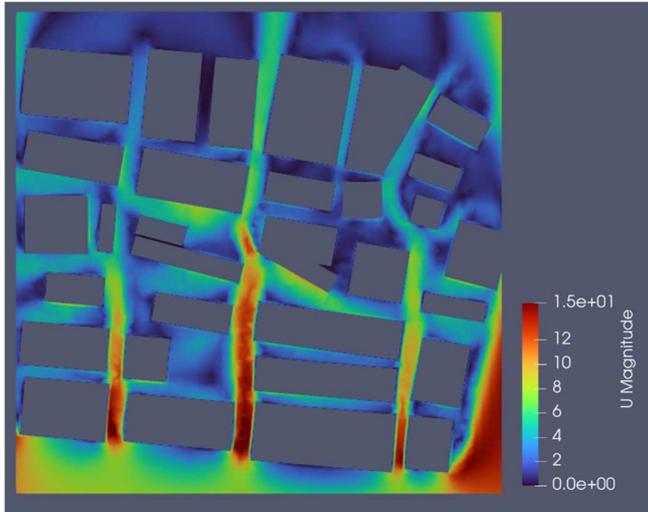
(a) Drone's route in windless environment



(b) CFD case 1 (wind from SE to NW, $w_v = 5\text{m/s}$)



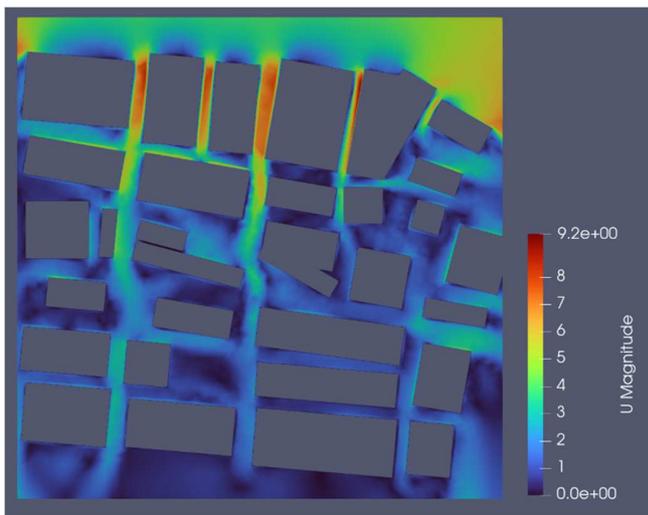
(c) Drone's route in case 1



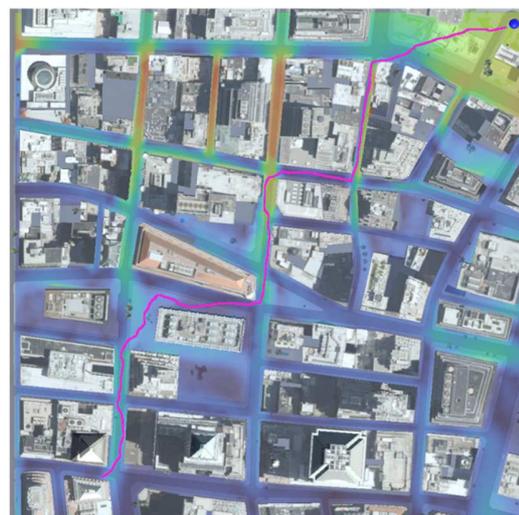
(d) CFD case 2 (wind from SE to NW, $w_v = 10\text{m/s}$)



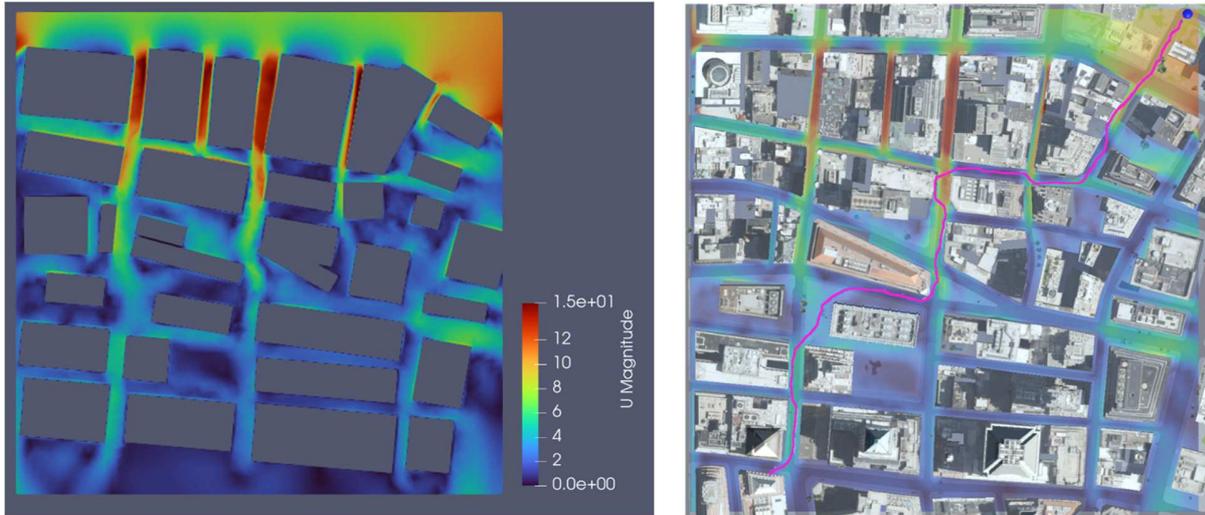
(e) Drone's route in case 2



(f) CFD case 3 (wind from NW to SE, $w_v = 5\text{m/s}$)



(g) Drone's route in case 3



(h) CFD case 4 (wind from NW to SE, $w_v = 10\text{m/s}$)

(i) Drone's route in case 4

952

Fig. 14 Validation case study results

953

6. CONCLUSIONS

954

In this paper, we introduced a cutting-edge approach for drone navigation within urban environments, incorporating the influence of winds. This method relies on multi-objective deep reinforcement learning, LSTM, and a simulation environment. We highlight that conventional drone navigation and obstacle avoidance systems often neglect the impact of environmental factors, particularly the lateral wind surrounding tall buildings in urban areas. Acknowledging the potential simplicity of our empirical test, it's essential to emphasize our intent. While designed in a specific setting, our framework encapsulates some of the more intricate aspects of drone navigation in dynamic environments. By focusing on autonomous drones' ability to adapt to ever-changing aerodynamic conditions using minimalistic inputs, specifically from an onboard camera, we underscored the pivotal relationship between urban building layouts and wind zone distributions. Such an approach challenges the drone to maximize learning from sparse environmental cues, a testament to our methodology's robustness. Consequently, we devise a novel mathematical methodology to simulate wind distribution around buildings and generate a training environment using this method. Specifically, we employed an LSTM architecture combined with a PPO policy to train the drone to navigate and avoid building obstacles in a windless environment. Then, we introduced simulated wind zones into the environment, teaching the drone to execute detours. Finally, we validated the drone's navigation and obstacle avoidance capabilities within an environment featuring varying wind conditions, using an actual New York

971

972 City model integrated with CFD results. The findings demonstrate that our proposed system
973 effectively adapts to wind effects in unknown environments, successfully reaching its destination
974 without collisions. Moreover, the drone's flight path confirms the agent's capacity to plan an
975 optimal global route based on local information, eliminating the need for a global map. During
976 the training process, we can also ascertain that transferring a drone model trained with simulation
977 wind methodology to an unknown environment with CFD results is feasible. The agent learns to
978 control the drone under wind effects from the wind zone created by the wind simulation
979 methodology.

980 Despite the promising results demonstrated by our proposed system, several limitations
981 warrant further investigation. First, the current approach predominantly focuses on wind as the
982 primary environmental factor affecting drone navigation within urban areas. Other factors, such
983 as precipitation, temperature, signal coverage, air pressure, and the inherent decision to assign
984 equal weights to all objectives in tailoring specific objectives, are not accounted for. This could
985 significantly influence drone performance in diverse scenarios. Furthermore, our system's
986 validation in a virtual Manhattan district, although rigorous, introduced challenges due to its
987 heavy reliance on a camera and GPS. Complex environments characterized by densely built
988 structures and varying topographies can sometimes obscure visual feedback, leading to the
989 misidentification of narrow gaps between structures. Navigating such intricate settings and
990 environments with dynamic obstacles, such as moving vehicles and pedestrians, remains a
991 significant challenge. Another noteworthy limitation is the simulation's use of GPS equipment,
992 mirroring a GPS function but lacking in replicating real-time GPS interactions, which are vital
993 for an authentic and comprehensive simulation experience. Addressing sensor noise, especially
994 in the context of environmental factors like electromagnetic interference and adverse weather
995 conditions, remains critical. The current reinforcement learning model, while promising, may
996 encounter limitations in more intricate environments or with diverse drone models. Emphasizing
997 the computationally demanding nature of integrating reinforcement learning with Computational
998 Fluid Dynamics (CFD) results, the system's scalability in various urban settings and drone
999 capabilities await further exploration. Future research should focus on optimizing the wind
1000 simulation methodology to mirror urban environment complexities better. It's also essential to
1001 incorporate a broader range of environmental factors, assigning them variable weights as
1002 appropriate. Addressing dynamic elements within urban settings remains a key area of focus.

1003 Furthermore, developing a more refined obstacle detection methodology is necessary to enhance
1004 the drone's navigation in complex terrains. While our research has demonstrated potential in
1005 urban drone navigation using deep reinforcement learning, it is important to acknowledge the
1006 limitations imposed by the use of simulated GPS data. Future studies may explore the integration
1007 of real GPS data or other advanced positioning technologies to enhance the applicability and
1008 accuracy of such navigational systems. Integrating Hardware-in-the-Loop (HIL) simulations
1009 could provide a more accurate reflection of and response to real-world GPS signals. Lastly,
1010 enhancing the scalability of the proposed system and streamlining the training process to reduce
1011 computational demands are critical steps for improving the practical applicability of this
1012 approach.

1013 It is essential to acknowledge a limitation that has emerged through our empirical test
1014 cases. While these cases were meticulously designed to demonstrate the effectiveness of our
1015 proposed methodology in a controlled environment with limited sensor input, their simplified
1016 nature may limit the direct applicability to more complex, real-world scenarios. We chose these
1017 test cases intentionally, aiming to highlight the core capabilities of our methodology in a clear
1018 and controlled setting. This approach was particularly relevant for urban drone navigation, where
1019 complex and expensive sensor systems are often impractical. By demonstrating effective
1020 performance under these conditions, we aimed to underscore the robustness and practical utility
1021 of our methodology. However, we recognize that the simplified nature of these test cases might
1022 not capture the full spectrum of challenges and variables present in more intricate urban
1023 environments. As such, while these test cases serve as a valuable proof of concept and a
1024 foundational framework for our methodology, they do not fully represent the complexity of real-
1025 world applications. This limitation does not diminish the value of our findings but rather
1026 provides a direction for future research. It underscores the need for further studies that apply our
1027 methodology in more diverse and complex scenarios. We believe that extending our research to
1028 these settings will not only validate but also enhance the applicability and scalability of our
1029 approach.

1030

1031 **ACKNOWLEDGEMENTS**

1032 This material is supported by the Air Force Office of Scientific Research (AFOSR) under grant
1033 FA9550-22-1-0492. Any opinions, findings, conclusions, or recommendations expressed in this
1034 article are those of the authors and do not reflect the views of the AFOSR.

1035 **REFERENCES**

1036 [1] A. Albert, F.S. Leira, L. Imsland. UAV Path Planning using MILP with Experiments. *Modeling,*
1037 *Identification and Control: A Norwegian Research Bulletin*, 38 (1) (2017), pp. 21-32.
1038 <https://doi.org/10.4173/mic.2017.1.3>

1039 [2] E. Aldao, L.M. González-de Santos, H. González-Jorge. LiDAR Based Detect and Avoid System
1040 for UAV Navigation in UAM Corridors. *Drones*, 6 (8) (2022), pp. 185.
1041 <https://doi.org/10.3390/drones6080185>

1042 [3] O. Andersson, M. Wzorek, P. Rudol, P. Doherty. Model-predictive control with stochastic
1043 collision avoidance using bayesian policy optimization. 2016 IEEE International Conference on
1044 Robotics and Automation (ICRA), IEEE, (2016), pp. 4597-4604.
1045 <https://doi.org/10.1109/ICRA.2016.7487661>

1046 [4] L. Apvrille, T. Tanzi, J.-L. Dugelay. Autonomous drones for assisting rescue services within the
1047 context of natural disasters. 2014 XXXIth URSI General Assembly and Scientific Symposium
1048 (URSI GASS), IEEE, (2014), pp. 1-4. <https://doi.org/10.1109/URSIGASS.2014.6929384>

1049 [5] M.Y. Arafat, M.M. Alam, S. Moh. Vision-Based Navigation Techniques for Unmanned Aerial
1050 Vehicles: Review and Challenges. *Drones*, 7 (2) (2023), pp. 89.
1051 <https://doi.org/10.3390/drones7020089>

1052 [6] M.A. Arshad, S.H. Khan, S. Qamar, M.W. Khan, I. Murtza, J. Gwak, A. Khan. Drone Navigation
1053 Using Region and Edge Exploitation-Based Deep CNN. *IEEE Access*, 10 (2022), pp. 95441-
1054 95450. <https://doi.org/10.1109/ACCESS.2022.3204876>

1055 [7] A.T. Azar, A. Koubaa, N. Ali Mohamed, H.A. Ibrahim, Z.F. Ibrahim, M. Kazim, A. Ammar, B.
1056 Benjdira, A.M. Khamis, I.A. Hameed. Drone deep reinforcement learning: A review. *Electronics*,
1057 10 (9) (2021), pp. 999. <https://doi.org/10.3390/electronics10090999>

1058 [8] T. Baca, D. Hert, G. Loianno, M. Saska, V. Kumar. Model predictive trajectory tracking and
1059 collision avoidance for reliable outdoor deployment of unmanned aerial vehicles. 2018 IEEE/RSJ
1060 International Conference on Intelligent Robots and Systems (IROS), IEEE, (2018), pp. 6753-
1061 6760. <https://doi.org/10.1109/IROS.2018.8594266>

1062 [9] D. Baskar, A. Gorodetsky. A Simulated Wind-field Dataset for Testing Energy Efficient Path-
1063 Planning Algorithms for UAVs in Urban Environment. *AIAA AVIATION 2020 FORUM*,
1064 (2020), p. 2920. <https://doi.org/10.2514/6.2020-2920>

1065 [10] Y. Bengio, J. Louradour, R. Collobert, J. Weston. Curriculum learning. *Proceedings of the 26th*
1066 *annual international conference on machine learning*, (2009), pp. 41-48.
1067 <https://doi.org/10.1145/1553374.1553380>

1068 [11] L. Biao, J. Cunyan, W. Lu, C. Weihua, L. Jing. A parametric study of the effect of building layout
1069 on wind flow over an urban area. *Building and Environment*, 160 (2019), pp. 106160.
1070 <https://doi.org/10.1016/j.buildenv.2019.106160>

1071 [12] B. Blocken, T. Stathopoulos, J. Carmeliet. Wind environmental conditions in passages between
1072 two long narrow perpendicular buildings. *Journal of Aerospace Engineering*, 21 (4) (2008), pp.
1073 280-287. [https://doi.org/10.1061/\(ASCE\)0893-1321\(2008\)21:4\(280\)](https://doi.org/10.1061/(ASCE)0893-1321(2008)21:4(280))

1074 [13] M. Bolognini, L. Fagiano. Lidar-based navigation of tethered drone formations in an unknown
1075 environment. *IFAC-PapersOnLine*, 53 (2) (2020), pp. 9426-9431.
1076 <https://doi.org/10.1016/j.ifacol.2020.12.2413>

1077 [14] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, J.J. Leonard. Past,
1078 present, and future of simultaneous localization and mapping: Toward the robust-perception age.
1079 *IEEE Transactions on robotics*, 32 (6) (2016), pp. 1309-1332.
1080 <https://doi.org/10.1109/TRO.2016.2624754>

1081 [15] A. Capolupo, S. Pindozi, C. Okello, N. Fiorentino, L. Boccia. Photogrammetry for
1082 environmental monitoring: The use of drones and hydrological models for detection of soil

- 1083 contaminated by copper. *Science of the Total Environment*, 514 (2015), pp. 298-306.
 1084 <https://doi.org/10.1016/j.scitotenv.2015.01.109>
- 1085 [16] P. Chakravarty, K. Kelchtermans, T. Roussel, S. Wellens, T. Tuytelaars, L. Van Eycken. CNN-
 1086 based single image obstacle avoidance on a quadrotor. 2017 IEEE international conference on
 1087 robotics and automation (ICRA), IEEE, (2017), pp. 6369-6374.
 1088 <https://doi.org/10.1109/ICRA.2017.7989752>
- 1089 [17] A. Devo, G. Mezzetti, G. Costante, M.L. Fravolini, P. Valigi. Towards generalization in target-
 1090 driven visual navigation by using deep reinforcement learning. *IEEE Transactions on Robotics*,
 1091 36 (5) (2020), pp. 1546-1561. <https://doi.org/10.1109/TRO.2020.2994002>
- 1092 [18] A. Devos, E. Ebeid, P. Manoonpong. Development of autonomous drones for adaptive obstacle
 1093 avoidance in real world environments. 2018 21st Euromicro conference on digital system design
 1094 (DSD), IEEE, (2018), pp. 707-710. <https://doi.org/10.1109/DSD.2018.00009>
- 1095 [19] N. El-Sheimy, Y. Li. Indoor navigation: State of the art and future trends. *Satellite Navigation*, 2
 1096 (1) (2021), pp. 1-23. <https://doi.org/10.1186/s43020-021-00041-3>
- 1097 [20] T. Elmokadem, A.V. Savkin. A hybrid approach for autonomous collision-free UAV navigation
 1098 in 3D partially unknown dynamic environments. *Drones*, 5 (3) (2021), pp. 57.
 1099 <https://doi.org/10.3390/drones5030057>
- 1100 [21] A. Entrop, A. Vasenev. Infrared drones in the construction industry: designing a protocol for
 1101 building thermography procedures. *Energy procedia*, 132 (2017), pp. 63-68.
 1102 <https://doi.org/10.1016/j.egypro.2017.09.636>
- 1103 [22] M. Esposito, M. Crimaldi, V. Cirillo, F. Sarghini, A. Maggio. Drone and sensor technology for
 1104 sustainable weed management: A review. *Chemical and Biological Technologies in Agriculture*,
 1105 8 (1) (2021), pp. 1-11. <https://doi.org/10.1186/s40538-021-00217-8>
- 1106 [23] Y. Fan, S. Chu, W. Zhang, R. Song, Y. Li. Learn by observation: Imitation learning for drone
 1107 patrolling from videos of a human navigator. 2020 IEEE/RSJ International Conference on
 1108 Intelligent Robots and Systems (IROS), IEEE, (2020), pp. 5209-5216.
 1109 <https://doi.org/10.48550/arXiv.2008.13193>
- 1110 [24] D. Floreano, R.J. Wood. Science, technology and the future of small autonomous drones. *nature*,
 1111 521 (7553) (2015), pp. 460-466. <https://doi.org/10.1038/nature14542>
- 1112 [25] L.R. García Carrillo, A.E. Dzul López, R. Lozano, C. Pégard. Combining stereo vision and
 1113 inertial navigation system for a quad-rotor UAV. *Journal of intelligent & robotic systems*, 65 (1-
 1114 4) (2012), pp. 373-387. <https://doi.org/10.1007/s10846-011-9571-7>
- 1115 [26] S. Hochreiter, J. Schmidhuber. Long short-term memory. *Neural computation*, 9 (8) (1997), pp.
 1116 1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- 1117 [27] V.J. Hodge, R. Hawkins, R. Alexander. Deep reinforcement learning for drone navigation using
 1118 sensor data. *Neural Computing and Applications*, 33 (2021), pp. 2015-2033.
 1119 <https://doi.org/10.1007/s00521-020-05097-x>
- 1120 [28] C. Huang, X. Zhou, X. Ran, J. Wang, H. Chen, W. Deng. Adaptive cylinder vector particle
 1121 swarm optimization with differential evolution for UAV path planning. *Engineering Applications*
 1122 *of Artificial Intelligence*, 121 (2023), pp. 105942. <https://doi.org/10.1016/j.engappai.2023.105942>
- 1123 [29] M.L. Imrane, A. Melingui, J.J.B. Mvogo Ahanda, F. Biya Motto, R. Merzouki. Artificial
 1124 potential field neuro-fuzzy controller for autonomous navigation of mobile robots. *Proceedings of*
 1125 *the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 235
 1126 (7) (2021), pp. 1179-1192. <https://doi.org/10.1177/0959651820974831>
- 1127 [30] H. Jasak. OpenFOAM: open source CFD in research and industry. *International Journal of Naval*
 1128 *Architecture and Ocean Engineering*, 1 (2) (2009), pp. 89-94. <https://doi.org/10.2478/IJNAOE-2013-0011>
- 1129
 1130 [31] K. Javanroodi, M. Mahdavinejad, V.M. Nik. Impacts of urban morphology on reducing cooling
 1131 load and increasing ventilation potential in hot-arid climate. *Applied energy*, 231 (2018), pp. 714-
 1132 746. <https://doi.org/10.1016/j.apenergy.2018.09.116>

- 1133 [32] S. Jung, S. Hwang, H. Shin, D.H. Shim. Perception, guidance, and navigation for indoor
1134 autonomous drone racing using deep learning. *IEEE Robotics and Automation Letters*, 3 (3)
1135 (2018), pp. 2539-2544. <https://doi.org/10.1109/LRA.2018.2808368>
- 1136 [33] L.P. Kaelbling, M.L. Littman, A.W. Moore. Reinforcement learning: A survey. *Journal of*
1137 *artificial intelligence research*, 4 (1996), pp. 237-285. <https://doi.org/10.48550/arXiv.cs/9605103>
- 1138 [34] M.T.R. Khan, M. Muhammad Saad, Y. Ru, J. Seo, D. Kim. Aspects of unmanned aerial vehicles
1139 path planning: Overview and applications. *International Journal of Communication Systems*, 34
1140 (10) (2021), pp. e4827. <https://doi.org/10.1002/dac.4827>
- 1141 [35] J. Kwak, Y. Sung. Autonomous UAV flight control for GPS-based navigation. *IEEE Access*, 6
1142 (2018), pp. 37947-37955. <https://doi.org/10.1109/ACCESS.2018.2854712>
- 1143 [36] S. Lee, D. Har, D. Kum. Drone-assisted disaster management: Finding victims via infrared
1144 camera and lidar sensor fusion. 2016 3rd Asia-Pacific World Congress on Computer Science and
1145 Engineering (APWC on CSE), IEEE, (2016), pp. 84-89. [https://doi.org/10.1109/APWC-on-](https://doi.org/10.1109/APWC-on-CSE.2016.025)
1146 [CSE.2016.025](https://doi.org/10.1109/APWC-on-CSE.2016.025)
- 1147 [37] T. Lee, S. McKeever, J. Courtney. Flying free: A research overview of deep learning in drone
1148 navigation autonomy. *Drones*, 5 (2) (2021), pp. 52. <https://doi.org/10.3390/drones5020052>
- 1149 [38] J. Linchant, J. Lisein, J. Semeki, P. Lejeune, C. Vermeulen. Are unmanned aircraft systems (UAS
1150 s) the future of wildlife monitoring? A review of accomplishments and challenges. *Mammal*
1151 *Review*, 45 (4) (2015), pp. 239-252. <https://doi.org/10.1111/mam.12046>
- 1152 [39] Y. Liu, Q. Wang, H. Hu, Y. He. A novel real-time moving target tracking and path planning
1153 system for a quadrotor UAV in unknown unstructured outdoor scenes. *IEEE Transactions on*
1154 *Systems, Man, and Cybernetics: Systems*, 49 (11) (2018), pp. 2362-2372.
1155 <https://doi.org/10.1109/TSMC.2018.2808471>
- 1156 [40] Y. Lu, Z. Xue, G.-S. Xia, L. Zhang. A survey on vision-based UAV navigation. *Geo-spatial*
1157 *information science*, 21 (1) (2018), pp. 21-32. <https://doi.org/10.1080/10095020.2017.1420509>
- 1158 [41] K. Minoda, F. Schilling, V. Wüest, D. Floreano, T. Yairi. Viode: A simulated dataset to address
1159 the challenges of visual-inertial odometry in dynamic environments. *IEEE Robotics and*
1160 *Automation Letters*, 6 (2) (2021), pp. 1343-1350. <https://doi.org/10.1109/LRA.2021.3058073>
- 1161 [42] B. Mishra, D. Garg, P. Narang, V. Mishra. Drone-surveillance for search and rescue in natural
1162 disaster. *Computer Communications*, 156 (2020), pp. 1-10.
1163 <https://doi.org/10.1016/j.comcom.2020.03.012>
- 1164 [43] V. Mnih, K. Kavukcuoglu, D. Silver, A.A. Rusu, J. Veness, M.G. Bellemare, A. Graves, M.
1165 Riedmiller, A.K. Fidjeland, G. Ostrovski. Human-level control through deep reinforcement
1166 learning. *nature*, 518 (7540) (2015), pp. 529-533. <https://doi.org/10.1038/nature14236>
- 1167 [44] A. Mousavian, D. Anguelov, J. Flynn, J. Kosecka. 3d bounding box estimation using deep
1168 learning and geometry. *Proceedings of the IEEE conference on Computer Vision and Pattern*
1169 *Recognition*, (2017), pp. 7074-7082. <https://doi.org/10.48550/arXiv.1612.00496>
- 1170 [45] G. Muñoz, C. Barrado, E. Çetin, E. Salami. Deep reinforcement learning for drone delivery.
1171 *Drones*, 3 (3) (2019), pp. 72. <https://doi.org/10.3390/drones3030072>
- 1172 [46] A. Nagabandi, I. Clavera, S. Liu, R.S. Fearing, P. Abbeel, S. Levine, C. Finn. Learning to adapt
1173 in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint*
1174 *arXiv:1803.11347* (2018). <https://doi.org/10.48550/arXiv.1803.11347>
- 1175 [47] J.C. Nathanael, C.H.J. Wang, K.H. Low. Simulation of Wind Field in a Building Complex for
1176 Evaluation of the Wind Effect Along UAS Flight Path. *AIAA AVIATION 2023 Forum*, (2023),
1177 p. 4096. <https://doi.org/10.2514/6.2023-4096>
- 1178 [48] A. Patrik, G. Utama, A.A.S. Gunawan, A. Chowanda, J.S. Suroso, R. Shofiyanti, W. Budiharto.
1179 GNSS-based navigation systems of autonomous drone for delivering items. *Journal of Big Data*,
1180 6 (2019), pp. 1-14. <https://doi.org/10.1186/s40537-019-0214-3>
- 1181 [49] H.X. Pham, H.M. La, D. Feil-Seifer, L.V. Nguyen. Autonomous uav navigation using
1182 reinforcement learning. *arXiv preprint arXiv:1801.05086* (2018).
1183 <https://doi.org/10.1109/SSRR.2018.8468611>

- 1184 [50] V. Puri, A. Nayyar, L. Raja. Agriculture drones: A modern breakthrough in precision agriculture.
 1185 Journal of Statistics and Management Systems, 20 (4) (2017), pp. 507-518.
 1186 <https://doi.org/10.1080/09720510.2017.1395171>
- 1187 [51] S. Ragi, E.K. Chong. UAV path planning in a dynamic environment via partially observable
 1188 Markov decision process. IEEE Transactions on Aerospace and Electronic Systems, 49 (4)
 1189 (2013), pp. 2397-2412. <https://doi.org/10.1109/TAES.2013.6621824>
- 1190 [52] A. Ramezani Dooraki, D.-J. Lee. A multi-objective reinforcement learning based controller for
 1191 autonomous navigation in challenging environments. Machines, 10 (7) (2022), pp. 500.
 1192 <https://doi.org/10.3390/machines10070500>
- 1193 [53] A. Rejeb, K. Rejeb, S.J. Simske, H. Treiblmaier. Drones for supply chain management and
 1194 logistics: a review and research agenda. International Journal of Logistics Research and
 1195 Applications (2021), pp. 1-24. <https://doi.org/10.1080/13675567.2021.1981273>
- 1196 [54] A. Restas. Drone applications for supporting disaster management. World Journal of Engineering
 1197 and Technology, 3 (03) (2015), pp. 316. <http://dx.doi.org/10.4236/wjet.2015.33C047>
- 1198 [55] M.F. Sani, G. Karimian. Automatic navigation and landing of an indoor AR. drone quadrotor
 1199 using ArUco marker and inertial sensors. 2017 international conference on computer and drone
 1200 applications (IConDA), IEEE, (2017), pp. 102-107.
 1201 <https://doi.org/10.1109/ICONDA.2017.8270408>
- 1202 [56] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov. Proximal policy optimization
 1203 algorithms. arXiv preprint arXiv:1707.06347 (2017). <https://doi.org/10.48550/arXiv.1707.06347>
- 1204 [57] A. Shahoud, D. Shashev, S. Shidlovskiy. Visual navigation and path tracking using street
 1205 geometry information for image alignment and servoing. Drones, 6 (5) (2022), pp. 107.
 1206 <https://doi.org/10.3390/drones6050107>
- 1207 [58] A. Shantia, R. Timmers, Y. Chong, C. Kuiper, F. Bidoia, L. Schomaker, M. Wiering. Two-stage
 1208 visual navigation by deep neural networks and multi-goal reinforcement learning. Robotics and
 1209 Autonomous Systems, 138 (2021), pp. 103731. <https://doi.org/10.1016/j.robot.2021.103731>
- 1210 [59] S.M. Shavarani, M.G. Nejad, F. Rismanchian, G. Izbirak. Application of hierarchical facility
 1211 location problem for optimization of a drone delivery system: a case study of Amazon prime air
 1212 in the city of San Francisco. The International Journal of Advanced Manufacturing Technology,
 1213 95 (2018), pp. 3141-3153. <https://doi.org/10.1007/s00170-017-1363-1>
- 1214 [60] S.-Y. Shin, Y.-W. Kang, Y.-G. Kim. Reward-driven U-Net training for obstacle avoidance drone.
 1215 Expert Systems with Applications, 143 (2020), pp. 113064.
 1216 <https://doi.org/10.1016/j.eswa.2019.113064>
- 1217 [61] A. Suleiman, Z. Zhang, L. Carlone, S. Karaman, V. Sze. Navion: A 2-mw fully integrated real-
 1218 time visual-inertial odometry accelerator for autonomous navigation of nano drones. IEEE
 1219 Journal of Solid-State Circuits, 54 (4) (2019), pp. 1106-1119.
 1220 <https://doi.org/10.1109/JSSC.2018.2886342>
- 1221 [62] Q. Sun, M. Li, T. Wang, C. Zhao. UAV path planning based on improved rapidly-exploring
 1222 random tree. 2018 Chinese control and decision conference (CCDC), IEEE, (2018), pp. 6420-
 1223 6424. <https://doi.org/10.1109/CCDC.2018.8408258>
- 1224 [63] G. Tong, N. Jiang, L. Biyue, Z. Xi, W. Ya, D. Wenbo. UAV navigation in high dynamic
 1225 environments: A deep reinforcement learning approach. Chinese Journal of Aeronautics, 34 (2)
 1226 (2021), pp. 479-489. <https://doi.org/10.1016/j.cja.2020.05.011>
- 1227 [64] F. Vanegas, K.J. Gaston, J. Roberts, F. Gonzalez. A framework for UAV navigation and
 1228 exploration in GPS-denied environments. 2019 IEEE Aerospace Conference, IEEE, (2019), pp. 1-6.
 1229 <https://doi.org/10.1109/AERO.2019.8741612>
- 1230 [65] D. Ventura, M. Bruno, G.J. Lasinio, A. Belluscio, G. Ardizzone. A low-cost drone based
 1231 application for identifying and mapping of coastal fish nursery grounds. Estuarine, Coastal and
 1232 Shelf Science, 171 (2016), pp. 85-98. <https://doi.org/10.1016/j.ecss.2016.01.030>

- 1233 [66] L. von Stumberg, V. Usenko, J. Engel, J. Stückler, D. Cremers. From monocular SLAM to
1234 autonomous drone exploration. 2017 European Conference on Mobile Robots (ECMR), IEEE,
1235 (2017), pp. 1-8. <https://doi.org/10.1109/ECMR.2017.8098709>
- 1236 [67] S. Wang, Z. Mao, C. Zeng, H. Gong, S. Li, B. Chen. A new method of virtual reality based on
1237 Unity3D. 2010 18th international conference on Geoinformatics, IEEE, (2010), pp. 1-5.
1238 <https://doi.org/10.1109/GEOINFORMATICS.2010.5567608>
- 1239 [68] Z. Wang, Y. Wu, Q. Niu. Multi-sensor fusion in automated driving: A survey. *Ieee Access*, 8
1240 (2019), pp. 2847-2868. <https://doi.org/10.1109/ACCESS.2019.2962554>
- 1241 [69] N. Wen, X. Su, P. Ma, L. Zhao, Y. Zhang. Online UAV path planning in uncertain and hostile
1242 environments. *International journal of machine learning and cybernetics*, 8 (2017), pp. 469-487.
1243 <https://doi.org/10.1007/s13042-015-0339-4>
- 1244 [70] Z. Xue, T. Gonsalves. Vision based drone obstacle avoidance by deep reinforcement learning. *AI*,
1245 2 (3) (2021), pp. 366-380. <https://doi.org/10.3390/ai2030023>
- 1246 [71] M. Zhang, M. Zhang, Y. Chen, M. Li. IMU data processing for inertial aided navigation: A
1247 recurrent neural network based approach. 2021 IEEE International Conference on Robotics and
1248 Automation (ICRA), IEEE, (2021), pp. 3992-3998.
1249 <https://doi.org/10.1109/ICRA48506.2021.9561172>

1250