

# Synergistic Reinforcement and Imitation Learning for Vision-driven Autonomous Flight of UAV Along River

Zihan Wang\*, Jianwen Li\*, and Nina Mahmoudian

**Abstract**—Vision-driven autonomous flight and obstacle avoidance of Unmanned Aerial Vehicles (UAVs) along complex riverine environments for tasks like rescue and surveillance requires a robust navigation policy, which is yet difficult to obtain due to the shortage of trainable riverine environment simulators. To easily verify the vision-based navigation controller performance for the river following task before real-world deployment, we developed a trainable photo-realistic dynamics-free riverine simulation environment using Unity. In this paper, we address the shortcomings that vanilla Reinforcement Learning (RL) algorithm encounters in learning a navigation policy within this partially observable, non-Markovian environment. We propose a synergistic approach that integrates RL and Imitation Learning (IL). Initially, an IL expert is trained on manually collected demonstrations, which then guides the RL policy training process. Concurrently, experiences generated by the RL agent are utilized to re-train the IL expert, enhancing its ability to generalize to unseen data. By leveraging the strengths of both RL and IL, this framework achieves a faster convergence rate and higher performance compared to pure RL, pure IL, and RL combined with static IL algorithms. The results validate the efficacy of the proposed method in terms of both task completion and efficiency. The code and trainable environments are available<sup>1</sup>.

## I. INTRODUCTION

Autonomous flight of UAVs over inland waterways has been investigated in applications like water sampling [1], inspection of irrigation system [2], and cooperative drone-boat navigation tasks [3]–[5]. However, few studies addressed the along-river autonomous navigation tasks of UAVs, which can be applied in rescue, surveillance, and inspection missions over inland waterways. Seasonal variations in river channel geometry and water levels render manually planned GPS waypoints less reliable for precise river following missions. Moreover, in cluttered riverine environments, UAVs must navigate while avoiding obstacles such as trees and bridges, all while maintaining a safe altitude above the water.

A camera can be used as the major perception source for the UAV to take advantage of information embedded in images and the recent developments in image processing. To reduce the cost and safety issues for real-world experiments, we developed a riverine simulation environment using Unity [6] to verify the feasibility of controlling a UAV to fly along a river with only photo-realistic vision input.

\* Authors contributed equally.

This work was supported by ONR N00014-20-1-2085

Zihan Wang, Jianwen Li, and Nina Mahmoudian are with the School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA wang5044, li3602, ninam@purdue.edu

<sup>1</sup><https://github.com/lijianwen1997/Synergistic-Reinforcement-and-Imitation-Learning>

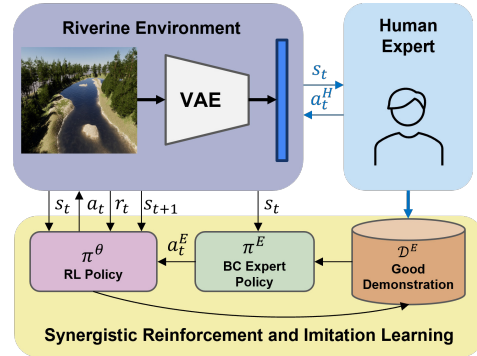


Fig. 1: System architecture. A human expert collects good trajectories before training and the transitions are represented by blue arrows.  $a_t^H$  denotes the human expert’s action. The transitions during training are represented by black arrows.  $a_t^E$  denotes the IL expert’s action.

River following can be classified as a track following task, which poses significant training challenges as the environment does not naturally conform to a Markov Decision Process (MDP) due to the agent’s partial observability and the non-Markovian reward function [7]. Specifically, the first-person camera view of the UAV agent in the riverine scene may lack sufficient water pixels to guide the agent along the river, or it may contain bifurcated river channels, confusing the agent’s navigation decisions. Moreover, the UAV agent receives positive rewards only upon visiting new river segments, with no reward for idling above already visited segments. Consequently, the immediate reward depends on both past observations and actions, rather than solely on the current ones, rendering the task a Partially Observable Markov Decision Process (POMDP) with non-Markovian rewards. This challenge has been addressed in various ways, including formulating it as a coverage control problem in submodular RL frameworks [8]. To further tackle this, combined approaches of RL and Learning from Demonstration (LfD) have been proposed [9]–[11]. However, how to effectively leverage the complementary advantages of imitative learning from demonstrations and reinforcement learning from experience for the track following task remains an area requiring further research and exploration.

To investigate into this problem, we pre-train an IL expert using human demonstrations to provide basic guidance, then combine a RL algorithm with the online-updating IL algorithm that accumulates and learns from decent agent experiences, to mimic the progressive expert behaviors while

learning. In our experiments, Proximal Policy Optimization (PPO [12]) is chosen as the RL agent, Behavior Cloning (BC [13]) the IL expert. This training framework is examined to boost the learning speed and improve final agent performance in the track following tasks, outperforming PPO, StaticBC, DynamicBC and PPO with StaticBC guidance baselines.

In summary, the main contributions of this paper are:

- The creation of the trainable photo-realistic riverine simulation environment, which is publicly available, and serves as a testbed for vision-driven autonomous navigation of UAVs in the river following task.
- The development of synergistic reinforcement and imitation learning framework that interactively trains RL agent and IL expert to reduce the convergence time and increase the overall performance.

## II. RELATED WORK

Autonomous river following in a relatively simple curved river segment has been tested using a water segmentation mask from HSV thresholded images [14], [15]. This water mask extraction technique and the control method both require careful manual tuning, especially for complex riverine environments with obstacles and multiple river channels. To alleviate laborious manual tuning of the river following controller, modern image embedding methods and RL algorithms can be investigated in a photo-realistic riverine simulator.

Game engines as simulation environments have brought more advantages for vision-driven RL training, due to their flexibility in customizing environmental parameters, their low cost as a testbed before real-world deployment, and most importantly, their photo-realism to reduce the Sim2Real gap in air [16], [17], land [18], [19] (indoor/outdoor), and even underwater applications [20]. To achieve a safer and more robust autonomous navigation controller in a diverse and adaptive riverine domain, researchers utilized Unreal Engine to train the agent for the river following task [21], [22]. Their proof of the superiority of the Variational Auto Encoder (VAE) [23] for image encoding, over end-to-end controller training, is adopted in our work. However, the reset criteria of their river simulation environment are manual-based, the agent action is limited to a 2D plane, and the environment is not publicly available. To overcome these difficulties, our Unity-based riverine simulator can be automatically trained and tested for the river following task, supports 3D action space and collision detection, and is made public as an OpenAI gym [24] environment, thanks to Unity ML-Agents Toolkit [6].

Reinforcement Learning and Imitation Learning are powerful techniques in robotics. RL enables robots to learn optimal behavior through trial and error, mastering tasks like obstacle avoidance [25], path planning [26], and autonomous driving [27] in dynamic environments. RL can be time-consuming and costly, requiring extensive interactions with the environment. On the other hand, IL uses expert demonstrations to quickly train robots enabling them to mimic complex skills efficiently, especially when human expertise


is available [28]. IL is limited by its dependence on static training datasets, making it difficult to generalize to new scenarios.

Researchers have begun integrating RL and IL to enhance training efficiency and performance. For instance, DAgger [29] enhances Behavior Cloning by training on a dynamically updating dataset mirroring the observations the trained policy is expected to encounter. However, it necessitates continual online querying of optimal actions from a human expert. Other approaches include frameworks that use IL’s guidance to train RL. Some approaches use IL demonstrations solely during the initialization phase [30], [30], while others regulate the discrepancy between the agent’s policy and the expert’s [31], [32]. However, for these methods that use a static IL’s guidance throughout the training, further improvements of RL agents are impeded by the constraints imposed by the expert policy.

To overcome these limitations, we introduce an interactive framework called “Synergistic Reinforcement and Imitation Learning.” This framework trains an RL agent using IL expert advice while simultaneously refining the expert policy through online sampling of high-quality demonstrations generated by the RL agent. By facilitating collaborative improvement between IL and RL methodologies, this approach enhances training efficiency and overall performance.



(a) Comparison of an image captured in Wildcat Creek, Indiana, USA (left) and image from Unity riverine environment (right) that are alike in overall components layout and texture appearance.



(b) Orthographic images of the riverine simulation environment and the CliffCircular environment.

Fig. 2: Real and simulated riverine environments, and Cliff-Circular environment (with two extra random cliff grids) for river/track following task.

## III. SIMULATION

### A. Photo-realistic Riverine Environment

The riverine simulation environment is built in Unity using River Auto Material 2019 and Mountain Environment . The

river spline tool is provided by the former package to support fast, flexible, and scriptable river building with various realistic water textures and shaders. The latter package provides a huge library of scanned forest assets like trees, branches, logs, and rocks, which helps to enrich the biomes of the riverine scene, as well as makes it more alike the real-world riverine scene components and appearance, as shown in Fig. 2a.

Aside from the realistic river and terrain trees, the riverine simulation environment features several additional aspects that not only mimic the general objects in real-world rivers, creeks, and streams, but also make the training of river following agents more difficult: 1) river channels with varied widths and depths, 2) river tributaries that diverge from the main water stream, 3) bridges over the river that can obstruct the navigable path of UAVs, and 4) sharp and slow turns.

Several environment design choices are made for better photo-realism and more robust training. Instead of continuous velocity control, discrete waypoint control for UAV movement is used, which eliminates the efforts of controller dynamics adaptation in real-world riverine environments with various wind speeds, water surface effects, etc. Thus we will name the UAV agent as the camera agent in the contents hereafter. The circular shape of the river in our simulated environment is designed for better camera view coherence without seeing too many unrealistic un-modeled parts (sky, terrains, etc.) as the camera agent moves to any valid pose along the river. The annulus design of the riverine environment guarantees consistent episodic rewards achievable regardless of the initial reset position of the camera agent at the beginning of each episode.

TABLE I: Multi-discrete action values of the camera agent.

Action	0	1	2
Up-down Translation (m)	1	0	-1
Horizontal Rotation (deg)	10	0	-10
Forward-backward Translation (m)	1	0	-1
Leftward-rightward Translation (m)	0.5	0	-0.5

### B. Riverine Camera Agent

The simulation agent, created using Unity ML-Agents Toolkit [6], is a camera that looks forward with a fixed pitch angle (20 degrees in our experiments) with respect to the drone frame.

The camera agent is controlled by a four-dimensional multi-discrete action that maps from the command inputs of the common remote controller of UAVs: 1) translation up and down - throttle, 2) horizontal rotation counter-clockwise and clockwise - yaw rate, 3) translation forward and backward - pitch, 4) translation leftward and rightward - roll. Each action branch has three discrete values  $\{0, 1, 2\}$ , where 1 means no action on that branch. The specific values of movement caused by actions are in Table I.

### C. Riverine Gym Environment

A river can be abstracted as a 3D Catmull-Rom spline [33] that is interpolated from several control points. The river




Fig. 3: Demonstrative diagram of valid activity space and acceptable yaw range of camera agent in the river following task.  $h_1 = 1m, h_2 = 15m, \alpha = 150^\circ$  in our experiments.

surface is formed if we stretch this spline bi-directionally along some direction. This closed spline is lifted up to form a 3D bounding volume to limit the height of the camera agent, as well as its horizontal position. The tangents along the spline provide failure judgment when the camera agent's yaw angle is far from either tangent angle of the nearest point on this spline with the projected agent position. This prohibits the turnaround motion of the agent since the built river is circular so that the agent can finish the whole-river navigation without any turnaround. Lingering around or doing continuous meaningless movements with no reward improvement (over 50 steps in our experiments) will also trigger environment reset, and so does collision with bridges. If the environment "done" is triggered during training, the camera agent will be reset at a random point inside the spline-based bounding volume, with a random yaw near either spline tangent of the nearest spline point. The above criteria for the river following task can be better visualized in Fig. 3.

The river central spline can be simplified as a sequence of connected line segments with nearly the same length. To have a decent baseline RL algorithm to compare, and to intuitively sense the learning capabilities of different learning algorithms, shaped reward is given to the camera agent if the projection of its positions on the spline before and after an action covers any un-visited line segments. The line segment where the camera agent's current position projects onto is denoted as visited. If the camera agent has visited all the river spline segments, it will get the highest reward, then the environment will be reset. For collision detection, a mesh collider is used for bridges, and a box collider with a side length of 0.5m is used for the camera agent. The overall reward function for the autonomous river following task is given below.

$$r(s, a) = \begin{cases} -1, & \text{agent fails} \\ r_{shaped} & \text{agent covers un-visited spline segment} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Agent fails when 1) out of spline bounding volume, 2) yaw diverts too much, 3) collides with bridge, and 4) no progress over certain steps.  $r_{shaped}$  is the percentage of newly-visited segments  $n$  after the action  $a$  over all segments number  $N$ , multiplied by a constant value, 10 in this case:  $r_{shaped} = 10 * n/N$ . All negative reward conditions denote the reset of the environment (Eq. 1).

#### D. CliffCircular Environment

CliffCircular (Fig. 2b) is derived from the CliffWalking environment [34], [35] and tailored for the track following task. Like the camera agent in the riverine environment, the agent in CliffCircular must navigate along walkable grids without stepping onto cliffs until it circumnavigates the central cliff block. Additionally, CliffCircular presents partial observability and non-Markovian characteristics due to the local observation of cliffs surrounding the agent’s current grid, the flexibility in determining the circulation direction, and the nature of the track following task, which necessitates consideration of historical information.




Fig. 4: CliffCircular environment diagram. Central track grids, agent location and observation grids are laid on top of the color layers of obstacle and walkable grids.

CliffCircular is a 12x12 grid world, illustrated in Fig. 4. The observation space is the flattened 5x5 grids around the agent indicating cliff. Action space includes five discrete operations: no operation, up, right, down and left. The +1 reward will be given if agent’s projection on the central track grid is newly visited, -100 if agent falls off cliff, otherwise 0. Episode will be reset if cliff is triggered or all central track grids are visited. At the beginning of each episode, the agent’s location and two additional cliff grids are randomly initialized. This setup encourages the policy to learn a safe and robust track following strategy. A TimeLimit environment wrapper of length 128 is used.

## IV. METHODOLOGY

The proposed method, illustrated in Fig. 1 and detailed in Algorithm 1 for solving the river following task, relies on several components: a demonstration dataset provided by a human expert, a BC expert offering guidance, and a PPO agent responsible for environment exploration, dataset updates, and policy learning. The BC expert is firstly trained on the demonstration dataset,  $\mathcal{D}^E : \{\tau_i\}_N$  for  $N$  successful trajectories, resulting in the final form of StaticBC and initial form of DynamicBC baselines. Here StaticBC is the expert

policy that is freezed during the learning process of the PPO agent, whereas DynamicBC is the expert policy that will be trained on the good trajectories experienced by the PPO agent. Each trajectory is a sequence of state-action pairs  $\tau = s_1, a_1, \dots, s_t, a_t$ . The goal of the expert policy  $\pi^E$  is to maximize the log-likelihood over  $\mathcal{D}^E$ .

$$\pi^E = \underset{\pi}{\operatorname{argmax}} \mathbb{E}_{(a_t, s_t) \in \mathcal{D}^E} [\log \pi(a_t | s_t)] \quad (2)$$

To enhance the performance of the expert policy during PPO training, trajectories with episodic rewards exceeding the reward threshold  $T_{reward}$  are added to dataset  $\mathcal{D}^E$ .  $T_{reward}$  is set to 4 in the riverine environment, 18 in the CliffCircular environment. Evaluation of the PPO agent occurs every 10000 training steps in the CliffCircular environment and every 5000 steps in the riverine environment, with more frequent updates to the expert policy in the latter due to its higher complexity. In the simpler CliffCircular environment with multi-discrete observation, duplicate transitions are removed from dataset  $\mathcal{D}^E$  to mitigate bias. The expert policy undergoes retraining using the latest 2000 steps from the dataset for 20 epochs.

---

#### Algorithm 1 Synergistic Reinforcement and Imitation Learning

---

Collect the transition dataset  $\mathcal{D}^E$  by the human expert  
 Train the IL policy  $\pi^E$   
 Initialize policy network  $\phi$ , value network  $\psi$  of RL agent  
**repeat**  
   Observe state  $s_t$   
   Generate action  $a_t$  using  $\pi^\phi$   
   Execute the action  $a_t$  in the environment  
   Observe the next state  $s_{t+1}$  and reward  $r_t$   
   Store transition  $(s_t, a_t, s_{t+1}, r_t)$  in the agent buffer  $\mathcal{D}^R$   
**if** time to update RL agent **then**  
   Sample a mini-batch  $\mathcal{B}$  from  $\mathcal{D}^R$   
   Generate actions  $a_t^E$  using  $\pi^E$   
   Calculate the cross entropy  $H(\pi^E, \pi^\psi)$  (Eq. 4)  
   Calculate loss  $\mathcal{L}$  (Eq. 5)  
   Update value network  $\psi$  and policy network  $\phi$   
**end if**  
**if** time to evaluate RL agent **then**  
   Sample trajectories using  $\pi^\phi$   
   **if**  $\tau$  has a reward above  $T_{reward}$  **then**  
     Append the trajectory  $\tau$  in  $\mathcal{D}^E$   
     Remove repeated transitions from the  $\mathcal{D}^E$   
     Update the IL policy  $\pi^E$  (Eq. 2)  
   **end if**  
**end if**  
**until** Convergence or maximum training steps

---

The target of PPO is to obtain a policy, defined as a distribution over actions conditioned on states  $\pi(a_t | s_t)$ , to

maximize the long-term discounted cumulative rewards

$$\max_{\pi} \mathbb{E}_{\tau \sim p_{\pi}(\tau)} \left[ \sum_{t=0}^T \gamma^t r(s_t, a_t) \right] \quad (3)$$

where  $\tau$  is a trajectory and  $p_{\pi}(\tau)$  is the distribution of the trajectory under policy  $\pi$ , and  $T$  is the time horizon.

To encourage the PPO policy to be close to the policy learned from the decent demonstrations, we add a penalty term in the loss function

$$\mathcal{L}_{action} = H(\pi^E, \pi^{\psi}) = - \sum \pi^E(a_t | s_t) \log \pi^{\psi}(a_t | s_t) \quad (4)$$

where  $H(\pi^E, \pi^{\psi})$  is the cross entropy that measures the difference between the expert policy  $\pi^E$  and the RL policy  $\pi^{\psi}$ .

Combining the above loss functions, PPO+DynamicBC becomes

$$\mathcal{L} = w_1 \mathcal{L}_{policy} + w_2 \mathcal{L}_{value} + w_3 \mathcal{L}_{action} \quad (5)$$

where we keep the policy loss  $\mathcal{L}_{policy}$  and value loss  $\mathcal{L}_{value}$  of PPO intact. The weights  $w_1$ ,  $w_2$ , and  $w_3$  correspond to three losses, with  $w_1$  and  $w_2$  set to 1, and  $w_3$  initially set to 1. The PPO policy is evaluated every 10000 training steps for 20 epochs. If the PPO outperforms BC in mean episodic reward,  $w_3$  is reduced to 0.2 to encourage exploration. PPO is updated every 1024 training steps.

## V. SIMULATION VERIFICATION

In the riverine environment, observation images are resized to 128 x 128 for both encoding and RL training, and the encoded image is a vector of 1024 float numbers. VAE model is obtained by training 2000 randomly collected images from the environment. The encoded mean vector (without Gaussian noise) is used as the observation input to RL algorithms. 50 trajectories each with 50 steps are collected by a human expert to form the dataset  $\mathcal{D}^E$ . After 50 epochs of training, the expert policy  $\pi^E$  has a success rate of 93.27% to predict the expert action from the demonstration and is able to achieve a mean reward of 0.60 in the riverine environment. This expert policy will be used to generate reference actions during the training of RL agent.

In the CliffCircular environment, 100 trajectories are collected by an expert, which achieves a mean episode reward of approximately 0. The BC expert, pre-trained on this dataset, attains a reward of  $-0.24$ , indicating a suboptimal policy that is capable of accomplishing the track following objective but commits few errors occasionally. Code related to RL and IL is based on StableBaseline3 [36] and Imitation Learning Baseline Implementations [37].

## VI. RESULTS AND ANALYSIS

### A. Training Results

The proposed method (PPO+DynamicBC), together with baseline methods (StaticBC, DynamicBC, PPO, and PPO+StaticBC), are used to train the discrete-action agents in both environments. All the methods have been trained for

14000 steps in the riverine environment and 10000 steps in the CliffCircular environment for 5 different seeds.

Fig. 5 illustrates the training results of different methods in two environments. The training curve is derived by averaging results from 5 different seeds and is depicted as a solid line representing the mean reward, with a shaded region around it indicating the standard error. The episodic reward is computed by averaging over 100 past episodes, with updates occurring every 1024 steps. Additionally, the green horizon line signifies the performance achieved by a pre-trained (BC) model.

In the CliffCircular environment, the PPO baseline slowly improves and surpasses the pre-trained BC baseline after 85000 steps. The PPO+DynamicBC initially returns rewards similar to PPO+StaticBC baseline. However, PPO+DynamicBC’s learning curve shows a higher increasing speed after 10000 steps and finally reaches a mean episodic reward of 17.90 at 85000 steps, while the maximum possible reward is 20. The best mean episodic reward of PPO+StaticBC is less than 11.31. After 10000 steps of training, PPO+DynamicBC’s mean episode length is 28.50, while PPO+StaticBC and PPO’s mean episode lengths are 32.62 and 88.40, respectively. This suggests that the proposed method has acquired a more efficient policy for completing the task in fewer steps, whereas PPO+StaticBC has developed a less optimal policy, requiring an additional 4.12 steps to complete the task. Evidently, PPO encounters challenges in learning an effective policy within this POMDP track following environment with non-Markovian rewards.

In the river following environment, PPO baseline slowly improves throughout the reinforcement learning phase but is still lower than the pre-trained BC baseline. The PPO+DynamicBC initially returns rewards similar to PPO+StaticBC baseline. Both methods can get a reward at or above the BC after 25000 steps. However, after the PPO agent’s performance surpasses the BC, PPO+DynamicBC’s learning curve shows a higher increasing speed and finally reaches a mean episodic reward of 2.0 at 105000 steps, while the best mean episodic reward of PPO+StaticBC is less than 1.5. PPO+StaticBC also shows multiple decrease trends during training, mainly because of the constraint of the static expert policy. PPO+DynamicBC shows one decrease trend (110000 to 130000) however it recovered to 1.7 after 10000 steps. On average, PPO+DynamicBC’s score is 1.5 higher than PPO, 0.6 higher than BC, and 0.4 higher than PPO+StaticBC.

### B. Evaluation Results

As shown Fig. 6, the reward and navigation efficiency are compared through violin plots. To generate the episodic reward distribution, all methods with the best weights are evaluated for 50 epochs in the same seed. In the Cliff-Circular environment, BC has both high and low scores, resulting in a relatively low mean reward. PPO without guidance learns how to avoid collisions but cannot figure out the task objective. PPO+StaticBC is able to increase the chance of finishing the task and decrease the chance of




Fig. 5: Training results of different approaches in the Cliff-Circular environment and river following environment.

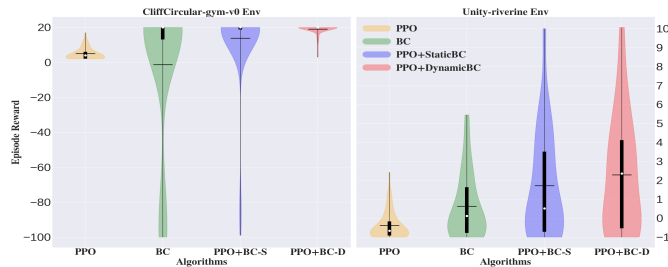


Fig. 6: Episode reward distributions of four algorithms on CliffCircular-gym-v0 and riverine environments. The horizontal bar represents the mean, the vertical bar represents 25% and 75% percentiles, and the white dot is the median. Statistics from 50 episodes for each algorithm.

collision compared to StaticBC. PPO+DynamicBC achieves the best result. It learns an optimal policy to finish the track following task while avoiding cliffs. In the riverine environment, BC and the vanilla PPO can never achieve any reward larger than 6 while both algorithms PPO+StaticBC and PPO+DynamicBC can achieve higher episodic reward of 10. Noticeably, PPO+DynamicBC achieves the highest mean and median episodic rewards compared to all other methods.

This verifies the superiority of PPO+DynamicBC algorithm since the IL (BC) model is evolving itself while guiding RL (PPO) model, which acts as a better regularizer compared to PPO+StaticBC algorithm. In contrast, PPO and BC baseline algorithms have inferior performance due to the shortage of runtime expert guidance and limited generalization ability, respectively. Moreover, this synergistic combination of PPO and BC is validated to largely boost the training speed and final model performance.

To qualitatively compare the navigation efficiency and failure modes of both dynamic and static algorithms, we plot their trajectories in the riverine environment when starting at the same locations (Fig. 7). The static algorithm agent can only navigate about 70% of river segments in both experiments, before it failed due to exceeding river spline volume (distracted by the tributary 2b) and bumping into the bridge, respectively. This is better visualized by the agent view near blue circles. The dynamic algorithm agent not only finishes whole-river navigation in both experiments but also with fewer steps, which justifies its robustness and efficiency. This can be visualized by the leading checkpoints (star in the




Fig. 7: Trajectories of two algorithms PPO+StaticBC and PPO+DynamicBC starting from the same pose in the Unity riverine environment, in two evaluation experiments. Star symbols signify the agent’s position every 50 steps. Trajectory color gets darker as the agent moves further away from the starting position.

figure) of the dynamic algorithm agent, meaning the agent moves a longer effective distance with fewer meaningless movements within the same 50 steps than the static agent. Besides, the dynamic algorithm handles better around river corners with smoother trajectory and larger safe margins to river banks. It also rejects tributary when seeing the main stream ahead by timely heading correction. Both cases can be viewed in figures near the red circles in Fig. 7.

## VII. CONCLUSION

In this paper, we develop a photo-realistic riverine simulation environment to train a UAV agent to autonomously follow the river while avoiding obstacles. A novel framework that synergistically combines RL and IL is proposed, and evaluated in both a riverine environment and a customized grid world track following environment. The performance of the proposed method (PPO+DynamicBC), compared with the baseline methods (PPO, StaticBC, DynamicBC and PPO+StaticBC) in both environments, achieves the best qualitative and quantitative results. In the riverine environment, it achieves at least 0.4 more rewards than all other methods on average, and the highest best and median rewards. In the CliffCircular environment, this algorithm finishes the track following task with nearly full score (episode reward 19.04 out of 20), with significantly lower reward variance than all baseline algorithms.

Further research can improve the proposed method in several aspects. Firstly, instead of relying solely on the episodic reward threshold to filter trajectories, specific transitions can also be evaluated by some advantage measure and the valuable ones can be merged to the dataset for IL expert refinement. This could reduce the demonstration buffer size, bias of this dataset, and the overall training time. Secondly, IL expert’s policy might not consistently outperform that of the RL agent during training process. Therefore, when to utilize the action loss for updating the RL agent could be further investigated. Thirdly, observation augmentation (with historical observations or actions) techniques can be implemented to push the partially observable non-Markovian decision process to MDP for convergence guarantee in the track following tasks.

## REFERENCES

- [1] M. E. Hodgson, N. I. Vitzilaios, M. L. Myrick, T. L. Richardson, M. Duggan, K. R. I. Sanim, M. Kalaitzakis, B. Kosaraju, C. English, and Z. Kitzhaber, "Mission planning for low altitude aerial drones during water sampling," *Drones*, vol. 6, no. 8, p. 209, 2022.
- [2] S. I. Ullah and I. Fullfilment, "Vision-based autonomous mapping & obstacle avoidance for a micro-aerial vehicle (mav) navigating canal," 2019.
- [3] T. Huang, Z. Chen, W. Gao, Z. Xue, and Y. Liu, "A usv-uav cooperative trajectory planning algorithm with hull dynamic constraints," *Sensors*, vol. 23, no. 4, p. 1845, 2023.
- [4] J. Li, G. Zhang, and B. Li, "Robust adaptive neural cooperative control for the usv-uav based on the lvs-lva guidance principle," *Journal of Marine Science and Engineering*, vol. 10, no. 1, p. 51, 2022.
- [5] A. Gonzalez-Garcia, A. Miranda-Moya, and H. Castañeda, "Robust visual tracking control based on adaptive sliding mode strategy: Quadrotor uav-catamaran usv heterogeneous system," in *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2021, pp. 666–672.
- [6] A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Elion, C. Goy, Y. Gao, H. Henry, M. Mattar *et al.*, "Unity: A general platform for intelligent agents," *arXiv preprint arXiv:1809.02627*, 2018.
- [7] D. Neider, J.-R. Gaglione, I. Gavran, U. Topcu, B. Wu, and Z. Xu, "Advice-guided reinforcement learning in a non-markovian environment," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, 2021, pp. 9073–9080.
- [8] M. Prajapat, M. Mutný, M. N. Zeilinger, and A. Krause, "Submodular reinforcement learning," *arXiv preprint arXiv:2307.13372*, 2023.
- [9] B. Kang, Z. Jie, and J. Feng, "Policy optimization with demonstrations," in *International conference on machine learning*. PMLR, 2018, pp. 2469–2478.
- [10] M. Jing, X. Ma, W. Huang, F. Sun, C. Yang, B. Fang, and H. Liu, "Reinforcement learning from imperfect demonstrations under soft expert guidance," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 04, 2020, pp. 5109–5116.
- [11] K. Lin, D. Li, Y. Li, S. Chen, Q. Liu, J. Gao, Y. Jin, and L. Gong, "Tag: Teacher-advice mechanism with gaussian process for reinforcement learning," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [12] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [13] D. A. Pomerleau, "Efficient training of artificial neural networks for autonomous navigation," *Neural computation*, vol. 3, no. 1, pp. 88–97, 1991.
- [14] A. Taufik, S. Okamoto, and J. H. Lee, "Multi-rotor drone to fly autonomously along a river using a single-lens camera and image processing," *International Journal of Mechanical Engineering*, vol. 4, no. 6, pp. 39–49, 2015.
- [15] A. Taufik, "Multi-rotor drone to fly autonomously along a river and 3d map modeling of an environment around a river," 2016.
- [16] Y. Song, S. Najj, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A flexible quadrotor simulator," in *Conference on Robot Learning*. PMLR, 2021, pp. 1147–1157.
- [17] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.
- [18] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 097–11 107.
- [19] E. Marchesini and A. Farinelli, "Discrete deep reinforcement learning for mapless navigation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 688–10 694.
- [20] P. Zieliński and U. Markowska-Kaczmar, "3d robotic navigation using a vision-based deep reinforcement learning model," *Applied Soft Computing*, vol. 110, p. 107602, 2021.
- [21] P. Wei, R. Liang, A. Michelmore, and Z. Kong, "Vision-based 2d navigation of unmanned aerial vehicles in riverine environments with imitation learning," *Journal of Intelligent & Robotic Systems*, vol. 104, no. 3, p. 47, 2022.
- [22] J. R. Liang, *Vision-Based Unmanned Aerial Vehicle Navigation in Virtual Complex Environment using Deep Reinforcement Learning*. University of California, Davis, 2021.
- [23] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [24] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [25] J. Li, J. Chavez-Galaviz, K. Azizzadenesheli, and N. Mahmoudian, "Dynamic obstacle avoidance for usvs using cross-domain deep reinforcement learning and neural network model predictive controller," *Sensors*, vol. 23, no. 7, p. 3572, 2023.
- [26] S. Aradi, "Survey of deep reinforcement learning for motion planning of autonomous vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 2, pp. 740–759, 2020.
- [27] S. Feng, H. Sun, X. Yan, H. Zhu, Z. Zou, S. Shen, and H. X. Liu, "Dense reinforcement learning for safety validation of autonomous vehicles," *Nature*, vol. 615, no. 7953, pp. 620–627, 2023.
- [28] A. Hussein, E. Elyan, M. M. Gaber, and C. Jayne, "Deep imitation learning for 3d navigation tasks," *Neural computing and applications*, vol. 29, pp. 389–404, 2018.
- [29] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [30] J. Kober, B. Mohler, and J. Peters, "Imitation and reinforcement learning for motor primitives with perceptual coupling," in *From motor learning to interaction learning in robots*. Springer, 2010, pp. 209–225.
- [31] Z. Huang, J. Wu, and C. Lv, "Efficient deep reinforcement learning with imitative expert priors for autonomous driving," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [32] V. G. Goecks, G. M. Gremillion, V. J. Lawhern, J. Valasek, and N. R. Waytowich, "Integrating behavior cloning and reinforcement learning for improved performance in dense and sparse reward environments," in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 465–473.
- [33] E. Catmull and R. Rom, "A class of local interpolating splines," in *Computer aided geometric design*. Elsevier, 1974, pp. 317–326.
- [34] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. d. Cola, T. Deleu, M. Goulão, A. Kallinteris, A. KG, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, "Gymnasium," Mar. 2023. [Online]. Available: <https://zenodo.org/record/8127025>
- [35] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [36] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: <http://jmlr.org/papers/v22/20-1364.html>
- [37] A. Gleave, M. Taufeeque, J. Rocamonde, E. Jenner, S. H. Wang, S. Toyer, M. Ernestus, N. Belrose, S. Emmons, and S. Russell, "imitation: Clean imitation learning implementations," arXiv:2211.11972v1 [cs.LG], 2022. [Online]. Available: <https://arxiv.org/abs/2211.11972>