

Adaptive Risk-Tendency: Nano Drone Navigation in Cluttered Environments with Distributional Reinforcement Learning

Cheng Liu¹, Erik-Jan van Kampen¹, Guido C.H.E. de Croon¹

Abstract—Enabling the capability of assessing risk and making risk-aware decisions is essential to applying reinforcement learning to safety-critical robots like drones. In this paper, we investigate a specific case where a nano quadcopter robot learns to navigate an apriori-unknown cluttered environment under partial observability. We present a distributional reinforcement learning framework to generate adaptive risk-tendency policies. Specifically, we propose to use lower tail conditional variance of the learnt return distribution as intrinsic uncertainty estimation, and use exponentially weighted average forecasting (EWAf) to adapt the risk-tendency in accordance with the estimated uncertainty. In simulation and real-world empirical results, we show that (1) the most effective risk-tendency varies across states, (2) the agent with adaptive risk-tendency achieves superior performance compared to risk-neutral policy or risk-averse policy baselines. Code and video can be found in this repository: <https://github.com/tudelft/risk-sensitive-rl.git>

I. INTRODUCTION

Reinforcement learning (RL) is promising in solving sequential decision making problems such as robotic navigation with obstacle avoidance as it seeks long-term optimal policies [1], [2]. Recent advances in deep reinforcement learning (DRL), which combines deep neural networks with RL, have shown the capability of achieving super human performance in diverse complex environments [3], [4], [5]. The majority of current DRL methods are designed for maximizing the expectation of accumulated future returns, omitting to consider the risk of rare catastrophic events. However, when it comes to applying RL to safety-critical robots like drones, instead of aiming at achieving a high expected return, dealing with risks and making decisions under uncertainty is crucial and remains a challenge.

A natural way to risk-sensitive RL is considering the worst-case of the stochastic return rather than its expectation, but this may lead to over-conservative policies [6]. Recent works propose to model the distribution of the future return and to generate multiple policies with different risk-sensitivities by changing levels of a risk metric [7]. While [7] captures the stochasticity in accumulated returns by approximating the mean and variance of a gaussian distribution, *distributional RL* reconstructs the true intrinsic distribution of future returns [8], [9], [10]. A major merit of distributional RL is that it can generate multiple policies with different levels of risk-tendencies [11], [10], [12].

Distributional RL has been applied to safety-critical applications such as autonomous driving at occluded intersections [14] and mobile-robot indoor navigation [15]. These

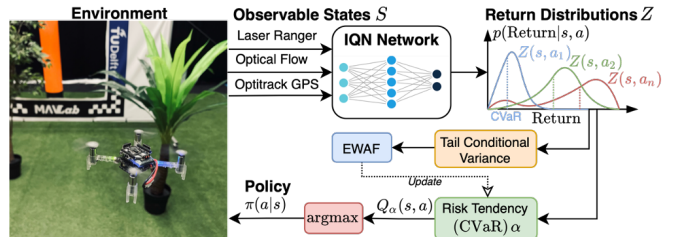


Fig. 1: ART-IQN framework that combines IQN with intrinsic uncertainty estimation and risk-tendency adaption, enabling a Crazyflie [13] nano drone navigating through a cluttered environment safely and efficiently.

methods learn a policy that can vary its risk-tendency during training, but they still rely on a fixed risk-tendency for each deployment task. However, a proficient pilot would not be as cautious while cruising in fair weather as when landing in stormy weather. In other words, the ideal degree of risk-tendency varies as a function of not only the task but also the real-time feedback from the environment. A step towards building intelligent robots is *adapting risk-tendency* on the fly automatically.

To achieve this goal, we propose Adaptive Risk-Tendency Implicit Quantile Network (ART-IQN) that can adapt risk-tendency by reacting to the context. We propose to let intrinsic uncertainty [16] (estimated by lower tail conditional variance) set the way in which the agent acts - adapting risk-tendency by forecasting the intrinsic uncertainty. The effectiveness of ART-IQN is validated on safety-critical tasks - autonomous drone navigation in cluttered environments with constrained sensors (shown in Fig. 1). Both in simulation and real-world experiments, our method shows superior performance in the trade-off between navigation efficiency and safety in comparison with risk-neutral and risk-averse baselines. Our main contributions are:

- an automatic adaption in risk-tendency on the fly in accordance with intrinsic uncertainty estimation;
- a drone navigation algorithm based on distributional RL, that can learn a variety of risk-sensitive policies;
- a sim-to-real RL framework and a light-weight simulation environment, enabling seamless policy generalization from simulation to reality.

II. RELATED WORK

A. Risk and Uncertainty in RL-based Navigation

RL-based robot navigation methods surged recently due to the capability of generalization and robustness[17], [18].

¹ Delft University of Technology, Email: c.liu-10@tudelft.nl, e.vankampen@tudelft.nl, g.c.h.e.decroon@tudelft.nl.

Several navigation and obstacle avoidance algorithms based on RL have also emerged to address risks and uncertainties in the environment. For instance, [19] proposed to use a neural network to predict collision probability at future steps for obstacle avoidance tasks, utilizing MC-dropout [20] and bootstrapping [21] to estimate the uncertainty of the model prediction. Additionally, [22] enabled estimation of the regional increase of uncertainty in novel dynamic scenarios by introducing LSTM [23] to add memory of historical motion of the robot. [24] resorts to a model-free policy network as action selector, a GRU [25] to predict uncertainty in local observation and uses the prediction variance to adjust the variance of the stochastic policy.

However, these methods either use Model Predictive Control (MPC) [26] as action selector, which is still computationally heavy for the micro controllers present on nano copters, or they require an additional predictor model to estimate the uncertainty. In our method, the risk measure and uncertainty estimation are easily and efficiently implemented based on distributional RL (more details in Section III), which requires minimal additional computational resources. Besides, our method can be applied to other tasks with more complex objective functions and constraints while MPC will consume much more resources.

B. Distributional Reinforcement Learning

Distributional RL has gained momentum recently, which takes into account the whole distribution of value functions rather than the expectation [7], [8], [9], [10]. Since the whole distribution contains more information beyond the first moment, one can utilize it to make more informed decisions that lead to higher rewards. Recent literature shows that similar mechanisms also exist in human brains [27].

A forerunner of distributional RL is categorical DQN [8], which uses categorical distribution with fixed supports to approximate the probability density function (PDF) of the return. A more flexible way to approximate the distribution is quantile regression [28]. For instance, the quantile regression DQN (QR-DQN) algorithm [9] learns the distribution by approximating the quantile function (QF) with fixed quantile values. The implicit quantile network (IQN) algorithm [10] further improved the flexibility and approximation accuracy compared to QR-DQN by learning quantile values from quantile fractions sampled from a uniform distribution $\mathcal{U}[0, 1]$. This is achieved with a deep neural network representing the QF by mapping quantile fractions to quantile values under Wasserstein distance [10], a loss metric which indicates the minimal cost for transporting mass to make two distributions identical.

There are also applications of distributional RL to safety-critical environments in the literature. [14] incorporates IQN to solve an autonomous driving task at intersections by combing risk-averse IQN with safety guarantees. Based on [12], [15] proposes a method enabling a mobile robot navigating office scenarios with multiple risk-sensitivities.

Even though risk-tendency can be altered without retraining a policy, those methods require a fixed risk-tendency

for each deployment task. Our algorithm is able to adjust its risk-tendency by reacting to dynamic uncertainty levels rather than following a fixed manually set risk-tendency.

III. METHODOLOGY

A. Problem Statement

We formulate the drone navigation task as Partially Observable Markov Decision Process (POMDP) [29].

1) *POMDP Setup*: The POMDP can be defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{P}, R, \gamma)$, where \mathcal{S} , \mathcal{A} and \mathcal{O} represent the state, action and observation spaces. The drone interacts with the environment in discrete timesteps. At each timestep t , it receives the observation $o_t \in \mathcal{O}$ from the environment and performs an action $a_t \in \mathcal{A}$ based on its policy function $\pi_t(a_t|o_t)$, which causes a transition of the state from s_t to $s_{t+1} \sim \mathcal{P}(\cdot|s_t, a_t)$, generating a reward $r_t = R(s_t, a_t)$ and a new observation $o_{t+1} \sim \mathcal{O}(\cdot|s_{t+1}, a_t)$. Following policy π , the discounted sum of future rewards is denoted by the random variable $Z^\pi(s_t, a_t) = \sum_{k=0}^{\infty} \gamma^k R(s_{t+k}, a_{t+k})$ with $\gamma \in (0, 1)$ as the discount factor. Standard RL aims at maximizing the expectation of Z^π , which is known as the action-value function $Q^\pi(s_t, a_t) = \mathbb{E}[Z^\pi(s_t, a_t)]$.

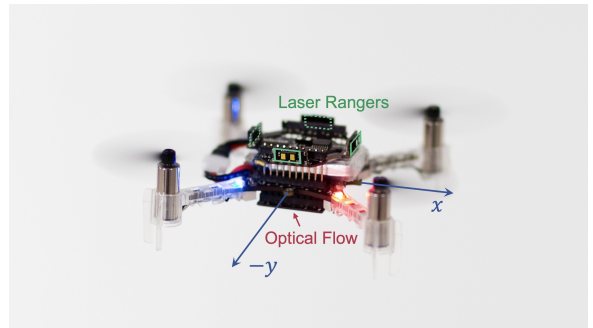


Fig. 2: A Crazyflie with 4 lasers to detect obstacles. [Picture by Guus Schoonewille, reprinted TU Delft]

2) *States and Observations*: We use Crazyflie nano quadrotor as our experiment platform. As shown in Fig. 2, the Crazyflie is equipped with four lasers in the drone's positive and negative x and y axis to detect obstacles. It also has an optical flow camera to estimate velocity for low-level flight control. Given a navigation task, \mathcal{S} contains information about the drone itself, the goal and obstacles. The state can be parameterized as $s_t = \langle \mathbf{p}, d_g, \mathbf{d}_o \rangle$, where \mathbf{p} is drone's global position, $d_g = \|\mathbf{p} - \mathbf{p}_g\|_2$ is the distance from the drone to the goal, \mathbf{d}_o is a vector consisting of distances from the drone to surrounding obstacles.

Due to the constrained sensors onboard, \mathcal{S} is not fully observable to the drone. Instead, the drone receives a partial observation which is formulated as a tuple $o_t = \langle \mathbf{p}, d_g, \mathbf{d}_l \rangle$ where \mathbf{d}_l denotes laser reflections. The laser detects obstacles at a maximum range of 4 meters. \mathbf{p} and d_g are given by a global motion capturing system in real-world experiments.

3) *Action Space*: To incorporate our algorithm, \mathcal{A} consists of discretized velocities with multiple magnitudes and directions. The velocity magnitudes that can be chosen are m discretized values exponentially spaced in $(0, v_m]$, in

which v_m is the maximum velocity. Since only obstacles that intersect with the laser beams will be detected, there are 4 possible moving directions evenly spaced between $[0, 2\pi)$.

4) *Reward Function*: The reward function is manually designed to award the drone for reaching the goal as fast as possible, while penalizing for collisions or getting close to obstacles:

$$R(s_t, a_t) = \begin{cases} 50 & d_g < d_f \\ 5(d_o - d_s) & r_d < d_o < d_s \\ -25 & d_o < r_d \\ -0.1 & \text{otherwise,} \end{cases} \quad (1)$$

where d_f is the goal-reaching threshold, d_o is the distance from drone to the closest obstacle, d_s is the safety margin, r_d represents the radius of drone.

B. Adaptive Risk-Tendency Implicit Quantile Network

To adjust the risk-tendency on the fly dynamically, we propose the Adaptive Risk Tendency Implicit Quantile Network (ART-IQN) algorithm. We introduce the key components of ART-IQN as shown in Fig. 1, which are (1) the risk-sensitive IQN, (2) the intrinsic uncertainty estimation and (3) the EWF uncertainty forecasting.

1) *Implicit Quantile Network*: In distributional RL, the distributional Bellman equation [8] can be defined as

$$Z^\pi(s, a) \stackrel{D}{=} R(s, a) + \gamma Z^\pi(s', a'), \quad (2)$$

where $\stackrel{D}{=}$ denotes equality in distribution, state s' and action a' at next timestep are distributed according to $s' \sim \mathcal{P}(s, a)$, $a' \sim \pi(\cdot|s')$.

We represent $Z^\pi(s, a)$ implicitly by its quantile function as in IQN [10] because of the easiness to combine IQN with risk metrics to adapt risk-tendency, which will be further explained in Section III-B.2. Concretely, the quantile function is approximated by a neural network with learnable parameters θ . We express such implicit quantile function as $Z_\theta^\pi(s, a; \tau)$, where $\tau \in [0, 1]$ is the quantile level. To optimize θ , quantile regression [28] is used with quantile Huber-loss as a surrogate of the Wasserstein distance [10].

A neural network with parameters θ' is used as the target distribution approximator and the temporal difference (TD) at sample (s, a, r, s') is computed as

$$\delta_{\tau, \tau'} = r + \gamma Z_{\theta'}^\pi(s', a'; \tau') - Z_\theta^\pi(s, a; \tau), \quad (3)$$

for τ, τ' independently sampled from the uniform distribution, i.e. $\tau, \tau' \sim \mathcal{U}[0, 1]$.

The τ -quantile Huber-loss is defined as

$$\rho_\kappa(\delta; \tau) = |\tau - \mathbb{I}\{\delta < 0\}| \frac{\mathcal{L}_\kappa(\delta)}{\kappa}, \text{ with} \quad (4)$$

$$\mathcal{L}_\kappa(\delta) = \begin{cases} \frac{1}{2}\delta^2 & \text{if } |\delta| \leq \kappa \\ \kappa(|\delta| - \frac{1}{2}\kappa) & \text{otherwise,} \end{cases}$$

where \mathbb{I} is an indicator operator. The threshold κ provides smooth gradient-clipping. We approximate the quantile loss by sampling N independent quantiles τ and N' independent targets τ' . The loss function to update θ is

$$\mathcal{L}(\theta) = \frac{1}{N \cdot N'} \sum_{i=1}^N \sum_{j=1}^{N'} \rho_\kappa(\delta_{\tau_i, \tau'_j}; \tau_i). \quad (5)$$

By backpropagating $\mathcal{L}(\theta)$ with respect to θ , the Wasserstein distance is minimized between the current return distribution $Z^\pi(s, a)$ and the target $R(s, a) + \gamma Z^\pi(s', a')$.

2) *Risk-sensitive Policy and Risk Metric*: Distributional RL is inherently risk-sensitive by combining *risk metrics* [6] to create a *distorted expectation* [12] on the return distribution. A distorted expectation is a risk weighted expectation of the distribution under a specific distortion function - which indicates a non-decreasing function $\beta : [0, 1] \rightarrow [0, 1]$ satisfying $\beta(0) = 0$ and $\beta(1) = 1$. The distorted expectation of Z under β is defined as $Q_\beta = \int_0^1 F_Z^{-1}(\tau) d\beta(\tau)$, where $F_Z^{-1}(\tau)$ is the quantile function or cumulative density function. According to [10], any distorted expectation can be represented as a weighted sum over the quantiles. A corresponding sample-based risk-sensitive policy is obtained by approximating Q_β by K samples of $\tilde{\tau} \sim \mathcal{U}[0, 1]$:

$$\pi_\beta(s) = \arg \max_{a \in \mathcal{A}} \frac{1}{K} \sum_{k=1}^K Z_{\beta(\tilde{\tau}_k)}(s, a). \quad (6)$$

Altering the sampling principle for τ creates various risk-sensitive policies. Specifically, we consider Conditional Value-at-Risk (CVaR) [30], a *coherent risk metric* [6] as our distortion function for its convenience and diversity in generating risk-tendency levels. CVaR is applied to IQN by modifying $\tilde{\tau} \sim \mathcal{U}[0, 1]$ to $\tilde{\tau} \sim \mathcal{U}[0, \alpha]$, where α is the CVaR value. We get risk-averse policies as α decreases to near zero and reduce back to risk-neutral when $\alpha = 1$.

3) *Lower Tail Conditional Variance for Intrinsic Uncertainty Estimation*: One major source of risk comes from intrinsic uncertainty, which is due to stochasticity of the environment or partial observability. Opposite to *epistemic uncertainty* [16], intrinsic uncertainty is independent of the agent's knowledge about the task.

In distributional RL, a more or less spread out return distribution indicates the level of intrinsic uncertainty [31]. Inspired by [32], where the decaying upper tail conditional variance of the return distribution is used for more efficient exploration, we use the lower tail conditional variance as the intrinsic uncertainty estimation for risk-tendency adaption. The lower half tail conditional variance is equivalent to the *right truncated variance* (RTV):

$$\text{RTV} = \frac{2}{N} \sum_{i=1}^{\frac{N}{2}} (F_Z^{-1}(\tau_i) - F_Z^{-1}(\tau_{\frac{N}{2}}))^2, \quad (7)$$

in which τ_i are the $\frac{i}{N}$ -th quantile levels. Intuitively, RTV is biased towards negative returns. We calculate RTV with respect to the median rather than the mean due to its statistical robustness [32], [33]. Note that $F_Z^{-1}(\tau)$ is implicitly approximated by $Z_\theta^\pi(\cdot; \tau)$ in our method.

4) *EWAF for Risk-Tendency Adaption*: In our framework, the risk-tendency can be cast in the choice of CVaR. To formulate CVaR as a function of RTV, we propose to model CVaR by an exponentially weighted categorical distribution. Specifically, consider a categorical distribution C with two logits: $C_i = \exp(w_i) / \sum_i \exp(w_i)$, in which $w_i \in \mathbb{R}, i = 1, 2$. By letting $\alpha = C_1$, the CVaR is restricted to a range of $(0, 1)$ and can be adjusted by altering the logit weights. Concretely, at each timestep t , the CVaR is adapted by updating w_i with feedback f by a step size η : $w_1 = w_1 - \eta f, w_2 = w_2 + \eta f$. We set $f = \text{RTV}_t - \text{RTV}_{t-1}$ as an indicator of intrinsic uncertainty feedback. To avoid the CVaR approaching zero, an additional term: $\sum_i \exp(w_i)/b$ is added both to the denominator of C_i for $i = 1, 2$ and to the numerator of C_1 , which results in a CVaR range of $(\frac{1}{b+1}, 1)$. For example, $\alpha \in (0.1, 1)$ when $b = 9$.

The overall ART-IQN that can adapt its risk-tendency by reacting to intrinsic uncertainty variations is explained in Algorithm 1. In principle, when RTV is increasing and the current CVaR is relatively large, the agent will behave more risk-unwillingly by choosing smaller CVaR.

Algorithm 1: ART-IQN for Drone Navigation

Input: Post-training IQN θ ; $\mathcal{A}, K, N, w_1, w_2, b, \eta$
Initialize state s , CVaR α
while $d_g > d_f$ and $t \leq H$ **do**
 Observe $o_t \leftarrow \langle \mathbf{p}, d_g, \mathbf{d}_l \rangle$ from s_t
 Get quantile function $Z_\theta(o_t, a; \tau) = \text{IQN}_\theta(o_t; \tau)$
 Distorted sampling $\tilde{\tau}_k \sim \mathcal{U}[0, \alpha], k = \{1, \dots, K\}$
 Take action $a_t = \arg \max_{a \in \mathcal{A}} \frac{1}{K} \sum_{k=1}^K Z_\theta(o_t, a; \tilde{\tau}_k)$
 Calculate right truncated variance:
 $\text{RTV}_t = \frac{2}{N} \sum_{i=1}^{\frac{N}{2}} (Z_\theta(o_t, a_t; \tau_i) - Z_\theta(o_t, a_t; \tau_{\frac{N}{2}}))^2$
 Obtain feedback $f_t = \text{RTV}_t - \text{RTV}_{t-1}$
 Forecasting $w_1 = w_1 - \eta f_t, w_2 = w_2 + \eta f_t$
 Adapt CVaR $\alpha = \frac{(b+1) \exp(w_1) + \exp(w_2)}{(b+1) \sum_i \exp(w_i)}, i = 1, 2$

end

C. Training and Evaluation Pipelines

1) *Environment*: We design an OpenAI gym [34] like 2D environment for training with state, observation and action spaces defined in Section III-A.2. To achieve fast simulation, the drone is modelled as a point mass and the velocity command is assumed to be immediately executed. The state is updated every $T = 0.1s$ in simulation time. We utilize *domain randomization* [35] to train policies that can generalize to diverse scenarios. Specifically, the goal distance is uniformly sampled $d_g \sim \mathcal{U}[5, 7](m)$ with drone initialized at a fixed start point for each training episode. The number of obstacles, the shape and position of each obstacle are randomly generated for each episode as demonstrated in Fig. 3. Laser beams and obstacle outlines are modelled as line segments for simplicity. Gaussian noise $\mathcal{N}(\mu, \sigma)$ with mean $\mu = 0.0$ and standard deviation $\sigma = 0.01$ is added to the measurement of each laser to simulate a noisy sensor.

2) *Training Process*: We follow *curriculum learning* to train the agent - the complexity of the environment increases

as training process goes on. The first stage of training is implemented with relative small number of randomized obstacles, $n_{obs} \in [0, 5]$. After training several episodes (until a navigation success rate of 0.8 is reached), the complexity of the environment is increased by adding more obstacles, $n_{obs} \in [6, 12]$. To make sure the agent accumulates a diverse range of experiences under a variety of risk-tendencies, the CVaR value is uniformly sampled $\alpha \sim \mathcal{U}(0, 1]$ at each episode. An episode is terminated after a collision or if the goal is not reached within H timesteps. The whole training process is ended after the average return is empirically converged. It took ≈ 3.5 hours on a 2.2 GHz Intel i7 Core CPU, achieving a success rate of 0.88.



Fig. 3: Randomly generated environments for training.

The agent is modelled as a fully connected network that has 3 hidden layers with 512 units per layer. Each fully-connected layer is followed with a ReLU [36] activation function except the output layer. Adam [37] is used as our optimizer with lr as the learning rate. At each update step for every D episodes, a batch size of B samples are drawn from the experience replay buffer with size E . The hyper-parameters used are listed in Table I.

TABLE I: HYPER-PARAMETERS

Hyper-parameter symbols and values							
lr	2×10^{-4}	v_m	1 [m/s]	D	5	N, N'	16
E	5×10^4	r_d	0.05 [m]	K	64	m	3
γ	0.99	d_f	0.1 [m]	B	32	b	9
T	0.1 [s]	d_s	0.2 [m]	H	200	η	0.5

3) *Evaluation in Simulation*: To show the efficacy of our algorithm, ART-IQN is compared with IQN with multiple risk-tendencies $\alpha = \{0.1, 0.25, 0.5, 0.75, 1.0\}$. In addition, we also trained a DQN [3] agent as our baseline following the same training procedure. The average episodic return, success rate, collision rate and average navigation time are compared among agents across various environments. Specifically, the evaluation is executed on three sets of environments with $n_{obs} = \{2, 6, 12\}$ for each agent. Each set has 100 diverse randomized environments.

IV. RESULTS

A. Simulation Results

1) *Quantitative Analysis*: Table II gives quantitative results by comparing: (1) DQN, (2) IQN with different risk tendencies and (3) ART-IQN. For all the environments, DQN performs similar to risk-neutral IQN. For $n_{obs} = 2$, all the agents finish the navigation task at a high success rate and a low collision rate since the environment is comparatively easy. For $n_{obs} = 6$, while IQN agent with a lower CVaR value maintain a lower collision rate compared to the one with a higher CVaR value, the task finishing time is longer and the timeout rate is higher. In contrast, ART-IQN achieves

TABLE II: QUANTITATIVE SIMULATION RESULTS

n_{obs}	CVaR	Average episodic return (mean \pm std)			Success rate			Collision rate			Navigation time [s]		
		2	6	12	2	6	12	2	6	12	2	6	12
DQN[3]	-	37.45 \pm 9.01	25.60 \pm 13.91	17.16 \pm 15.85	0.85	0.65	0.56	0.13	0.31	0.39	5.23	8.11	8.34
IQN[10]	0.1	39.42 \pm 7.28	33.49 \pm 12.16	20.42 \pm 12.43	0.87	0.74	0.59	0.09	0.15	0.17	5.41	12.52	18.52
	0.25	40.36 \pm 8.07	32.37 \pm 12.02	21.61 \pm 12.88	0.88	0.72	0.67	0.10	0.17	0.25	5.31	10.23	13.23
	0.5	38.47 \pm 8.56	29.90 \pm 13.98	19.76 \pm 14.72	0.87	0.73	0.66	0.11	0.19	0.28	5.30	8.46	09.46
	0.75	37.62 \pm 9.23	25.72 \pm 13.71	18.32 \pm 14.22	0.84	0.69	0.62	0.13	0.22	0.31	5.25	8.57	08.60
	1.0	38.70 \pm 7.89	23.39 \pm 14.86	16.29 \pm 15.47	0.86	0.67	0.57	0.12	0.30	0.41	5.12	7.89	08.05
ART-IQN	-	39.85 \pm 7.23	36.43 \pm 13.51	24.88 \pm 12.31	0.87	0.77	0.70	0.11	0.17	0.15	5.32	8.26	11.76

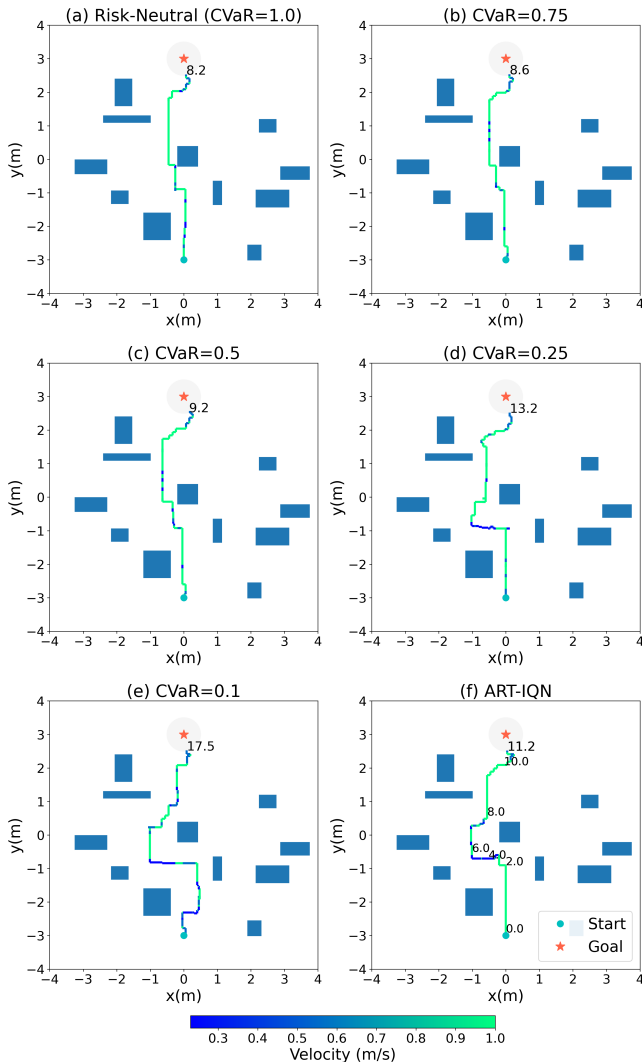


Fig. 4: Drone behavior and navigation time comparison between agents with different risk-tendencies.

a success rate of 0.77, maintains low collision rate and an average navigation time of only 0.37s longer compared to risk-neutral policy. For $n_{obs} = 12$, the average return and success rate of all the agents decrease. This drop can be explained by the severe partial observability the agents face. While there are more cases of timeout with lower CVaR values and more collisions with a higher one, ART-IQN performs best in success rate and collision rate with a decent navigation time.

2) *Qualitative Analysis*: We demonstrate the behavior of agents in Fig. 4 by considering a typical environment

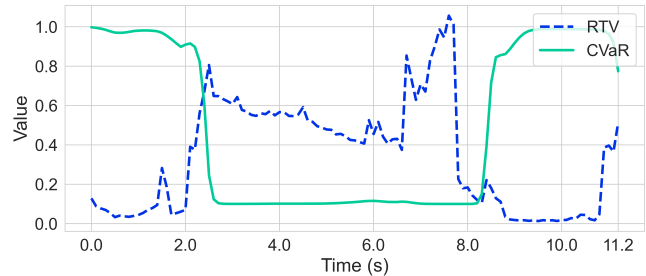


Fig. 5: RTV and adaptive CVaR. ART-IQN can adapt its CVaR value accordingly with RTV as an estimation of intrinsic uncertainty in the environment.

encountered in evaluation. In Fig. 4 (a), risk-neutral IQN achieves the goal as fast as possible, ignoring the risk of getting too close to obstacles that leads to a higher rate of collision. On the other hand, risk-averse policies, especially the one with $\alpha = 0.1$ as shown in Fig. 4 (e), generate safer policies but sacrifice in navigation efficiency. In contrast, as shown in Fig. 4 (f), ART-IQN acts adaptively - avoiding obstacles cautiously in the middle area where there is more uncertainty encountered, and flying at a higher speed when it's more certain about the current observation.

Fig. 5 shows the RTV and adaptive CVaR along the trajectory in Fig. 4 (f). The drone starts with CVaR $\alpha = 1.0$, which is risk-neutral when it does not get to know the environment. During 0s to 2s, the drone flies at its maximum speed risk-neutrally as the current uncertainty is low. Around 2s to 8s, the intrinsic uncertainty estimated by RTV increases and stays at a high level, resulting a decrease in CVaR values, which corresponds to the risk-averse behavior in the middle area of Fig. 4 (f). When the RTV drops and maintains a low level from 8s until the episode ends, the CVaR increases, achieving a risk-neutral policy to reach the goal point efficiently. It is clear that, as the drone navigates through the environment, it adjusts its risk-tendency to be risk-averse when the uncertainty increases and to be risk-neutral when there is less uncertainty in the environment.

B. Real-World Experiments

1) *Hardware Setup*: The Crazyflie nano drone we used for real-world experiments is shown in Fig. 2. It has dimensions $92 \times 92 \times 29mm$ and weighs 27.5g. Policy is performed on a laptop, which gives velocity command and communicates with Crazyflie via a radio-to-USB dongle, and then followed by an onboard closed-loop attitude controller on the Crazyflie. The velocity sample period T is set to be the same as in simulation as shown in Table

TABLE III: REAL-WORLD EXPERIMENT RESULTS

Environment	CVaR	# Success / # Collision			Navigation time (mean \pm std) [s]		
		1	2	3	1	2	3
IQN	0.1	9 / 0	3 / 0	3 / 0	14.41 \pm 3.67	16.73 \pm 0.49	21.52 \pm 1.76
	1.0	8 / 1	2 / 1	0 / 3	10.12 \pm 3.29	12.54 \pm 0.58	-
ART-IQN	-	9 / 0	3 / 0	3 / 0	12.32 \pm 3.46	13.37 \pm 0.47	17.95 \pm 1.85

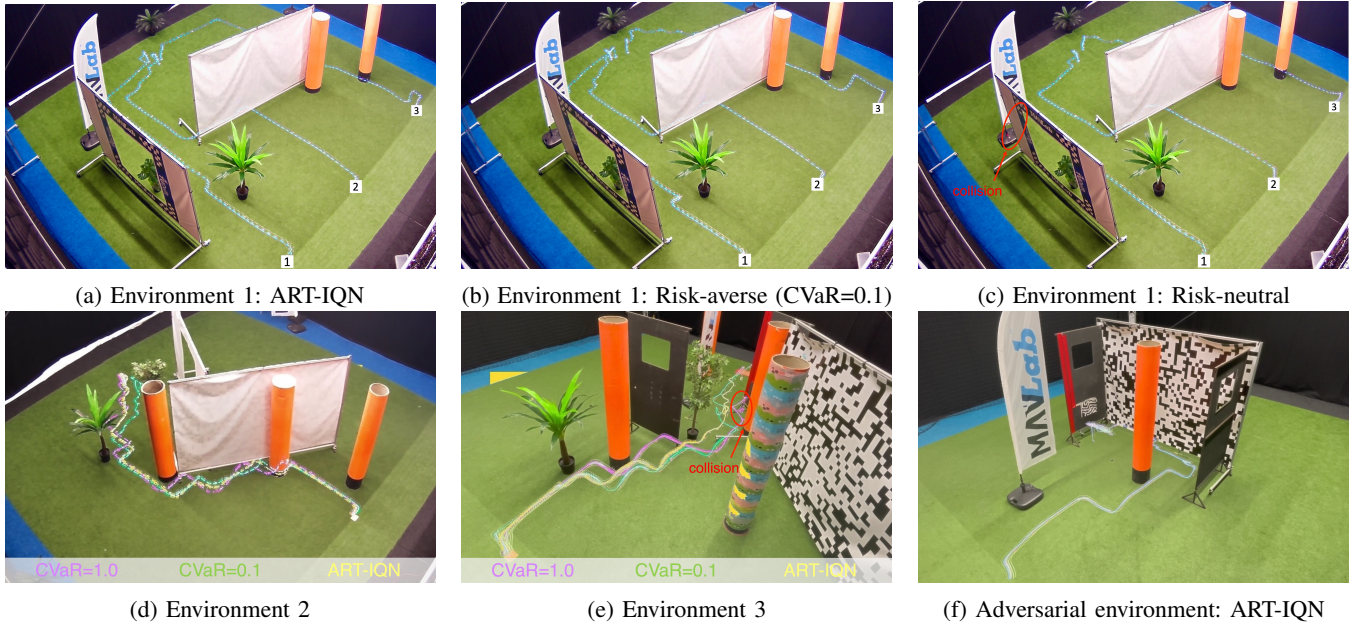


Fig. 6: Image frames for real-world experiments. The trajectory is plotted by the blue LED on the drone. For environment 2 and 3, the trajectories generated by different agents are distinguished with recolored LEDs.

I. The drone navigation task is implemented in 3 different $10 \times 10m$ cluttered environments. As shown in Fig. 6, we put artificial trees, boards and cylinders as obstacles in the environment that the agent has not seen in simulation to testify the generalization ability of our method. Reflective markers are attached on four propeller hubs to let the motion capturing system track the global position of the drone.

2) *Evaluation Results:* We test IQN with $\alpha = \{0.1, 1.0\}$ and ART-IQN for comparison. In environment 1, each agent is initialized at 3 different starting points for 3 runs, resulted a total number of 27 runs. In environment 2 and a more complex environment 3, the drone takes off at the same starting point. As shown in Table III, all runs succeeded except risk-neutral policies. Although risk-neutral IQN achieves fastest navigation in succeeded runs, it ignores the risk in the environment causing the drone to collide with obstacles. Risk-averse IQN succeeds in all experiments without collisions, but with a loss of navigation efficiency. In contrast, ART-IQN navigates through all the environments safely and efficiently.

Fig. 6 (a)-(e) demonstrate diverse behaviors among different risk-tendencies. Unlike risk-neutral IQN, both ART-IQN and risk-averse IQN keep a safe distance to obstacles to avoid collisions. The advantage of ART-IQN compared to risk-averse IQN is mainly reflected in the shorter navigation time as in Table III. In addition, we designed a more difficult environment to study if IQN agents could generalize to adversarial situations. However, all agents including ART-

IQN stick around in the corner of the U-shape obstacle as shown in Fig. 6 (f). It is highly possible that U-shape obstacles have not been seen often by the agent and can be solved by generating similar situations during training.

V. CONCLUSION

In conclusion, focusing on the autonomous drone navigation under partial observability, we propose an adaptive risk-tendency algorithm based on distributional RL to adapt risk-tendency accordingly with the estimated intrinsic uncertainty. Our algorithm uses EWF to adjust risk-tendency represented by the CVaR, with lower tail conditional variance as an estimation of the intrinsic uncertainty. We show the effectiveness of our algorithm both in simulation and real-world experiments. Empirical results show that our algorithm can adaptively balance the efficiency-safety trade-off.

However, the step size η to update CVaR is currently pre-specified, it would be worthy to optimize it. Furthermore, the navigation task can be performed more fluidly by combining adaptive risk-tendency with other RL methods like DSAC [12] to deal with continuous action spaces. Despite that, ART-IQN could serve as a first step to develop risk-tendency adaptation methodologies for distributional RL applications especially in risk-sensitive settings.

ACKNOWLEDGMENT

The authors would like to thank Jinke He for the discussions and Bart Duisterhof, Yingfu Xu for the real-world experiment setup.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. The MIT Press, 2018. [Online]. Available: <http://incompleteideas.net/book/the-book-2nd.html>
- [2] C. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3/4, pp. 279–292, 1992.
- [3] V. e. a. Mnih, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015. [Online]. Available: <http://dx.doi.org/10.1038/nature14236>
- [4] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *4th International Conference on Learning Representations, ICLR, 2016*. [Online]. Available: <http://dblp.uni-trier.de/db/conf/iclr/iclr2016.html#LillicrapPHHETS15>
- [5] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *J. Mach. Learn. Res.*, vol. 17, pp. 39:1–39:40, 2016. [Online]. Available: <http://jmlr.org/papers/v17/15-522.html>
- [6] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," *CoRR*, vol. abs/1710.11040, 2017. [Online]. Available: <http://arxiv.org/abs/1710.11040>
- [7] Y. C. Tang, J. Zhang, and R. Salakhutdinov, "Worst cases policy gradients," in *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, ser. Proceedings of Machine Learning Research, vol. 100. PMLR, 2019, pp. 1078–1093. [Online]. Available: <http://proceedings.mlr.press/v100/tang20a.html>
- [8] M. G. Bellemare, W. Dabney, and R. Munos, "A distributional perspective on reinforcement learning," in *ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 449–458. [Online]. Available: <http://proceedings.mlr.press/v70/bellemare17a.html>
- [9] W. Dabney, M. Rowland, M. G. Bellemare, and R. Munos, "Distributional reinforcement learning with quantile regression," *CoRR*, vol. abs/1710.10044, 2017. [Online]. Available: <http://arxiv.org/abs/1710.10044>
- [10] W. Dabney, G. Ostrovski, D. Silver, and R. Munos, "Implicit quantile networks for distributional reinforcement learning," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1104–1113. [Online]. Available: <http://proceedings.mlr.press/v80/dabney18a.html>
- [11] N. A. Urpi, S. Curi, and A. Krause, "Risk-averse offline reinforcement learning," *CoRR*, vol. abs/2102.05371, 2021. [Online]. Available: <https://arxiv.org/abs/2102.05371>
- [12] X. Ma, Q. Zhang, L. Xia, Z. Zhou, J. Yang, and Q. Zhao, "Distributional soft actor critic for risk sensitive learning," *CoRR*, vol. abs/2004.14547, 2020. [Online]. Available: <https://arxiv.org/abs/2004.14547>
- [13] bitcraze, *Crazyflie 2.1 Specifications*. [Online]. Available: <https://www.bitcraze.io/products/crazyflie-2-1/>
- [14] D. Kamran, T. Engelgeh, M. Busch, J. Fischer, and C. Stiller, "Minimizing safety interference for safe and comfortable automated driving with distributional reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic, September 27 - Oct. 1, 2021*. IEEE, 2021, pp. 1236–1243. [Online]. Available: <https://doi.org/10.1109/IROS51168.2021.9636847>
- [15] J. Choi, C. R. Dance, J. Kim, S. Hwang, and K. Park, "Risk-conditioned distributional soft actor-critic for risk-sensitive navigation," in *IEEE International Conference on Robotics and Automation, ICRA 2021, Xi'an, China, May 30 - June 5, 2021*. IEEE, 2021, pp. 8337–8344. [Online]. Available: <https://doi.org/10.1109/ICRA48506.2021.9560962>
- [16] W. R. Clements, B. Robaglia, B. V. Delft, R. B. Slaoui, and S. Toth, "Estimating risk and uncertainty in deep reinforcement learning," *CoRR*, vol. abs/1905.09638, 2019. [Online]. Available: <http://arxiv.org/abs/1905.09638>
- [17] V. Francois-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," 2018, cite arxiv:1811.12560. [Online]. Available: <http://arxiv.org/abs/1811.12560>
- [18] L. Tai, G. Paolo, and M. Liu, "Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 31–36.
- [19] G. Kahn, A. Villafior, V. Pong, P. Abbeel, and S. Levine, "Uncertainty-aware reinforcement learning for collision avoidance," *CoRR*, vol. abs/1702.01182, 2017. [Online]. Available: <http://arxiv.org/abs/1702.01182>
- [20] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *ICML 2016, New York City, NY, USA, June 19-24, 2016*, ser. JMLR Workshop and Conference Proceedings, vol. 48. JMLR.org, 2016, pp. 1050–1059. [Online]. Available: <http://proceedings.mlr.press/v48/gal16.html>
- [21] I. Osband, C. Blundell, A. Pritzel, and B. V. Roy, "Deep exploration via bootstrapped DQN," in *Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, 2016*, pp. 4026–4034. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/hash/8d8818c8e140c64c743113f563cf750f-Abstract.html>
- [22] B. Lütjens, M. Everett, and J. P. How, "Safe reinforcement learning with model uncertainty estimates," in *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*. IEEE, 2019, pp. 8662–8668. [Online]. Available: <https://doi.org/10.1109/ICRA.2019.8793611>
- [23] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [24] T. Fan, P. Long, W. Liu, J. Pan, R. Yang, and D. Manocha, "Learning resilient behaviors for navigation under uncertainty," in *2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020*. IEEE, 2020, pp. 5299–5305. [Online]. Available: <https://doi.org/10.1109/ICRA40945.2020.9196785>
- [25] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," *CoRR*, vol. abs/1409.1259, 2014. [Online]. Available: <http://arxiv.org/abs/1409.1259>
- [26] E. F. Camacho and C. B. Alba, *Model predictive control*. Springer science & business media, 2013.
- [27] W. Dabney, Z. Kurth-Nelson, N. Uchida, C. K. Starkweather, D. Hassabis, R. Munos, and M. Botvinick, "A distributional code for value in dopamine-based reinforcement learning," *Nat.*, vol. 577, no. 7792, pp. 671–675, 2020. [Online]. Available: <https://doi.org/10.1038/s41586-019-1924-6>
- [28] M. Buchinsky, "Recent advances in quantile regression models: A practical guideline for empirical research," *The Journal of Human Resources*, vol. 33, no. 1, pp. 88–126, 1998. [Online]. Available: <http://www.jstor.org/stable/146316>
- [29] M. T. Spaan, "Partially observable markov decision processes," in *Reinforcement Learning*. Springer, 2012, pp. 387–414.
- [30] C. Filippi, G. Guastaroba, and M. G. Speranza, "Conditional value-at-risk beyond finance: a survey," *Int. Trans. Oper. Res.*, vol. 27, no. 3, pp. 1277–1319, 2020. [Online]. Available: <https://doi.org/10.1111/itor.12726>
- [31] T. Moskovitz, J. Parker-Holder, A. Pacchiano, and M. Arbel, "Deep reinforcement learning with dynamic optimism," *CoRR*, vol. abs/2102.03765, 2021. [Online]. Available: <https://arxiv.org/abs/2102.03765>
- [32] B. Mavrin, H. Yao, L. Kong, K. Wu, and Y. Yu, "Distributional reinforcement learning for efficient exploration," in *ICML 2019, 9-15 June 2019, Long Beach, California, USA*, vol. 97. PMLR, 2019, pp. 4424–4434. [Online]. Available: <http://proceedings.mlr.press/v97/mavrin19a.html>
- [33] P. Huber, J. Wiley, and W. InterScience, *Robust statistics*. Wiley New York, 1981.
- [34] OpenAI, *OpenAI gym environments*. [Online]. Available: https://gym.openai.com/envs/#classic_control
- [35] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 23–30.
- [36] A. F. Agarap, "Deep learning using rectified linear units (relu)," *ArXiv*, vol. abs/1803.08375, 2018.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>