

Continuous Drone Control using Deep Reinforcement Learning for Frontal View Person Shooting

Nikolaos Passalis · Anastasios Tefas

Abstract Drones, also known as Unmanned Aerial Vehicles (UAVs), can be used to aid various aerial cinematography tasks. However, using drones for aerial cinematography requires the coordination of several people increasing the cost and reducing the shooting flexibility, while also increasing the cognitive load of the drone operators. To overcome these limitations we propose a deep reinforcement learning (RL) method for continuous fine-grained drone control, that allows for acquiring high-quality frontal view person shots, is proposed. To this end, a head pose image dataset is combined with 3D models and face alignment/warping techniques to develop an RL environment that realistically simulates the effects of the drone control commands. An appropriate reward-shaping approach is also proposed to improve the stability of the employed continuous RL method. Apart from performing continuous control, it was demonstrated that the proposed method can be also effectively combined with simulation environments that support only discrete control commands, improving the control accuracy, even in this case. The effectiveness of the proposed technique is experimentally demonstrated using several quantitative and qualitative experiments.

Keywords Deep Reinforcement Learning · Continuous Control · Aerial Cinematography · Drone Control

1 Introduction

The wide availability of drones, which are also known as Unmanned Aerial Vehicles (UAVs), led to the development of versatile autonomous systems that can be used to aid several challenging tasks, ranging from promptly responding to medical emergencies [7], to detecting fires in forests [21], and protecting the wildlife [15]. The high versatility of drones and their ability to capture spectacular aerial shots make them especially suitable for performing aerial cinematography tasks [27]. However, flying a drone in this setting requires the coordination of several people. A pilot has to control each drone, while a camera operator has to control the shooting camera on each drone. At the same time, the director has to coordinate several pilots and camera operators, if multiple drones are used, and ensure the quality of the captured shots. This situation increases the cost of using drones for aerial cinematography and severely limits the shooting flexibility by putting a significant cognitive load on the director and drone/camera operators.

N. Passalis and A. Tefas
Department of Informatics
Aristotle University of Thessaloniki
Thessaloniki 541 24, Greece Tel.: +30-231-099-1932
E-mail: {passalis, tefas}@csd.auth.gr

The aforementioned limitations led to the development of various techniques for assisting drone-based cinematography, ranging from techniques for automated planning for multi-view drone shooting [27], and crowd avoidance for complying with drone legislation [43], to autonomous drone control systems [11], and various techniques for steering drone video shooting [32]. These techniques automate various parts of the shooting process, reducing the cost of aerial cinematography and the cognitive load of human operators. At the same time, the computational complexity of the developed methods must be taken into account, since drones usually have limited computational resources (processing power and memory), leading to the development of various approaches to lower the computational requirements of the aforementioned methods [29, 48].

Currently, we are still far from developing fully autonomous drones that would be able to automatically shot professional-grade footage according to a predefined plan, as designed by the director, while detecting salient events for performing opportunist shooting. Instead, this paper focuses on automating part of the shooting process. More specifically, we aim to appropriately control the height and the relative position of the drone with respect to the person of interest to acquire a clear frontal shot. To this end, a framework for developing deep reinforcement learning (RL) methods, that are able to learn accurate control policies for the *continuous fine-grained* control of a drone from raw pixel inputs, is presented. Note that even though deep RL has been applied for solving several challenging and delicate control tasks [8, 14, 47], that would usually require the difficult and laborious fine-tuning of traditional control methods [3], little work has been done for developing deep RL techniques for cinematography-oriented drone control.

The contributions of this paper are briefly summarized below. First, a realistic simulation environment is developed using the Head Pose Image Database (HPID) [13]. However, the HPID only contains a discrete number of poses for each person, supporting only discrete control commands. To overcome this limitation and effectively simulate the effect of continuous control commands, the simulation environment is combined with 3D models, along with the face alignment and warping technique proposed in [50]. This allows for training RL methods that will perform continuous fine-grained control. Then, an appropriate reward-shaping approach is proposed to improve the stability of the employed continuous RL method. Apart from performing continuous control, it is demonstrated that the method introduced in this paper can be also effectively combined with simulation environments that support only discrete control commands, improving the control accuracy. Finally, the proposed approach is compared both to an RL method that performs discrete control, as well as to a traditional controller that directly uses the output of a deep model that performs pose estimation. It is experimentally demonstrated that the proposed approach improves the control accuracy over these methods.

This paper is an extended version of our previous work [33], where a method capable of performing drone control in discrete steps for frontal view shooting was proposed. This paper significantly extends our previous work by proposing a method capable of performing continuous fine-grained drone control using deep RL. The proposed method is called “Continuous Drone Control” (CDC) through the rest of this paper. First, the limitations of the simulation environment that was used in [33] and in [49] are lifted by using an advanced 3D modeling method that can simulate the effect of continuous control commands [50]. This allows for developing novel fine-grained continuous control methods for frontal view shooting using deep RL. Second, the proposed approach is appropriately extended and evaluated using a policy gradient method, instead of Q-learning, allowing for directly performing continuous control without the need of discretizing the action space. Finally, CDC can be also readily used in simulation environments that only support discrete action spaces. Indeed, it is experimentally demonstrated that CDC can improve the control accuracy over control methods that perform discrete control, even when environments that only support discrete control are used. To further boost RL research, that critically relies on the availability of simulation environments, we provide an open-source implementation of the developed simulation environments as well as of the developed techniques at https://github.com/passalis/continuous_drone_frontal_rl.

The rest of the paper is structured as follows. The related work is briefly discussed and compared to the proposed approach in Section 2. Then, CDC is introduced and described in detail in Section 3. The experimental evaluation is provided in Section 4. Finally, future work is discussed and conclusions are drawn in Section 5.

2 Related Work

Several approaches can be used to appropriately control a drone/camera according to a specific cinematography-oriented objective. Perhaps the most widely used approach is to project a set of known 2D landmark points, e.g., nose, eyes, mouth, etc., into the 3D space and solve the corresponding Perspective-n-Point (PnP) problem to estimate the pose of the head of the person of interest [19,30]. Then, a PID controller can be used to appropriately control the drone to acquire the desired frontal shot [1,34]. However, this approach requires accurately detecting several 2D facial landmark points, which can be especially difficult if a low-resolution image input is used. Note that this is usually the case, since drones frequently use a low-resolution video feed for performing the on-board processing tasks, even when high resolution cameras are available, due to their limited processing power and memory. Apart from these, careful calibration of the system is required and it is non-trivial to extend this approach for estimating the pose of other objects, since the landmark points have to be appropriately redefined and a corresponding detector must be developed.

Some of the aforementioned drawbacks can be addressed by using deep pose estimation algorithms that are directly trained to estimation the pose of various objects using a training set of data that contains images along with their pose annotations [12,36,41]. Even though such deep learning approaches have been shown to increase the pose estimation accuracy and lead to more robust pose estimators, there is no guarantee that they will be *optimal* for performing control tasks, while, on the other hand, developing optimal control algorithms has a long history in various engineering fields [5,20]. Quite recently, it was shown that combining the great learning capacity of deep learning models with reinforcement learning (RL) techniques can indeed lead to the development of robust control policies, that can work under stochastic and noisy environments, producing spectacular results, that often outperform humans on sophisticated control tasks [24,26,44]. This also led to the development several recent deep RL approaches for various robotics applications [9,10,22,33,38]. However, using a deep RL approach for learning accurate control policies from raw pixel inputs is not straightforward, requiring the development of the appropriate simulation environments along with the careful encoding of the problem objective into a reward function. Apart from that, *reward shaping* can be also required to ensure that the method will indeed perform as it was intended and ensure the stable convergence of the learning process [28]. To the best of our knowledge, this is the first work where deep RL is used to perform accurate continuous fine-grained drone control from raw pixel input for frontal view shooting using a dedicated simulation environment that employs an advanced 3D modeling method that can simulate the effect of continuous control commands.

3 Proposed Method

The proposed CDC method is presented in this Section. First, a brief introduction to RL is provided along with the used notation. Then, the developed simulation environments are presented. Finally, the proposed reward shaping scheme, as well as the complete methodology for developing a deep RL agent for continuous drone control are derived and discussed in detail. Note that the term *agent* in this paper refers to parameterized software agents that learn to individually perform various actions in order to maximize the reward obtained from the environment and should not be confused with multiagent systems [51].

3.1 Deep Reinforcement Learning

Given an environment for which the *Markov property* is satisfied (the future state depends only on the current state and the selected action), RL can be modeled using Markov Decision Processes (MDPs). A MDP can be defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}(\cdot), \mathcal{R}(\cdot), \gamma)$, where:

1. \mathcal{S} is the set of possible states for the environment.
2. $\mathcal{A} = \{a_1, a_2, \dots, a_{N_a}\}$ is the set of possible actions, where N_a is the number of possible actions. Note that the action space can be continuous: $\mathcal{A} \subseteq \mathbb{R}^m$, where m is the number of controls.
3. $\mathcal{P}(s_{t+1}|s_t, a_t)$ is the probability that the environment will transit from the state s_t to the state s_{t+1} given that the action a_t has been performed.
4. $\mathcal{R}(s_t, a_t, s_{t+1}) = r_{t+1}$ is the *immediate* reward that the agent receives when performing the action a_{t+1} and transitioning from state s_t to state s_{t+1} .
5. γ is a discount factor that defines the importance of immediate rewards versus future rewards, where typically $0 \leq \gamma < 1$. For $\gamma = 0$ the agent is *short-sighted*, while higher values increase the importance of future rewards.

An agent starts at state s_0 and selects the next action according to a *policy* $\pi(s)$ that defines the action that the agent will perform. Note that the policy $\pi(\cdot)$ can be stochastic and it is usually described as a probability distribution over the available actions: $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$. The agent's behavior defines a sequence of state-action-rewards $s_0, a_1, r_1, s_1, a_2, r_2, s_2, a_3, r_3, s_3, \dots, s_T$ that describes the history of the agent for a given episode that ends after T steps. The discounted accumulated return for an episode consisting of T steps is defined as:

$$R = \sum_{t=1}^T \gamma^{t-1} r_t, \quad (1)$$

where each action is selected according to the policy $\pi(s)$. RL aims to learn the optimal policy π^* that provides the maximum expected return:

$$\pi^* = \arg \max_{\pi} \mathbb{E}[R|\pi]. \quad (2)$$

Several approaches have been proposed for tackling this problem [37]. For example, Q-learning aims to learn the optimal action-value function $Q^*(s, a)$, which is used to express the expected reward of performing the action a from the state s , given that an optimal policy is then followed. Then, the optimal policy can be directly derived as:

$$\pi^*(s) = \arg \max_a Q^*(s, a). \quad (3)$$

Even though simple look-up tables can be used to represent and store the optimal values (Q-values), this approach quickly becomes impractical as the size of the state set \mathcal{S} increases. To overcome this limitation, deep neural networks can be used to approximate the action-value function (deep Q-learning). The neural network is updated after each time-step using stochastic gradient descent [17], to minimize an appropriately defined loss function $\mathcal{L}(\cdot)$ that measures the error between the current estimation $Q(s_t, a_t, \mathbf{W}_t)$ and the updated estimation, where \mathbf{W}_t denotes the parameters of the neural network after t optimization steps. The difference between the current estimation and the updated estimation is defined as:

$$\delta = Q(s_t, a_t, \mathbf{W}_t) - (r + \gamma \max_a Q(s_{t+1}, a, \mathbf{W}_t)), \quad (4)$$

and can be minimized using any appropriate loss functions, e.g., the squared loss $\mathcal{L}(\delta) = 0.5\delta^2$.

Even though deep Q-learning can provide state-of-the-art solutions to several RL problems it is not easy to directly adapt it for solving *continuous* control problems, while it does not directly optimize the policy of the agent (instead it indirectly learns the optimal policy by estimating the

Q-values). Such continuous control problems can be more efficiently solved using a different kind of RL algorithms, the policy gradients methods [24,25], that directly learn the policy of the agent. A special category of these methods work by using two different estimators: the first one, called *actor*, is directly trained to estimate the optimal policy, while the second one, which is called *critic*, is used to estimate the *advantage* that a state/selected action provides and then used to appropriately train the actor. Again, deep neural networks can be used to implement the actor and critic models. It is worth noting that usually the critic model is trained similarly to the Q-learning approach by observing the rewards/punishment the agent obtains. Next, the actor is appropriately updated to increase the probability of selecting the actions that lead to maximizing the expected reward. A recently proposed deep RL method that can be used to train agents to performs continuous control, the Deep Deterministic Policy Gradient method [24], is employed for the experiments conducted in this paper. The interested reader is also referred to [2,23,37], for a more in-depth review of RL.

3.2 Proposed Drone Simulation Environment

The Head Pose Image Database [13], abbreviated as “HPID” thereof, was used to develop the simulation environment. HPID contains 2,790 face images of 15 subjects in various poses taken in a constrained environment. More specifically, for each person face images with various head tilts (vertical angle) and pans (horizontal angle) were taken. For the tilt, head photos at -90° , -60° , -30° , -15° , 0° , 15° , 30° , 60° , and 90° were taken, while for the pan images that depict the head at -90° , -75° , -60° , -45° , -30° , -15° , 0° , 15° , 30° , 45° , 60° , 75° and 90° were used. The full pan range exists only for the images with tilt angles between -60° and 60° . Therefore, we restrict the available tilt angles used in the simulator to the aforementioned range.

The developed environment for discrete control supports 5 different actions (assuming that the drone moves on the sphere):

1. *stay*: do not perform any action (to be used when a clear frontal view has been obtained),
2. *left*: move the drone left by $15^\circ \rightarrow$ pan decreases by 15° ,
3. *right*: move the drone right by $15^\circ \rightarrow$ pan increases by 15° ,
4. *up*: move the drone upwards by $15^\circ/30^\circ$ (depending on the available annotations) \rightarrow tilt decreases by $15^\circ/30^\circ$, and
5. *down*: move the drone downwards by $15^\circ/30^\circ$ (depending on the available annotations) \rightarrow tilt increases by $15^\circ/30^\circ$.

During these movements, we assume that the camera is appropriately controlled to keep the face centered, e.g., using a PID controller [34]. This is a quite realistic assumption, since deep face detection algorithms are able to successfully detect faces with high accuracy, while even simple PID-based control approaches work quite well for this task [34]. If an agent requests a control command that exceeds the limits of the simulator, e.g., an angle larger than 90° , then the simulator remains at its last valid state. The developed simulator uses these images to simulate the movement of a drone in a part of a sphere defined by the center of the head of the subject. Some example face images of one person are shown in Fig. 1. This environment can be directly used to simulate the movement of a drone in discrete steps: 15° steps for the pan and in $15^\circ/30^\circ$ steps for the tilt.

Furthermore, this environment can be also used in *pseudo-continuous* mode by translating the continuous control commands into multiple discrete steps. In this work, we assume that a maximum of three consecutive discrete steps can be performed in one continuous control step. Therefore, the continuous control output of the agent is quantized into 4 regions: no action, one control step, two control steps (the same action is repeated twice), and three control steps (the same action is repeated three times). This allows to significantly reduce the number of controls that the agent has to handle, since instead of having an agent that can perform 49 (7 actions per control axis) different combinations of actions (when the tilt and pan are simultaneously controlled), only two continuous control outputs are needed (one for the tilt and one for the pan).

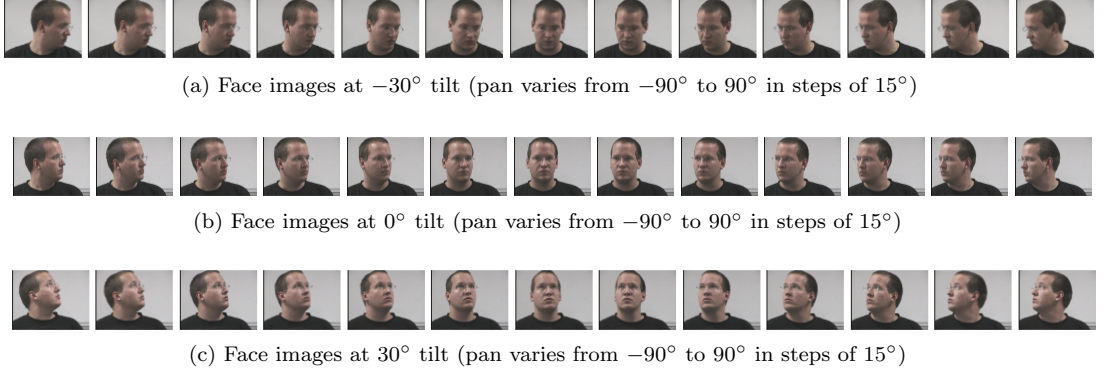


Fig. 1: Discrete Simulation environment: The drone moves in a sphere (approximately) centered at the face of the subject. The control commands (left/right/up/down) are simulated using the appropriate images from HPID.




Fig. 2: Continuous Simulation Environment: Some of the 3D generated images for pan that ranges from -30° to 90° using steps of 3° (note that the simulator resolution is 1°). Tilt was set to 0° .

However, the approach described above cannot be used to perform fine-grained continuous control, since it is still limited by the available face poses. To overcome this limitation and provide a more realistic simulation environment, the face alignment and warping technique proposed in [50], was used to obtain 3D generated images of each person for various pans/tilts. To this end, we directly used an open-source implementation of this technique, provided by the authors of [50], available at <https://github.com/junyanz/FaceDemo>. One frontal image from each person, which was randomly selected from the two frontal images provided by the HPID dataset, was used to generate the texture for the 3D model. In the developed continuous control environment the pan ranges from -45° to 45° in 1° steps (this can be assumed to be the control accuracy), while the tilt ranges from -20° to 20° in 1° steps. One face at various pans, as provided by the developed simulation environment for continuous control, is depicted in Fig. 2. Note that only the cropped face is fed to the RL agent. This can be easily ensured in a real application using face detection methods [42, 45, 46], and then appropriately cropping the part of the image that contains the face. It is worth noting that an agent trained on the discrete/pseudo-continuous environment can be used to perform coarse-grained control and approximately center the face image, while the agent trained on the fine-grained environment can be then employed to fine-tune the drone position accordingly and make small corrections. All the developed environments (discrete, pseudo-continuous and continuous) are *OpenAI Gym*-compatible [4] and provided at https://github.com/passalis/continuous_drone_frontal_rl.

3.3 Continuous Drone Control with Deep Reinforcement Learning

The complete pipeline of the proposed deep RL technique for drone control is described in this subsection. The RL agent interacts with the developed environment and *observes* its state (acquired shot) as shown in Fig. 1 and 2. To simplify the learning process, a face detector can be employed to detect and crop the face image [42] (or face annotations can be used if they are available). After appropriately cropping the image, it was resized to 64×64 pixels. Therefore, the RL agent at the t -th time step observes a tensor $\mathbf{x}_t \in \mathbb{R}^{64 \times 64 \times 3}$ that corresponds to the cropped face image.

Defining a meaningful reward function is critical for the fast and stable convergence of RL algorithms. Even though RL can deal, to some extent, with sparse and time-delayed rewards, we experimentally found out that rewarding (or punishing) the agent after each action can significantly speed up the learning process. The reward of the agent is related to the current control error. The control error at the t -th step is defined as:

$$e_t = \frac{1}{2} \left(\left(\frac{x_t}{x_t^m} \right)^2 + \left(\frac{x_p}{x_p^m} \right)^2 \right), \quad (5)$$

where x_t is the current tilt (in degrees) and x_p is the current pan (in degrees), while x_t^m and x_p^m is the absolute maximum tilt/pan that can be observed. Note that it is trivial to modify this error function to acquire non-frontal shots around the desired tilt/pan range. Then, the reward function is defined as:

$$r_t^{(raw)} = 1 - e_t. \quad (6)$$

However, this reward function provides a small positive rewards even for the wrong control movements. Therefore, a threshold e_{thres} can be used for rewarding the agent, leading to a thresholded reward function:

$$r_t^{(thres)} = \begin{cases} 0, & \text{if } e_t > e_{thres} \\ 1 - e_t/e_{thres}, & \text{otherwise} \end{cases}, \quad (7)$$

If e_{thres} is set to 1, then the agent is rewarded at every time-step. Again, note that even though the reward is proportional to the control error e_t , this can slow down the learning process. In the conducted experiments the agent was rewarded only when it was near to acquiring a correct frontal shot, i.e., e_{thres} was set to 0.2. Then, to further boost the learning process, an extra small reward a_r or punishment a_p can be provided whenever the agent makes a correct or wrong movement:

$$r_t^{(bonus)} = \begin{cases} -a_p, & \text{if } e_t > e_{t-1} \\ a_r, & \text{if } e_t < e_{t-1} \\ 0, & \text{otherwise} \end{cases}. \quad (8)$$

Using a higher penalty for wrong actions ensures the stability of the control process and discourages control oscillations. The optimal values for these parameters can be chosen by performing cross-validation. In the conducted experiments it was established (by performing a simple line search) that the best control accuracy is obtained for $a_p = 0.4$ and $a_r = 0.2$. Finally, the reward function is defined as the sum between the raw reward $r_t^{(thres)}$, which depends on the control error e_t , and the “bonus” reward $r_t^{(bonus)}$, that encourages correct and stable control actions:

$$r_t^{(shaped)} = r_t^{(thres)} + r_t^{(bonus)}. \quad (9)$$

Three different agents were trained in this work: a) an agent capable of performing discrete control (baseline), b) an agent capable of performing pseudo-continuous control, and c) an agent capable of performing continuous control. The Double Q-learning approach was employed for training the first agent (discrete control) [16], while the Deep Deterministic Policy Gradient (DDPG)

Table 1: Neural network architecture. The size of the final output layer depends on the RL method that is used.

Layer Type	Activation Function	Output Shape
Input	-	$64 \times 64 \times 3$
Convolutional (5×5 , stride 2)	-	$30 \times 30 \times 16$
Batch Normalization	relu	$30 \times 30 \times 16$
Max Pooling (2×2)	-	$15 \times 15 \times 16$
Convolutional (3×3)	relu	$13 \times 13 \times 32$
Batch Normalization	-	$13 \times 13 \times 32$
Max Pooling (2×2)	-	$6 \times 6 \times 32$
Convolutional (3×3)	relu	$4 \times 4 \times 64$
Batch Normalization	-	$4 \times 4 \times 64$
Max Pooling (2×2)	-	$2 \times 2 \times 64$
Dense	relu	128
Dense	relu	64
Dense	-/tanh/-	1/2/5

method was used for training the other two agents [24]. The neural networks used for approximating the action-value function in Q-learning and the networks used for implementing the actor and critic in the DDPG method use the same architecture, which is shown in Table 1. Only the output layer of the networks is appropriately modified to accommodate the needs of each method, i.e., 5 output neurons are employed for the network used in the Q learning method (each neuron corresponds to the Q-value of the corresponding action for the given state), 1 output neuron is used for the network that implements the critic, while 2 output neurons (equipped with a *tanh* activation function) are used for the actor network (the output of each of the neurons is used to directly control the pan and tilt). Batch normalization is used after each convolutional layer [18], while the *relu* activation function is used for all the convolutional and dense layers (except from the final one). Also, note that the critic receives the output of the actor in its penultimate fully connected layer and predicts the Q-value of the action selected by the actor for the given input state, while the actor and critic share the same layers up to the first layer after the last convolutional layer. The latter allows to share the knowledge encoded between the actor and critic, while speeding up the learning process.

Learning Hyper-Parameters The size of the experience replay buffer is set to $N_{replay} = 5,000$ samples when the discrete/pseudo-continuous environment is used and to $N_{replay} = 10,000$ samples for the continuous environment. We used the smaller replay buffer size that allowed the models to converge smoothly. Using larger sizes is not expected to significantly affect the learning process. Batches of 128 samples were used for the optimization, while the target model was updated with a rate of 0.001 after each iteration. The Huber loss function was employed for training the critic network, as well as the network used in Q-learning:

$$\mathcal{L}(\delta) = \begin{cases} 0.5\delta^2, & \text{if } \delta < \delta_{thres} \\ |\delta| - 0.5\delta_{thres}, & \text{otherwise} \end{cases}, \quad (10)$$

where δ is the temporal difference error defined in Eq. (4). Huber loss is used instead of squared loss, since it is more robust to outliers, providing smoother gradients and stabilizing the learning process (δ_{thres} was set to 1). This is especially important when the total episode reward can accumulate to relatively large values, as in the conducted experiments. For updating the weights of the network, the RMSProp optimizer was used [40]. The learning rate was set to $\eta = 0.0001$ for the Double Q-learning method, while the learning rate for the actor was set to $\eta_a = 0.00001$ and to $\eta_c = 0.01$ for the critic. Setting the learning rates to the appropriate values were of crucial importance in order to ensure the stable convergence of the training process. To explore the solution space a

linear exploration policy was used. For the Q-learning method exploration starts with an initial rate of 1 and linearly decreases to 0.1 during the first 95,000 training steps. For the DDPG method a Gaussian process with mean 0 was used. The standard deviation of the process was linearly decreased from 0.5 (0.2 for the continuous environment) to 0.01 during the initial 95,000 steps (45,000 steps for the continuous environment). The agents were trained for 100,000 training steps for the discrete/pseudo-continuous environments and for 50,000 training steps for the continuous environment. The number of steps for each episode was set to 50. For each episode a random initial position was used for the drone. For training, head images from 10 persons of the HPID were used, while for the evaluation images from the rest 5 persons were used. The discount factor was set 0.95 for the discrete agents, but it was reduced to 0.5 for the continuous agents, since larger values led to stability issues during the training process. The keras-rl library was used to implement CDC [35].

4 Experimental Evaluation

The experimental evaluation is provided in this Section. For evaluating the method 500 random episodes were used, where for each episode the agent was allowed to perform 20 control actions. The discrete control method is abbreviated as “D-RL”, the pseudo-continuous drone control method as “P-CDC”, while the continuous drone control method as “CDC”. The learned agents were also compared to two other strategies: a) using a dummy agent that does not perform any control action (abbreviated as “Stay”) and b) using a deep CNN to perform pose regression and then appropriately control the drone (abbreviated as “Pose Regressor”). The Pose Regressor network use the same architecture as the CNN used for estimating the Q-values (Table 1) and it was trained to directly regress the tilt and pan of a face image using the same train/test setup, leading to a mean tilt error of 16.67° and a mean pan error of 13.71° (evaluation on the test set).

The evaluation results for the discrete and pseudo-continuous environments are reported in Table 2. Note even though these environments have different action spaces, they have identical states, since the same states are observed by the agents. Therefore, the results obtained with the D-RL and P-CDC are directly comparable, since the observations, rewards and control limits are the same. Both the simple (“raw”) and the appropriately shaped (“shaped”) reward functions are compared. Several conclusions can be drawn from the results reported in Table 2. First, using RL to train the agents leads to significantly better results than the “Stay” baseline agent, which corresponds to a dummy agent that does not perform any control action (the drone remains at its initial position). The deep pose regressor also leads to higher error, demonstrating the importance of learning optimal controllers for the task at hand instead of relying on hand-crafted methods to perform control. Furthermore, the proposed reward shaping approach always significantly improves the control accuracy, regardless the used RL method. The proposed pseudo-continuous control approach also leads to overall better results than the discrete control approach, reducing the mean control error from 6.23° to 5.12° (test evaluation). Finally, note that the agents are especially prone to overfitting (the P-CDC method learns to almost perfectly control the drone during the training).

These improvements are also demonstrated in the qualitative evaluation shown in Figures 3, 4, and 5. The first row corresponds to the D-RL method, while the second one to the P-CDC method. The P-CDC method manages to obtain a frontal view of the subjects significantly faster, since it is capable of adaptively and simultaneously controlling both axes (pan and tilt). For example, in Figure 3 the P-CDC method obtains a frontal view of the subject in just two control steps, while the D-RL method needs 7 control steps and it is more prone to oscillations around the center position. Similar behavior can be also observed in the other two figures, confirming the superior behavior of CDC for pseudo-continuous control.

The results for the fully continuous environment are shown in Table 3. Again, the shaped reward function leads to significantly better control accuracy than training with the raw reward function, highlighting the importance of using correctly designed reward functions. Note that the

Table 2: Drone control evaluation for frontal view person shooting using the discrete and pseudo-continuous environments. Both the simple (“raw”) and the appropriately shaped (“shaped”) reward functions are compared. The absolute mean tilt and pan error (to obtain a perfectly frontal shot) are reported.

Method	Eval. Type	Mean Absolute Tilt Error (°)	Mean Absolute Pan Error (°)
Stay	Train	29.94	48.24
Pose Regressor	Train	6.87	7.44
D-RL (raw)	Train	8.82	13.32
D-RL (shaped)	Train	1.47	2.58
P-CDC (raw)	Train	9.24	6.96
P-CDC (shaped)	Train	0.00	0.03
Stay	Test	29.37	48.93
Pose Regressor	Test	14.07	12.18
D-RL (raw)	Test	25.26	19.08
D-RL (hint)	Test	9.09	3.36
P-CDC (raw)	Test	12.81	5.88
P-CDC (hint)	Test	6.12	4.11



Fig. 3: Qualitative drone control evaluation: Comparing the D-RL approach (first row) to the P-CDC approach (second row). The state after each of the first 8 control actions are depicted. The test set was used for the evaluation. D-RL achieves 30°/45° tilt/pan error after two control steps (best achieved in 7 control steps), while P-CDC achieves 15°/15° tilt/pan error after two control steps (best achieved in 4 control steps).

Table 3: Drone control evaluation for frontal view person shooting using the continuous environment. Both the simple (“raw”) and the appropriately shaped (“shaped”) reward functions are compared. The absolute mean tilt and pan error (to obtain a perfectly frontal shot) are reported.

Method	Eval. Type	Mean Absolute Tilt Error (°)	Mean Absolute Pan Error (°)
Stay	Train	10.00	23.41
CDC (raw)	Train	2.62	6.86
CDC (shaped)	Train	0.18	0.72
Stay	Test	10.28	23.48
CDC (raw)	Test	3.25	5.75
CDC (shaped)	Test	0.51	2.00



Fig. 4: Qualitative drone control evaluation: Comparing the D-RL approach (first row) to the P-CDC approach (second row). The state after each of the first 8 control actions are depicted. The test set was used for the evaluation. D-RL achieves $0^\circ/75^\circ$ tilt/pan error after two control steps (best achieved in 7 control steps), while P-CDC achieves $0^\circ/0^\circ$ tilt/pan error after two control steps (best achieved in 2 control steps).



Fig. 5: Qualitative drone control evaluation: Comparing the D-RL approach (first row) to the P-CDC approach (second row). The state after each of the first 8 control actions are depicted. The test set was used for the evaluation. D-RL achieves $30^\circ/30^\circ$ tilt/pan error after two control steps (best achieved in 5 control steps), while P-CDC achieves $15^\circ/15^\circ$ tilt/pan error after two control steps (best achieved in 3 control steps).

agent learns to almost perfectly center the subjects using fine control commands, both during the training and the testing. This behavior is also confirmed in the test control sequences depicted in Fig. 6. The agent correctly controls the drone to obtain a clear frontal view, while usually small or no oscillations around the center are observed.

Learning Stability The rewards and mean Q-values during the training process are plotted in Fig. 7. Both the mean Q-value and reward steadily increase during the training process. However, note that the rewards fluctuate during the training process. The stochasticity of the training process greatly contributes to this behavior, since the environment is randomly initialized before starting each training episode. This behavior can be also attributed to the quite unstable nature of deep RL algorithms, as also discussed in [24, 26], which often require several heuristics, such as using a replay buffer, target network updates, etc., in order to ensure the smooth convergence of the training process. As a result, this observation led us to fully evaluate the learned agent every few training iterations (instead of just relying on the reward obtained on one specific episode). The evaluation results are shown in Fig. 8. Note that the behavior of the agent can vary significantly during the training process (especially when the DDPG method is used), even though a) the convergence



Fig. 6: Qualitative drone control evaluation using the continuous environment for six different test settings. The state after each of the first 11 control actions are depicted.

seems to be smooth and the evaluation was performed using (multiple episodes of) the training set (on which the agent was optimized). Therefore, to avoid using a sub-optimal model, the agent was evaluated every 5,000 training iterations and the best model (measured using the training set) was chosen as the final model. This process ensured that the best model that was found during the training process will be used. Indeed, it was experimentally confirmed that this process can increase the accuracy both for training and test evaluation, even though the best model was selected using the training set.

5 Conclusions

A method for developing deep RL agents for continuous fine-grained drone control to acquire high-quality frontal view person shots was presented in this paper. First, a realistic simulation environment was developed using the HPID. This environment was extended using 3D models, along with a face alignment and warping technique, to allow for simulating the effects of continuous control commands, overcoming the restrictions arising from the limited number of facial poses contained in this dataset. An appropriate reward-shaping approach was also proposed to improve the stability of the employed continuous RL method. Apart from performing continuous control, it was demonstrated that CDC can be also effectively combined with simulation environments that support only discrete control commands, improving the control accuracy, even in this case. Finally, the proposed approach was compared both to an RL agent that performs discrete control, as well as to a traditional controller that directly uses the output of a deep model that performs pose estimation. It was experimentally demonstrated that the proposed approach improves the control




Fig. 7: Convergence plots depicting the mean Q-value and the reward per training episode for the discrete control agent. The convergence of the D-RL agent using the shaped reward (first row) and the P-CDC agent using the shaped reward (second row) are depicted.

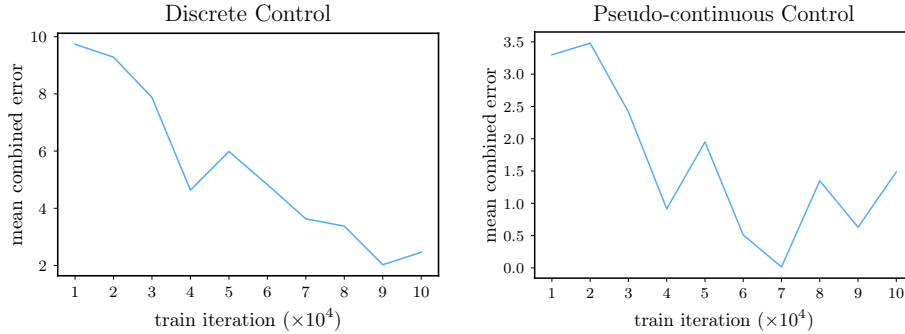


Fig. 8: Validation error during the training for the D-RL (shaped reward) and P-CDC (shaped reward).

accuracy over these methods, highlighting the importance of learning optimal controllers instead of relying on hand-crafted control techniques.

Several interesting future research directions exist. First, the impressive performance of the proposed technique paves the way for several other deep RL-based control methods for autonomous drone-based cinematography. However, it is worth noting that we are currently still far away from being able to directly deploy the developed methods in real world applications. In fact, even though

in our work we employed datasets with real images, instead of computer generated graphics, it is likely that much larger datasets should be created in order to ensure the successful deployment of these methods in real applications. Furthermore, CDC can be applied to more dynamic environments, e.g., moving targets can be used, and more sophisticated deep architectures, e.g., recurrent networks [6], can be used to accurately predict the movement of the target and control the drone. Furthermore, the use of end-to-end trainable systems, that simultaneously perform camera and drone control, can be developed and evaluated. Finally, the development of transfer learning techniques for RL [31, 39], that can combine the use of real datasets and simulators, which use computer-generated graphics, will allow for training RL techniques under a wider range of scenarios, while ensuring that they can be directly deployed in real-world applications.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflicts of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Acknowledgements This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE). This publication reflects the authors views only. The European Commission is not responsible for any use that may be made of the information it contains.

References

1. Ang, K.H., Chong, G., Li, Y.: Pid control system analysis, design, and technology. *IEEE Transactions on Control Systems Technology* **13**(4), 559–576 (2005)
2. Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A.: Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine* **34**(6), 26–38 (2017)
3. Åström, K.J., Hägglund, T., Astrom, K.J.: Advanced pid control (2006)
4. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym (2016)
5. Bryson, A.E.: Applied optimal control: optimization, estimation and control. CRC Press (1975)
6. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014)
7. Claesson, A., Fredman, D., Svensson, L., Ringh, M., Hollenberg, J., Nordberg, P., Rosenqvist, M., Djarv, T., Österberg, S., Lennartsson, J., et al.: Unmanned aerial vehicles (drones) in out-of-hospital-cardiac-arrest. *Scandinavian journal of trauma, resuscitation and emergency medicine* **24**(1), 124 (2016)
8. Duan, Y., Chen, X., Houthoofd, R., Schulman, J., Abbeel, P.: Benchmarking deep reinforcement learning for continuous control. In: *Proceedings of the International Conference on Machine Learning*, pp. 1329–1338 (2016)
9. Finn, C., Levine, S.: Deep visual foresight for planning robot motion. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2786–2793 (2017)
10. Finn, C., Yu, T., Fu, J., Abbeel, P., Levine, S.: Generalizing skills with semi-supervised reinforcement learning. *arXiv preprint arXiv:1612.00429* (2016)
11. Galvane, Q., Fleureau, J., Tariolle, F.L., Guillotel, P.: Automated cinematography with unmanned aerial vehicles. *arXiv preprint arXiv:1712.04353* (2017)
12. Garcia-Garcia, A., Orts-Escolano, S., Oprea, S., Garcia-Rodriguez, J., Azorin-Lopez, J., Saval-Calvo, M., Cazorla, M.: Multi-sensor 3d object dataset for object recognition with full pose estimation. *Neural Computing and Applications* **28**(5), 941–952 (2017)
13. Gourier, N., Hall, D., Crowley, J.L.: Estimating face orientation from robust detection of salient facial features. In: *ICPR International Workshop on Visual Observation of Deictic Gestures*. Citeseer (2004)
14. Gu, S., Holly, E., Lillicrap, T., Levine, S.: Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3389–3396 (2017)
15. Gynnild, A.: The robot eye witness: Extending visual journalism through drone surveillance. *Digital journalism* **2**(3), 334–343 (2014)

16. Hasselt, H.v., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, AAAI'16, pp. 2094–2100. AAAI Press (2016). URL <http://dl.acm.org/citation.cfm?id=3016100.3016191>
17. Haykin, S., Network, N.: A comprehensive foundation. *Neural networks* **2**(2004), 41 (2004)
18. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Proceedings of the International Conference on Machine Learning, pp. 448–456 (2015)
19. Kazemi, V., Josephine, S.: One millisecond face alignment with an ensemble of regression trees. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1867–1874 (2014)
20. Kirk, D.E.: Optimal control theory: an introduction. Courier Corporation (2012)
21. Krüll, W., Tobera, R., Willms, I., Essen, H., von Wahl, N.: Early forest fire detection and verification using optical smoke, gas and microwave sensors. *Procedia Engineering* **45**, 584–594 (2012)
22. Levine, S., Pastor, P., Krizhevsky, A., Quillen, D.: Learning hand-eye coordination for robotic grasping with large-scale data collection. In: Proceedings of the International Symposium on Experimental Robotics, pp. 173–184 (2016)
23. Li, Y.: Deep reinforcement learning: An overview. arXiv preprint arXiv:1701.07274 (2017)
24. Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 (2015)
25. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. In: Proceedings of the International Conference on Machine Learning, pp. 1928–1937 (2016)
26. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529 (2015)
27. Nägeli, T., Meier, L., Domahidi, A., Alonso-Mora, J., Hilliges, O.: Real-time planning for automated multi-view drone cinematography. *ACM Transactions on Graphics* **36**(4), 132 (2017)
28. Ng, A.Y., Harada, D., Russell, S.: Policy invariance under reward transformations: Theory and application to reward shaping
29. Nousi, P., Patsiouras, E., Tefas, A., Pitas, I.: Convolutional neural networks for visual information analysis with limited computing resources. In: Proceedings of the IEEE International Conference on Image Processing, pp. 321–325 (2018)
30. Ohayon, S., Rivlin, E.: Robust 3d head tracking using camera pose estimation. In: Proceedings of the International Conference on Pattern Recognition, vol. 1, pp. 1063–1066 (2006)
31. Pan, S.J., Yang, Q., et al.: A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* **22**(10), 1345–1359 (2010)
32. Passalis, N., Tefas, A.: Concept detection and face pose estimation using lightweight convolutional neural networks for steering drone video shooting. In: Proceedings of the European Signal Processing Conference, pp. 71–75 (2017)
33. Passalis, N., Tefas, A.: Deep reinforcement learning for frontal view person shooting using drones. In: Proceedings of the IEEE Conference on Evolving and Adaptive Intelligent Systems, pp. 1–8 (2018)
34. Passalis, N., Tefas, A., Pitas, I.: Efficient camera control using 2d visual information for unmanned aerial vehicle-based cinematography. In: Proceedings of the International Symposium on Circuits and Systems (to appear) (2018)
35. Plappert, M.: keras-rl. <https://github.com/matthiasplappert/keras-rl> (2016)
36. Ranjan, R., Patel, V.M., Chellappa, R.: Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
37. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction, vol. 1. MIT press Cambridge (1998)
38. Tang, L., Liu, Y.J., Tong, S.: Adaptive neural control using reinforcement learning for a class of robot manipulator. *Neural Computing and Applications* **25**(1), 135–141 (2014)
39. Taylor, M.E., Stone, P.: Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* **10**(Jul), 1633–1685 (2009)
40. Tieleman, T., Hinton, G.: Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. COURSE: Neural networks for machine learning **4**(2), 26–31 (2012)
41. Toshev, A., Szegedy, C.: Deeppose: Human pose estimation via deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1653–1660 (2014)
42. Triantafyllidou, D., Nousi, P., Tefas, A.: Fast deep convolutional face detection in the wild exploiting hard sample mining. *Big Data Research* (2017)
43. Tzelepi, M., Tefas, A.: Human crowd detection for drone flight safety using convolutional neural networks. In: Proceedings of the European Signal Processing Conference, pp. 743–747 (2017)
44. Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 16, pp. 2094–2100 (2016)
45. Vong, C.M., Tai, K.I., Pun, C.M., Wong, P.K.: Fast and accurate face detection by sparse bayesian extreme learning machine. *Neural Computing and Applications* **26**(5), 1149–1156 (2015)

46. Yang, S., Luo, P., Loy, C.C., Tang, X.: Wider face: A face detection benchmark. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5525–5533 (2016)
47. Zhang, F., Leitner, J., Milford, M., Upcroft, B., Corke, P.: Towards vision-based deep reinforcement learning for robotic motion control. arXiv preprint arXiv:1511.03791 (2015)
48. Passalis, N., Tefas, A.: Training Lightweight Deep Convolutional Neural Networks Using Bag-of-Features Pooling. *IEEE Transactions on Neural Networks and Learning Systems* **30**(6), 1705 – 1715 (2019)
49. Passalis, N., Tefas, A.: Deep reinforcement learning for controlling frontal person close-up shooting. *Neuro-computing* **335**, 37–47 (2019)
50. Zhu, J.Y., Agarwala, A., Efros, A.A., Shechtman, E., Wang, J.: Mirror mirror: Crowdsourcing better portraits. *ACM Transactions on Graphics (SIGGRAPH Asia 2014)* **33**(6) (2014)
51. Wooldridge, M.: An introduction to multiagent systems. John Wiley & Sons, 2009.