


# A Fully-Autonomous Aerial Robot for Search and Rescue Applications in Indoor Environments using Learning-Based Techniques

Carlos Sampedro<sup>1</sup>  · Alejandro Rodriguez-Ramos<sup>1</sup> · Hriday Bavle<sup>1</sup> · Adrian Carrio<sup>1</sup> · Paloma de la Puente<sup>1</sup> · Pascual Campoy<sup>1</sup>

## Abstract

Search and Rescue (SAR) missions represent an important challenge in the robotics research field as they usually involve exceedingly variable-nature scenarios which require a high-level of autonomy and versatile decision-making capabilities. This challenge becomes even more relevant in the case of aerial robotic platforms owing to their limited payload and computational capabilities. In this paper, we present a fully-autonomous aerial robotic solution, for executing complex SAR missions in unstructured indoor environments. The proposed system is based on the combination of a complete hardware configuration and a flexible system architecture which allows the execution of high-level missions in a fully unsupervised manner (i.e. without human intervention). In order to obtain flexible and versatile behaviors from the proposed aerial robot, several learning-based capabilities have been integrated for target recognition and interaction. The target recognition capability includes a supervised learning classifier based on a computationally-efficient Convolutional Neural Network (CNN) model trained for target/background classification, while the capability to interact with the target for rescue operations introduces a novel Image-Based Visual Servoing (IBVS) algorithm which integrates a recent deep reinforcement learning method named Deep Deterministic Policy Gradients (DDPG). In order to train the aerial robot for performing IBVS tasks, a reinforcement learning framework has been developed, which integrates a deep reinforcement learning agent (e.g. DDPG) with a Gazebo-based simulator for aerial robotics. The proposed system has been validated in a wide range of simulation flights, using Gazebo and PX4 Software-In-The-Loop, and real flights in cluttered indoor environments, demonstrating the versatility of the proposed system in complex SAR missions.

**Keywords** Autonomous robots · Search and rescue · Supervised learning · Reinforcement learning · Deep learning · Image-based visual servoing

## 1 Introduction

Research interests in Unmanned Aerial Vehicles (UAVs) have increased considerably in the last decade, especially regarding multirotor aerial robots, owing to their good maneuverability for solving complex missions. In this context, the research community has mostly focused on outdoor applications where the availability of Global Navigation Satellite System (GNSS) information facilitates the autonomous navigation of the UAV. Furthermore, in this kind of outdoor scenarios there are usually no disturbances

in the flight-path of the UAV related to ground obstacles. However, nowadays there is an increasing demand for solving complex indoor applications, such as surveillance, search and rescue, package delivery in large indoor facilities, etc. In this kind of applications, the absence of GNSS makes necessary the usage of the exteroceptive sensors mounted onboard the UAV for navigation and decision-making purposes.

The development of fully-autonomous missions in such indoor environments requires the research on smaller-size aerial robotic systems, which pose a great challenge on the system design, due to their limited payload and computational capabilities. These limitations impose the necessity to develop efficient and lightweight algorithms in order to find a proper balance between mission performance and computational resources available onboard the aerial robot. This balance becomes especially critical when particularly

---

✉ Carlos Sampedro  
carlos.sampedro@upm.es

<sup>1</sup> Centre for Automation and Robotics, Universidad Politécnica de Madrid (CSIC-UPM), Madrid, Spain

---

complex tasks have to be performed by the aerial robot without human supervision as more complex decision-making algorithms are required. The aforementioned constraints are present in most of the SAR scenarios, where the aerial robot usually requires to perform high-level tasks such as exploration of unknown scenarios, navigation with collision avoidance, target recognition and interaction, among others. In this paper, we focus our efforts on the development of an autonomous aerial robot capable of performing such high-level missions in an unsupervised manner (i.e. without human intervention) with a special interest in the versatility and ease of adaptation of the algorithms developed for object recognition and interaction.

Regarding the object recognition problem, classical computer vision algorithms are generally very dependent on the conditions of the environment where they operate (e.g. lighting conditions, variety of backgrounds, presence of clutter, etc.) which implies the need to readjust the parameters of the algorithm for each new environment in order to obtain a precise detection of the object and remove possible false positives in the image plane. Furthermore, this adjusting procedure can become an onerous task as it is usually carried out by trial and error tests. The use of machine learning techniques, when trained on meaningful datasets, allows overcoming these limitations, providing more versatile solutions which can be executed in a wide range of environments [7]. However, some of the recent machine learning models for object recognition [22, 30, 40] consist of considerably deep models with a large amount of parameters that have not been designed for operating onboard a robotic platform with hard computational constraints such as a UAV.

With respect to the task of interacting with the target for rescue operations, in this paper we characterize these tasks by means of IBVS methods which allow the interaction with targets in a wide variety of SAR missions. Within these missions, operation in SAR disaster scenarios has a special interest as it usually involves the interaction with the detected target by means of delivering some required items such as medicines, food, etc. [14, 21], where IBVS techniques can provide versatile and computationally efficient solutions. In this direction, classical IBVS methods usually require a tedious tuning stage of their parameters when changing to different operating conditions. This fact can eventually appear in SAR scenarios where different aerial robotic platforms may be used depending on the environment. Furthermore, classic IBVS methods can suffer from convergence and stability problems [10, 12].

Inspired by the aforementioned limitations, in this paper we propose a fully-autonomous UAV featured with learning-based techniques which can provide flexible and versatile solutions to indoor SAR missions. The main contributions of the proposed system are summarized here: i) A custom UAV has been built with a flexible system

architecture which allows the efficient coordination between the planning, situation awareness, feature extraction and executive systems for solving complex SAR missions. ii) This flexibility has allowed the integration of learning-based techniques for object recognition and object interaction. Concretely, several supervised learning classifiers, including computationally-efficient CNN models, have been trained and evaluated for target/background classification. In addition, a novel IBVS algorithm based on Deep Deterministic Policy Gradients (DDPG) [29] has been implemented and validated for solving IBVS tasks in rescue operations, comparing its performance with classic IBVS techniques. iii) An extensive evaluation of the previous learning-based components and the whole system in cluttered indoor SAR scenarios has been conducted in both simulated and real flights.

In order to obtain a reliable testbed for experimentation in SAR scenarios, in this paper we have adopted as the main use case the missions proposed for the 2016 International Micro Air Vehicle Competition<sup>1</sup> (IMAV), where several SAR problems had to be addressed. In the IMAV 2016 indoors competition, challenging high-level missions were designed, ranging from autonomous building entering and exiting, indoor exploration of unknown scenarios, object recognition, etc. In addition, UAVs were required to perform object interaction tasks, such as grasping a cylindrical item and releasing it into a cylindrical bucket, with the option of pre-loading the items previous to the takeoff maneuver. The latter scenario has been extensively studied in this work in order to provide our previous system [43] with more high-level functionalities through learning-based techniques.

The remainder of this paper is organized as follows: Section 2 introduces the related work; Section 3 describes the hardware configuration adopted in our aerial robotic platform. The system description is presented in Section 4. Section 5 presents the experiments performed in simulated and real scenarios, with their respective results, before we discuss them in Section 6 and finally, Section 7 concludes the paper, and points out future research directions.

## 2 Related Work

In the following paragraphs, some of the most relevant solutions for the autonomous operation of UAVs in SAR missions are covered in chronological order [1, 5, 14, 16, 41, 46, 48, 49]. We also refer to articles in which the delivery of specific items from UAVs in emergency situations has been studied [19, 21, 50]. Right after, we cover other relevant developments aimed towards the execution of autonomous

---

<sup>1</sup>IMAV 2016 official website: <http://www.imavs.org/2016/>

---

UAV missions outside the field of SAR applications [2, 3, 23, 53]. We conclude by referring to other works applying vision-based deep reinforcement learning to UAV navigation [38, 42] as well as other relevant uses of deep reinforcement learning in visual control tasks [28, 51, 52].

There is an increasing number of recent studies aiming at UAVs as a potentially useful complement to SAR applications. Early developments in high-level artificial intelligence applied to aerial robotics were introduced in Doherty et al. [14]. In particular, UAV autonomous missions were implemented for the SAR of injured civilians, with robots being able to scan designated areas, trying to identify injured civilians and attempting to deliver medical and other supplies to identified victims in realistic urban scenarios. The specific techniques from this work, used to detect humans at a high frame rate onboard an autonomous UAV, were described in detail in [41]. These techniques were applied in a real-world outdoor environment using visible and thermal infrared cameras. In their work, detected human positions were geolocated and a map of points of interest was built. The resulting map is proposed to plan medical supply delivery during a disaster relief effort.

In the context of SAR technology developments, the UAV Challenge-Outback Rescue has been established as an important international competition, where participants have been required to perform UAV SAR missions, which typically have involved executing autonomous take-off, navigation for an aerial search, and landing maneuvers. These exercises had associated image processing and control tasks needed to identify and deliver an emergency medical package to a mannequin simulating a lost person, placed in a  $4 \text{ km} \times 6 \text{ km}$  area. Among the solutions provided, one of the most relevant presented in [16] proposed a fixed-wing UAV featuring GPS-based navigation, ground image acquisition, and payload delivery, all implemented in a low-cost platform.

The work by Tomic et al. [49] introduced a modular and extensible software and hardware framework designed for the autonomous execution of SAR missions using aerial robots, which was successfully tested on a quadrotor platform. However, while using multiple sensors (four cameras and a laser scanner) the proposed system did not feature any collision avoidance capabilities.

In 2015, a pilot study was conducted by Abrahamsen [1] to assess the concept and feasibility of using a remotely piloted aircraft (RPA) system to support remote sensing in simulated major incident exercises. A custom-made, remotely controlled UAV with vertical takeoff and landing was equipped with visible and thermal infrared cameras, a laser beam, a mechanical gripper arm and an avalanche transceiver. Successful missions were executed for five simulated exercises, demonstrating that UAVs are suitable for carrying small payloads as well as useful tools to support

situation assessment and information exchange at major incident scenarios.

In [46], Scherer et al. tested another interesting modular architecture of a UAV system for SAR missions in an outdoor environment. The objective of the mission was to detect a ground target by means of color, text or shape, and to provide a live aerial video stream for remote monitoring. Their proposal consisted of a swarm of multi-copters coordinated to operate as a communications relay using a distributed control system. The system was implemented using the Robot Operating System (ROS) [39] and was capable of providing a real-time video stream from a UAV to one or more base stations using a wireless communication infrastructure. The proposed system supported a heterogeneous set of UAVs and image sensors and allowed the operator to select different levels of autonomy.

In the study published by Sun et al., a camera-based target detection and positioning system was developed and integrated into a fully autonomous fixed-wing UAV [48]. The system was capable of onboard and real-time target identification, post-target identification and localization, and aerial image collection for further mapping applications. Its performance was assessed using several simulated SAR missions, demonstrating its reliability and efficiency.

Deep learning was applied for supporting UAV SAR operations in [5]. In this work, a sequence of images of avalanche debris captured by a UAV was processed with a pretrained CNN model to extract discriminative features. A trained linear Support Vector Machine (SVM) was integrated at the output of the CNN to detect objects of interest. Moreover, they introduced a preprocessing method to increase the detection rate and a postprocessing method based on a Hidden Markov Model to improve the prediction performance of the classifier. Experimental results conducted on two different datasets at different levels of resolution showed that the detection performance increased when incrementing the resolution, at the cost of raising the computation time.

Developments for drone delivery of emergency items in search and rescue missions have as well been analyzed in the literature. Examples include: drug shipments [21], delivery of defibrillators [19] and life rings [50].

Other recent developments focused on providing UAVs with high levels of autonomy outside the field of application of SAR missions are discussed next. In the work of Bacharach et al. [3], a quadrotor helicopter equipped with a laser rangefinder was designed and implemented to autonomously explore and map unstructured and unknown indoor environments. The paper highlighted the difficulties of applying to UAVs algorithms that were originally developed for Unmanned Ground Vehicles (UGVs). Interesting solutions were described in this work, such as a multilevel sensing and control hierarchy, a high-speed laser

---

scan-matching algorithm, an Extended Kalman Filter (EKF) for data fusion, a high-level SLAM implementation, and an exploration planner. The manuscript showed experimental results demonstrating the helicopter’s ability to navigate accurately and autonomously in unknown environments. Algorithms originally conceived for UGVs were also exploited by Grzonka et al. [23] to increase UAV autonomy. In this case, they proposed a general navigation system that enabled a small-sized quadrotor platform to autonomously operate in indoor environments. A similar work was published by Achtelik et al. [2], which presented a software architecture providing a quadrotor helicopter with the capabilities to autonomously navigate, explore and locate objects of interest in unknown, unstructured indoor environments.

Results specific to autonomous navigation in indoor corridors were presented by Zingg et al. [53]. In their approach for wall collision avoidance, a depth map based on optical flow from images captured by an onboard omnidirectional fisheye camera was used. Inertial Measurement Unit (IMU) data was also used for compensating rotational effects of the optical flow.

Several implementations of visual control for UAVs can be found in the literature, but very few presenting it as one capability among several other ones in the context of an autonomous mission. Some of the aforementioned developments make use of computer vision geometry to determine navigation waypoints, but only [49] used visual information to provide control feedback in real time.

There are as well very few developments exploiting vision-based deep reinforcement learning for UAV navigation. In particular, Sadeghi et al. [42] introduced the  $CAD^2RL$  learning method, which allows collision-free navigation in a real indoor environment using synthetic data from 3D CAD models as the only training data. Another example is the work by Polvara et al. [38], who made use of Deep Q-Networks for the autonomous landing of a quadrotor.

Other interesting applications of deep reinforcement learning for solving visual control tasks outside the field of aerial robotics are mentioned next. Lee et al. [28] proposed to perform a visual servoing task by extracting deep features instead of using pixels or keypoints. The best features to solve the task were then selected using a Q-iteration algorithm. Also Zhang et al. [51] proposed to use vision-based deep reinforcement learning for controlling the motion of a three-joint robot manipulator. Finally, Zhu et al. [52] presented an efficient algorithm for visual navigation in indoor scenes using deep reinforcement learning. The algorithm was trained using high-quality 3D scenarios allowing for physical interaction with the objects in the scene.

In contrast to all the aforementioned developments, our work proposes a fully-autonomous UAV that is not only capable of autonomously navigating in indoor cluttered


environments with situational awareness, but can also interact with static and moving targets, which can be automatically detected and followed to precisely deliver items for rescue purposes. Furthermore, while much of the discussed literature focuses on accomplishing specific tasks only, and many developments are evaluated in computer simulations only, the solution proposed here focuses on complex missions involving multiple heterogeneous tasks and has been evaluated in detail in both simulated and real flights.

### 3 Hardware Configuration

Search and Rescue missions in indoor scenarios usually involve hard constraints relative to the clearance within indoor passages and the necessity of carrying items for rescue purposes. These constraints were also present in the IMAV 2016 competition, where the minimum clearance of the passages within the indoor environment was 1 m wide and the aerial robot was required to preload several items of 100 g each. Considering these limitations, it was necessary to build a custom UAV (see Fig. 1) relatively small in size with the adequate capacities for carrying the sensors and actuators required for localization, navigation, and object recognition and interaction.

Based on the aforementioned constraints, a custom UAV with relatively small dimensions (62 cm × 62 cm × 40 cm, including propellers of 28 cm) has been designed and built with a total takeoff weight of 3.2 kg, a maximum payload capacity of 1 kg and a maximum flight time of 12 min. The onboard computer consists of an Intel NUC6i5SYK featuring a 2.9 GHz Intel Core i5-6260U CPU. The avionics of the UAV are managed by a Pixhawk [34] autopilot, which integrates an Inertial Measurement Unit (IMU), a barometer and a magnetometer. The exteroceptive sensors mounted onboard consist of a Hokuyo laser rangefinder UTM-30LX with a horizontal field of view of 270° and an angular resolution of 0.25° with a maximum range of 30 m, an Intel Realsense R200 camera with an RGB image size of 640 × 480 pixels, a standard RGB 180° fisheye-lens bottom-looking camera with an image size of 640 × 480 pixels, and a Lightware SF10/A altimeter with a maximum range of 25 m (see Fig. 1a). The communication between the autopilot, proprioceptive, exteroceptive sensors and the onboard computer is performed over USB connections. In addition, appropriate electronic components and holding devices have been designed and integrated for object interaction tasks. For this purpose, two small curved hooks have been integrated into the UAV framework (see Fig. 1b). These hooks are controlled by two servo motors which are actuated via an Arduino board when a signal of *target locked* is commanded in order to release the items preloaded onboard the UAV.

**Fig. 1** Proposed aerial robot, showing 1) Hokuyo laser rangefinder UTM-30LX. 2) Intel NUC6i5SYK computer. 3) Pixhawk autopilot. 4) Lightware SF10/A altimeter. 5) Fisheye bottom camera. 6) Intel realsense R200 camera. 7) Arduino pro-mini. 8) Designed hooks for holding items



## 4 System Description

The system proposed in this paper has been built on top of our software architecture named Aerostack [45]. In this architecture, the components are organized in different layers of abstraction (see. Fig. 2) which provide high modularity and flexibility. This modularity has permitted the easy integration of the proposed learning-based components.

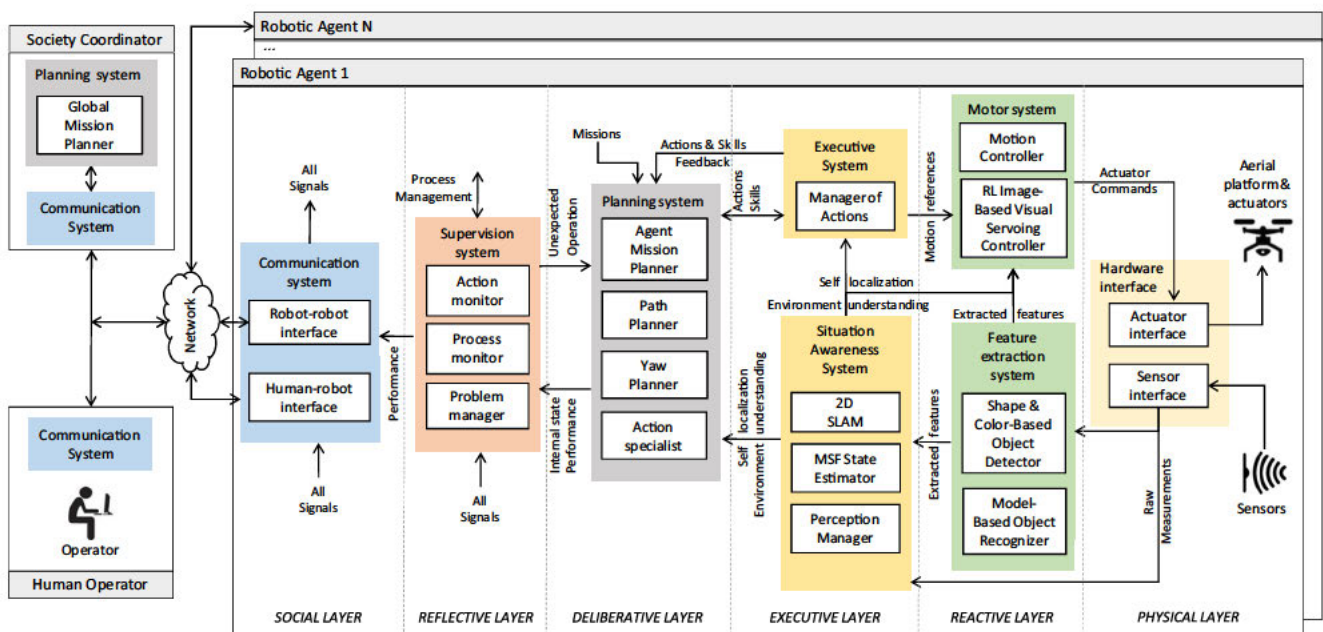
### 4.1 Planning System

#### 4.1.1 Mission Planner

The mission planner subsystem is composed of two main components, named Global Mission Planner (GMP) and Agent Mission Planner (AMP). In the following paragraphs we summarize the main capabilities of both components.

1. Global Mission Planner (GMP). In this component resides the higher level of intelligence in terms of mission planning. The design of the GMP allows the mission generation for a single agent or a swarm of agents. In the latter case, the GMP is able to distribute the mission between the agents in the swarm in order to optimize the accomplishment of such mission in terms of the explored area [44].

The main responsibility of the GMP is focused on the mission generation. For this purpose, its flexible design allows the specification of the mission in two operating modes: manual or automatic. In the manual mission generation mode, the definition of the mission is performed using an XML-based language, in which the human operator can define the mission task by task. In the automatic mission generation mode, the GMP interprets a high-level mission command provided by



**Fig. 2** System architecture. All the components developed in this work have been integrated into the different layers of the Aerostack framework [45]

---

a human operator (e.g. *find a target, explore, etc.*). Additional inputs in this mode comprise the dimensions of the area to be explored as well as the number of mission points. Using these inputs, the GMP is able to automatically generate mission points by applying a *K-means* clustering over points randomly distributed over the area to be explored. For a detailed explanation of this functionality we refer the reader to [44]. Once the global mission is prepared, it is distributed through the agent or agents in the swarm and received by the AMP.

One important functionality implemented in the GMP is its capability of concatenating several missions [43]. Based on this, the GMP has an active *list of missions* per agent. Once the last mission in the list is accomplished, the GMP is able to recover the previous mission in its corresponding state (i.e. current task). This functionality acquires an utmost importance in SAR missions, where the system has to react to the wide range of events that may occur during the execution of this type of missions.

2. Agent Mission Planner (AMP). This component is located at the agent's level and is responsible for scheduling task by task the received mission. For this purpose, the AMP acts as an interface between the GMP and the rest of components in the architecture.

Since the area to be explored is unknown *a priori*, the mission points generated by the GMP can fall within an obstacle. In order to address this problem, the AMP is capable of generating *safety points* when the current mission point falls within an obstacle. In order to generate a *safety point*, the AMP implements an iterative method in which random points lying on a safety circumference of predefined radius are generated. After several iterations of the algorithm, and if no *safety point* has been obtained in the current circumference, its radius is incremented and the iterative method continues. In this iterative procedure, the AMP communicates with the Path Planner in each iteration, until an obstacle-free point is obtained.

#### 4.1.2 Path Planner

The path planner component relies both on the use of a precise lidar sensor and a robust localization and mapping algorithm. The path planning algorithm utilized in the proposed architecture is based on an existent ROS navigation package [33], which was originally designed for differential-drive and holonomic-wheels robots, and has been adapted in this work to the Aerostack architecture, enabling its operation with multirotor UAVs. Furthermore, the original 2D functionality of the mentioned ROS planner package has been extended in order to provide

3D navigation capabilities by adding the remaining altitude coordinate as a constant value (by default) to each intermediate point in the path. However, the AMP can dynamically modify its value within the execution of a mission in order to fulfill specific requirements relative to the current environment.

The Path Planner component requires a 2D occupancy grid map (see Section 4.2.1) as well as a mission point generated by the AMP (in world frame of reference) for its normal operation. The 2D occupancy grid map is subsequently translated into a 2D cost map in which cost values are propagated out of occupied cells based on an inflation radius parameter. A detailed explanation of the algorithm and its components can be reviewed at [33].

#### 4.1.3 Yaw Planner

The yaw planner is in charge of associating a yaw angle to each waypoint within a 2D path based on the direction towards the commanded mission point or considering the AMP directives. Taking into consideration normal mission conditions, a specific policy has been defined:

- Middle waypoints: Orientation is set to a constant value according to the direction of navigation at each time step towards the commanded mission point.
- Last waypoint: Orientation is derived from the AMP directives extracted from the current task within the mission. This can be utilized for exploration purposes, in which different yaw angles can be specified at each mission point.

Following this policy, a UAV is considered to maximize the area covered by both the lidar and the image sensor's field of view, in order to plan throughout the optimum path and to avoid blind zones which can lead to a collision.

## 4.2 Situation Awareness System

### 4.2.1 2D Localization and Mapping

In this work, localization and mapping capabilities have been integrated by means of a state-of-the-art 2D SLAM algorithm [25], which has been extensively tested for ground robots in Urban Search and Rescue (USAR) missions, and can be also utilized in platforms that exhibit roll/pitch motion by transforming the laser scan into a local stabilized coordinate frame. Using the information coming from a 2D lidar sensor, in this algorithm a fast approximation of map gradients, and a multi-resolution grid map representation for mitigating local minima problems, are utilized in order to obtain a stable mapping and a robust scan matching. The latter can be further improved if a source of 3D pose is available and projected into the xy-plane

in order to initialize the optimization process of the scan matcher. The 2D pose obtained from the scan matching includes the xy-position of the aerial robot and its yaw angle referred in the world frame of reference. Finally, the map representation is encoded into a 2D occupancy grid map, including *occupied*, *non-occupied* and *non-explored* cell types, which is used by the Path Planner component for navigation purposes.

#### 4.2.2 Multi-sensor Fusion State Estimation

The objective of this component is to provide a full 6 DOF pose and the respective velocities of the UAV, enabling navigation using the Path Planner component described in Section 4.1.2. In this direction, our proposed architecture integrates two separate state estimation components which can be combined in order to provide a higher level situational awareness functionality or be used separately depending on the requirements of the mission.

**Flight Altitude State Estimator:** The 2D SLAM algorithm explained in Section 4.2.1, provides a 2D map of the environment, enabling obstacle detection and avoidance in a 2D plane at the given flight altitude of the UAV. Sensors such as laser altimeters or similar range sensors can lead to errors in the flight altitude estimation when flying above ground obstacles, as the measurements get referred to them instead of the ground surface.

In order to accurately estimate the flight altitude of the UAV in the presence of several ground obstacles, we propose an EKF-based algorithm which is able to estimate the flight altitude of the UAV as well as the elevation of the ground obstacles (more details can be found in [4]). This is achieved by fusing the measurements coming from the IMU, the barometer, and the laser range altimeter sensors.

The proposed state estimator considers a state vector  $\mathbf{x} \in \mathbb{R}^{10}$  based on the combination of four main components:  $\mathbf{x}_R$ ,  $\mathbf{x}_G$ ,  $\mathbf{x}_I$ , and  $\mathbf{x}_B$ , which represent the state of the robot, ground object, IMU sensor and barometer sensor respectively. The corresponding state of the aerial robot is defined by  $\mathbf{x}_R = \begin{pmatrix} \Omega_{xy}^T & \omega_{xy}^T & t_{z_R} & v_{z_R} & a_{z_R} \end{pmatrix}$ , where  $\Omega_{xy}^T = (\phi, \theta)^T$  are the roll and pitch Euler angles,  $\omega_{xy}^T = (\omega_x, \omega_y)^T$  represent the  $x$  and  $y$  angular velocities of the aerial robot in the UAV frame, and  $t_{z_R}$ ,  $v_{z_R}$  and  $a_{z_R}$  are the vertical coordinates of the position, velocity and acceleration of the aerial robot in the world frame of reference. Assuming that the robot changes its vertical acceleration and angular velocity slowly, we adopt a constant vertical acceleration and constant angular velocity as process model. The ground object state is defined by  $\mathbf{x}_G = t_{z_G}$ , where  $t_{z_G}$  is the altitude of the ground object in the world frame. Obstacles are set so that they always have positive altitude with

respect to the ground plane. Finally, the IMU and the barometer sensors contribute to the state with their corresponding biases, thus  $\mathbf{x}_I = b_{a_z}$  and  $\mathbf{x}_B = b_{b_z}$ , where  $b_{a_z}$  and  $b_{b_z}$  are the biases in the vertical acceleration and flight altitude measurements respectively.

**Robot State Estimator** This component is able to combine the measurements from the 2D SLAM (see Section 4.2.1) with the flight altitude estimator (see Section 4.2.2) or the IMU sensor in order to provide complete pose and velocity estimates of the UAV in the world frame of reference.

The Robot State Estimator is a standard ROS package [35] which implements an EKF-based estimator with state vector  $\mathbf{x}_R \in \mathbb{R}^{12}$ , being  $\mathbf{x}_R = (\Omega^T \ \mathbf{p}^T \ \mathbf{v}^T)$ , where  $\Omega^T = (\phi, \theta, \psi)^T$  are the roll, pitch and yaw Euler angles,  $\mathbf{p}^T = (p_x, p_y, p_z)^T$  represents the position of the robot, and  $\mathbf{v}^T = (v_x, v_y, v_z, \omega_x, \omega_y, \omega_z)$  is the vector containing the linear and angular velocities of the aerial robot.

It includes a non-linear measurement model able to fuse any robot pose or velocity measurements, provided by any number of sensors. In contrast to the Flight Altitude State Estimator, this EKF model does not incorporate in its state vector any biases present in the sensor measurements nor the altitude of the ground obstacles.

Our system integrates the previously described components for multi-sensor state estimation in a versatile manner, so they can be employed in different configurations depending on the selected hardware and mission requirements at hand.

#### 4.2.3 Perception Manager

The Perception Manager component is in charge of managing and centralizing the perception events that can occur during the execution of a mission (e.g. object *recognized*, *picked* and *released*).

In order to obtain a proper management of the perception events, this component integrates the information regarding the current situation of the states of the different objects which the UAV can interact with, together with the own internal state of the Perception Manager (e.g. *Exploring*, *Going For Picking Item*, *Going For Releasing Item*). An example of the initial configuration of such states for a SAR mission applied to the use case of IMAV 2016 is provided in Table 1. When a perception event is detected by the Perception Manager, (e.g. target recognized) the current state of the objects as well as its internal state are evaluated and updated. Based on this evaluation, the Perception Manager can request a *mission adaptation* event to the GMP. As an example, and taking the initial configuration presented in Table 1, if a *bucket* object is recognized in the current time instant, and since both corresponding *items* are already picked up, the Perception Manager will generate a

**Table 1** Example of the initial state of the objects in a Search and Rescue mission, where 0/1 represent a false/true statement

Object	State	Recognized	Picked	Released
Item A		0	1	0
Bucket A		0	0	0
Item B		0	1	0
Bucket B		0	0	0

mission adaptation event consisting of *Going For Releasing Item*.

Furthermore, another important functionality implemented in the Perception Manager consists in mapping the objects of interest (e.g. targets) in a SAR mission. For obtaining an adequate mapping of the objects that are being recognized during the execution of the mission, the Perception Manager provides a higher level of understanding over the recognition events coming from the Model-Based Object Recognizer component (see Section 4.3.2). For this purpose, the Perception Manager is responsible for referring the relative pose of the recognized objects in the frame of reference of the corresponding device (e.g. front camera) to the world frame of reference. This transformation is computed using Eq. 1. In addition, in order to provide a robust final pose of the recognized objects, the Perception Manager has an internal buffer for storing the pose of the objects being recognized over time. This buffer is utilized by a filter in which the poses away from the average pose are removed in an iterative process.


$${}^W T_O = {}^W T_U \cdot {}^U T_C \cdot {}^C T_O \quad (1)$$

where  $T$  stands for a  $4 \times 4$  homogeneous matrix transformation.  ${}^W T_O$  computes the global transformation from world to object frame of reference.  ${}^W T_U$  is the transformation from world to UAV frame of reference, which is given by the Robot State Estimator component.  ${}^U T_C$  represents a rigid transformation from UAV to front camera reference frame, and  ${}^C T_O$  provides the transformation from camera to object frame of reference, computed online using a Perspective- $n$ -Point (PnP) algorithm (see Fig. 3).

### 4.3 Feature Extraction System

#### 4.3.1 Shape and Color-Based Object Detector

This component of the system has been designed taking into account the hard computational constraints that are present when running real-time computer vision algorithms onboard an aerial robot. For this reason, this component is based on “low-cost” algorithms meant for increasing the effectiveness of posterior recognition stages. Its main



**Fig. 3** Frames of reference defined in the proposed solution. the dotted line in the object depicts the virtual tilted plane used for pose estimation purposes. Points  $p_0$  to  $p_3$  are projected into the image plane for IBVS purposes

objective is to generate candidate proposals to the Model-Based Object Recognizer component (see Section 4.3.2), and it is also utilized as the main object detection algorithm when the image processing is performed using the bottom camera of the UAV.

For addressing the aforementioned constraints, the functionalities implemented within this component are based on two main features: color and shape information. The former is based on heuristics extracted from the use case of IMAV 2016, in which the targets were defined as red and blue *buckets* or *items*. In order to segment the image based on color information, the HSV color space has been utilized, whose optimal ranges for each channel have been empirically derived. Shape information (e.g. aspect ratio) is mainly utilized in the image plane for differentiating between identical objects in terms of 3D appearance and color (e.g. a red *bucket* is defined as a cylinder of 0.25 m radius and 0.3 m height, while a red *item* consists of a cylinder of 0.1 m radius and 0.1 m height).

Using these functionalities, the Shape and Color-Based Object Detector acts as an effective preliminary stage in order to provide candidates to the Model-Based Object Recognizer component. Furthermore, using the combination of shape and color information, this component provides an intra-class classification of the detected object (e.g. *itemA*, *bucketA*, *itemB*, *bucketB*, where all these objects are cylinders and only differ in the color and shape as stated previously). In addition, it is important to remark the importance of this component in the posterior stage for computing the relative pose of the object being detected. As this component provides object proposals based on color information, the detected Region Of Interest (ROI) in the image fits accurately the contour of the object (see Figs. 17d and 18d for an example), which will lead to a better pose estimation as compared to other object proposal algorithms based on predefined ROI sizes such as sliding window



approaches, where the ROI can be slightly away from the object contour.

### 4.3.2 Model-Based Object Recognizer

Object detectors based on predefined knowledge of the object to be detected, such as color, shape, etc, can be very specific and prone to false positives. In order to provide a robust detection of the target, reducing its vulnerability to the environment conditions (e.g. lighting conditions) we implement this component, whose core is based on a supervised learning classifier for target/background segmentation.

The objective of this component is to recognize and locate the object of interest (target) within the image plane, by providing its corresponding ROI, and in addition it is responsible for the recognition of the target in terms of its 3D location with respect to the frame of reference of the camera. For achieving the aforementioned capabilities, the Model-Based Object Recognizer is composed of four main blocks:

- Object proposal: this block is in charge of generating the candidates within the image to be introduced to the classifier. In this work, we have utilized the candidates generated by the Shape and Color-Based Object Detector (see Section 4.3.1), which can implement an independent object detector component itself or be a part of a higher level recognizer, which is the case when the Model-Based Object Recognizer is operating.
- Feature Extractor: This module is only used in the case of the supervised learning classifiers considered in this work which are not CNN-based models, as CNNs perform an unsupervised feature extraction in the convolution layers. In each candidate ROI generated by the Object Proposal module, Histogram of Oriented Gradients (HOG) [13] features are computed, obtaining a descriptor vector of size 1728. The configuration of the HOG feature extractor is summarized here:
  - Window Size:  $56 \times 72$  pixels.
  - Cell Size:  $8 \times 8$  pixels.
  - Block Size:  $16 \times 16$  pixels ( $2 \times 2$  cells).
  - Block Stride: 8 pixels (50% of block overlapping).
  - Histogram configuration: 9 bins,  $20^\circ$  each (unsigned gradient).
- Classifier: this block implements a supervised learning classifier that has been trained for bucket/background classification. In this work, three supervised learning classifiers have been evaluated: L2 Regularized Logistic Regression (L2R-LR), Support Vector Machines (SVMs) with linear kernel (L-SVM), and CNNs models.

The formulation of the L2R-LR follows the implementation in [18], whose loss function is given by Eq. 2.

$$\min_{\omega} \frac{1}{2} \omega^T \omega + C \sum_{i=1}^l \log(1 + e^{-y_i \omega^T x_i}) \quad (2)$$

where  $\omega$  are the parameters to be learned by the classifier.  $C$  is the regularization parameter, and  $(x_i, y_i)$  is the instance-label pair of the  $i_{th}$  training sample.

The SVM classifier formulation has been defined using the implementation provided in [9] for the primal form (see Eq. 3).

$$\min_{\omega, b, \xi} \frac{1}{2} \omega^T \omega + C \sum_{i=1}^l \xi_i,$$

$$\text{subject to : } y_i (\omega^T \phi(x_i) + b) \geq 1 - \xi_i, \quad \xi_i \geq 0 \quad (3)$$

where  $\omega$  (weights),  $b$  (intercept term) and  $\xi_i$  (*slack variables*) are the parameters to be learned by the SVM classifier.  $C$  is the regularization parameter,  $(x_i, y_i)$  is the instance-label pair of the  $i_{th}$  training sample, and  $\phi(x_i)$  is a *feature mapping* function.

In the analogous dual form of the formulation presented in Eq. 3, a kernel function can be defined as:

$$K(x_i, x_j) = \phi(x_i)^T \phi(x_j) \quad (4)$$

where  $(x_i, x_j)$  are points in the input feature space, and  $\phi$  is a *feature mapping* function.

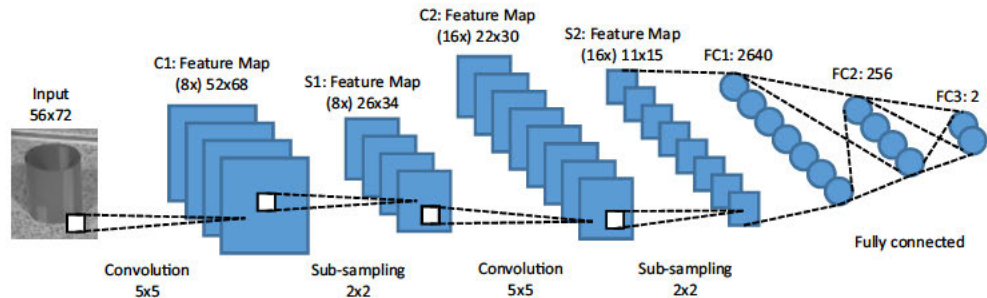
The kernel function can lead to different type of SVM classifiers. In this paper, we consider the SVM classifier with a linear kernel computed using Eq. 5.

$$K(x_i, x_j) = x_i^T x_j \quad (5)$$

Regarding the CNN classifier, its architecture consists of 7 layers: 2 convolutional layers, 2 max pooling layers and 3 fully-connected with one hidden layer of 256 units, using ReLU activation function [36] in each layer except the final one, in which a *softmax* activation function is utilized, being the input to the CNN model an image of  $56 \times 72$  pixels. After the evaluation conducted in Section 5.2.2, the selected supervised learning classifier is based on the architecture presented in Fig. 4.

- Pose estimator: In order to compute the relative pose of the target with respect to the frame of reference of the camera, the Pose Estimator block uses a PnP algorithm taking as input the corners of the previous computed ROI of the detected target (image plane), a set of *object points* (object frame of reference), and the intrinsic camera parameters. The final computed pose is selected so that it minimizes the reprojection error between the detected points in the image and

**Fig. 4** CNN classifier architecture. Architecture of the CNN model (see *CNN5* in Table 2) utilized for bucket/background classification



the projected *object points*, using Levenberg-Marquardt optimization and RANSAC for filtering the outliers. In order to compute the *object points* we assume that the UAV is always flying at a higher altitude than that of the object of interest. Based on this assumption, and taking into account the appearance of a cylindrical object in such situations, the selected *object points* are the four corners that define the tilted plane similar to the one depicted in Fig. 3. These four corners are defined in the object frame of reference given its 3D shape (height and diameter in the case of our target).

The Model-Based Object Recognizer is utilized as the main object recognition algorithm when the image processing is performed using the front camera. The implementation details for training and evaluating the Model-Based Object Recognizer are provided in Section 5.2

## 4.4 Motor System

### 4.4.1 Motion Controller

The Motion Controller utilized in our system architecture [37], is based on a cascaded control loop architecture where the external control loops act in position and the inner loops control the linear velocity of the UAV. The attitude control, which is executed by the Pixhawk autopilot, calculates the propeller rotation rates for the actuators.

Three main flight modes are available within the Motion Controller: trajectory, position or velocity, which can be switched during flight. In the experiments presented in this work, the proposed system is able to automatically switch from position to velocity control mode when the IBVS controller presented in Section 4.4.2 is activated in order to perform the release maneuver of a preloaded item.

Some of the capabilities of the Motion Controller that make it suitable for our purposes include a native saturation of the linear velocity commands in all flight modes and the capability of tuning the control gains after performing system identification of a simple dynamic model. These functionalities allow a smooth navigation in cluttered environments, decreasing potential motion blur in the images captured onboard the UAV.

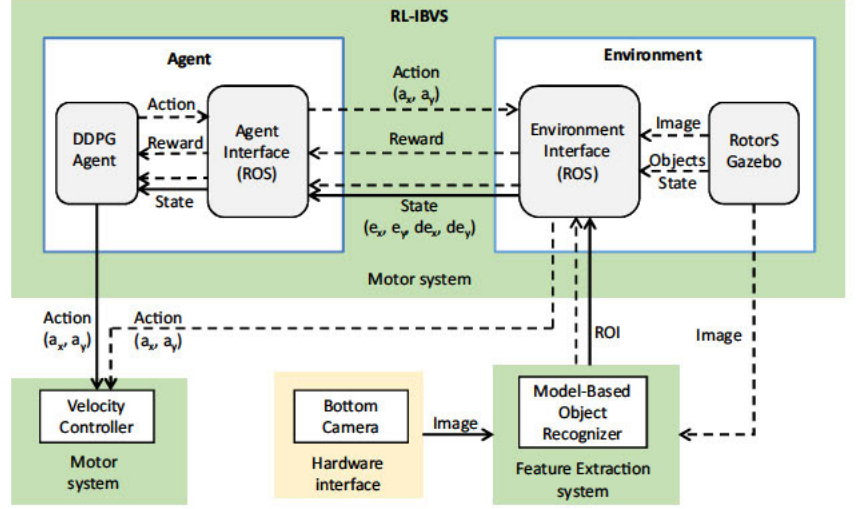
### 4.4.2 Image-Based Visual Servoing through Reinforcement Learning

The Image-Based Visual Servoing through Reinforcement Learning component (from now on RL-IBVS) is based on our own designed reinforcement learning framework which integrates an agent (i.e. an aerial robot executing a deep RL algorithm) which can interact with an environment based on RotorS Gazebo [20] simulator (see Fig. 5). The design of our reinforcement learning framework focuses on the two core levels of abstraction in reinforcement learning problems: the *agent* and the *environment*, providing a flexible and versatile interface for the development of new agents or environments. In the proposed RL framework, the interaction between the agent and the environment has been implemented following the guidelines provided by the OpenAI Gym [6] toolkit for reinforcement learning, using ROS for communication purposes.

In this work, we model the visual servoing task as a fully observable Markov Decision Process (MDP), where the final objective is to train the agent in order to find the policy  $\pi$  that maximizes the accumulated discounted reward  $R_t = \sum_{i=t}^T \gamma^{i-t} r(s_i, a_i)$ , given the state  $s_i \in \mathbb{R}^4$ , the action  $a_i \in \mathbb{R}^2$ , and a discount factor  $\gamma \in [0, 1]$ . The behavior of the agent is controlled by this policy which maps an observation to an action, and is usually evaluated by using the action-value function  $Q^\pi(s_t, a_t) = \mathbb{E}_\pi[R_t | s_t, a_t]$ , which computes the expected return taking action  $a_t$  at state  $s_t$  and then following policy  $\pi$  thereafter.

In the proposed RL-IBVS architecture, the agent implements the DDPG algorithm, which is a model-free, off-policy algorithm based on an actor-critic architecture. In this architecture, the *critic* neural network acts as a nonlinear function approximator for estimating the action-value function, while the *actor* neural network is in charge of learning the policy that maps a state into a continuous action. In the specific case of the RL-IBVS, the trained networks (*actor* and *critic*) that compose the DDPG model, consist of feed-forward neural networks with two hidden layers of 300 and 200 units. At test time, the only network utilized is the *actor* network, whose input consists of a 4-dimensional vector representing the state of the detected

**Fig. 5** Image-Based Visual Servoing through Reinforcement Learning (RL-IBVS) component. The dotted lines represent interactions between the components in training mode, while the continuous lines depict the interactions in test mode (e.g. real flights)



target in the image, with the output layer being composed of two units that provide the 2-dimensional action to be commanded to the Motion Controller of the UAV, in velocity control mode.

At each time step  $t$  the *agent* executes a continuous action  $a_t = (a_x, a_y)$  and receives an *observation* (next state) and a *reward* from the environment (see Fig. 5). Where  $a_x, a_y$  represent the linear velocities in the  $x$  and  $y$  direction commanded to the Motion Controller in global coordinates, generated by the deterministic policy  $\pi(s_t)$  computed by the *actor* network.

On the other hand, at each time step the *environment* receives the action generated by the *agent* and computes the next state  $s_t = (e_x, e_y, de_x, de_y)$ , and a scalar *reward*  $r_t$  (see Eq. 7). Where  $e_x, e_y$  represent the normalized error in position, measured in pixels, of the center of the detected target ROI with respect to the defined reference in the image (see Fig. 9), and  $de_x, de_y$  are the normalized derivatives of the error filtered using a low-pass filter. The normalization factor utilized includes the width and the height of the image for the  $x$  and  $y$  coordinates respectively. The reference in the image is computed by projecting into the image plane the four corners of the square circumscribed to the circumference that forms the cylinder cover (see points p0 to p3 in Fig. 3) at the desired position with respect to the UAV. Once the projected rectangle is obtained using the intrinsic camera parameters, its center is calculated, which will define the reference in the image. Since the hooks responsible for releasing the items are not aligned with the bottom camera axis, an offset is added to this reference center in the case of real flights to ensure the items will fall inside the bucket. It is worth remarking here that, although the 3D points of the object are projected considering the desired UAV altitude, our algorithm is independent of this altitude as only the center of the projected rectangle is utilized for computing the state.

The reward at each time step (see Eq. 7) is computed by considering the evolution of a *shaping* function [15] (see Eq. 6) in two consecutive time steps. This procedure permits to alleviate the temporal credit assignment problem while accelerating the learning process of the agent as a localized advice is provided to the agent in each transition via the *shaping* function [17, 27]. The latter has been carefully designed in this work for achieving a smooth desirable behavior of the UAV while approaching the target. For this purpose, as can be seen in Eq. 6, the behavior of the agent is penalized when the errors in the image, their velocities, or the commanded actions are high.

$$shaping_t = -\alpha\sqrt{e_x^2 + e_y^2} - \beta\sqrt{de_x^2 + de_y^2} - \epsilon\sqrt{a_x^2 + a_y^2} \quad (6)$$

$$r_t = shaping_t - shaping_{t-1} \quad (7)$$

where  $r_t$  is the reward computed at time step  $t$ , and  $\alpha, \beta, \epsilon$  are the gains that control the penalization in the error in position, velocity and action terms respectively. These gains have been empirically obtained, being  $\alpha = 100, \beta = 10$  and  $\epsilon = 1$ .

In order to remove large displacements of the detected ROI in the image plane due to high roll and pitch commands of the UAV, an image stabilization technique is utilized taking into account the roll and pitch information provided by the IMU. In this direction, the points of the actual ROI are transformed into a stabilized image frame by using Eq. 8. This procedure has proven to be crucial in order to obtain a proper behavior of the RL-IBVS.

$$p' = K R K^{-1} p \quad (8)$$

where  $p$  represents a point in the actual image plane,  $p'$  is the point in the stabilized image plane,  $K$  is the projection matrix and  $R$  is the rotation matrix which encapsulates

---

the roll and pitch information in the aerial robot frame of reference expressed in the camera frame.

Additionally, in order to provide invariance to the UAV's yaw angle during the mission, a 2D rotation based on the actual yaw angle is applied to the state before being introduced to the actor network (see Eq. 9).

$$s' = \begin{bmatrix} \cos(\psi) & \sin(\psi) \\ -\sin(\psi) & \cos(\psi) \end{bmatrix} s \quad (9)$$

where  $s'$  is the transformed state introduced to the *actor* network, and  $\psi$  represents the UAV's yaw angle.

The implementation details for training and evaluating the RL-IBVS are provided in Section 5.3, where in all the experiments the frequency of the RL-IBVS has been set to 20 Hz.

## 5 Experiments and Results

The aim of this section is to present the experiments that have been conducted in order to evaluate the different components that integrate the proposed system. To this aim, five main experiments have been designed. Two of these experiments are conducted in order to train and evaluate the proposed learning-based algorithms for performing object recognition and IBVS tasks. The remaining three experiments are focused on the evaluation of the proposed system in SAR scenarios. For this purpose, the first of these experiments is aimed towards the evaluation of the whole system in a SAR simulated scenario. Subsequently, a second real flight experiment is designed in order to evaluate separately the rescue strategy on a moving target, and finally, the last experiment is conducted in order to validate the whole proposed system in a real flight SAR mission inside a cluttered indoor scenario. A video demonstration of the reported experiments and results is provided with this manuscript in: <https://vimeo.com/235929544>.

### 5.1 Experimental Setup

All the developed algorithms have been integrated into the Aerostack architecture and implemented in C++ and Python, using ROS as the communication middleware. Deep learning models for object classification have been trained using Keras<sup>2</sup> library on a 2.6 GHz CPU Intel Core i7-6700HQ, whereas models utilized for reinforcement learning purposes have been trained using TensorFlow<sup>3</sup> on a GPU Nvidia GeForce GTX 970. Regarding the simulated flight experiments, the proposed setup uses RotorS Gazebo for evaluating the IBVS algorithms, and on the other hand,

the PX4 Software-In-The-Loop is integrated with Gazebo in order to provide a realistic evaluation of the whole system for SAR missions. Regarding the real flight experiments, two indoor scenarios have been designed. The scenario for the first real flight experiment consists of a 3 m × 4 m area conceived for evaluating the rescue strategy, in which an OptiTrack motion capture system has been utilized for recording the ground truth data relative to the UAV and the moving target. The second scenario covers a 11 m × 7 m area used for evaluating the whole system in a SAR mission. In all the presented experiments, no tethers or external power supplies were utilized. In addition, in all real flight experiments, the UAV was carrying preloaded items of 100 g each, which substantially increased the complexity of the missions.

### 5.2 Training and Evaluation of the Supervised Learning Classifiers for Object Recognition

#### 5.2.1 Image Dataset

In the SAR missions proposed for evaluating our system, the targets consist of cylindrical buckets of different colors positioned at random locations in the scenario. Currently, there are no publicly available datasets containing images of cylindrical buckets that could be used for training a deep learning classifier. For this purpose, a custom dataset has been created containing images from *background* (non bucket) class and *bucket* class. After a large process of data acquisition, a total of 875 images for the *bucket* class and 2250 images for the *background* class were collected.

From this original dataset, several data augmentation techniques have been applied in order to increase the number of images used for training and evaluating the classifier and to prevent overfitting problems. The data augmentation process consisted in the application of three main techniques: *random cropping*, *horizontal flipping*, and *noise addition*. The random cropping strategy was similar to the one presented in [26], obtaining four cropped images per original image, by selecting a random offset starting from each of the four corners in the original image. The horizontal flipping strategy consisted in mirroring the original image from left to right direction, which allowed doubling the number of images. Finally, the last data augmentation technique was based on adding a Gaussian noise to the original image with zero mean and a standard deviation of 10 pixels. From the 875 images of bucket class, we performed data augmentation over 375 of these images obtaining a total of 4500 images. The remaining 500 images were added to the augmented dataset, providing a total amount of 5000 images for bucket class. From the 2250 images belonging to background class, data augmentation was conducted over 550 of these images, obtaining a total

---

<sup>2</sup><https://faroit.github.io/keras-docs/1.2.2/>

<sup>3</sup><https://www.tensorflow.org/>

**Fig. 6** Examples of images used for training the supervised learning classifiers for *bucket/background* classification. **a** Bucket training examples. **b** Background (non bucket) training examples. Images of office furniture and different types of floors have been used as examples of the background class



of 3300 images. The remaining 1700 images were added to the augmented set, providing a total amount of 5000 images for background class.

Subsequently, in order to train and evaluate the supervised learning classifiers considered in this work, the final dataset consisting of 10000 images (see Fig. 6 for an example), was divided into *train* (70%), *validation* (15%) and *test* (15%) sets.

### 5.2.2 Training Methodology and Evaluation

In order to conduct a reliable comparison between the different supervised learning classifiers considered in this work, a similar methodology as the one presented in [8] has been applied: First, a total of 6 evaluation tests have been defined, where the *train*, *validation* and *test* sets were randomly picked for each evaluation test. In each of the 6 evaluation tests, a *5-fold* cross-validation procedure is performed in the case of the L2R-LR and L-SVM classifiers in order to select the optimal regularization parameters of the corresponding classifier (see Eqs. 2 and 3 in Section 4.3.2). The range of values considered for the regularization parameter cover a wide range starting from  $10^{-4}$  to  $10^2$  with a step of  $2 \times 10^p$ , where  $p$  ranges

from  $-4$  to  $1$  with a step of  $1$ . In the case of the CNN classifier, a cross-validation procedure over the parameters of this model can be intractable. Thus, in this case we have considered 8 CNN configurations (see Table 2) varying the number of *feature maps* in each convolutional layer, the size of the convolution filters, and the number of channels of the images introduced to the network (i.e. grayscale or RGB images).

All the CNN configurations have been trained using the Keras library with TensorFlow as backend. As explained in Section 4.3.2, the CNN models consisted of 7 layers: 2 convolutional layers, 2 max pooling layers and 3 fully-connected with one hidden layer of 256 units, using ReLU activation function [36] and *dropout* regularization [47] in each layer except the final one, in which a *softmax* activation function was used. The selected loss function was based on the categorical cross-entropy loss, using Adam optimizer [24] for its minimization during a training process of 60 epochs as maximum, and a mini-batch size of 128 images. In addition, the *early stopping* technique was applied for regularization purposes, taking into account the validation loss with a patience of 5 *epochs*.

The results obtained after conducting the 6 evaluation tests on each classifier are shown in Fig. 7a and b and

**Table 2** CNNs configurations evaluated

CNN conf.	Input dim.	#Feature Maps		Filter size		#param
		Conv1	Conv2	Conv1	Conv2	
CNN1	$56 \times 72 \times 1$	32	64	$3 \times 3$	$3 \times 3$	3,165,314
CNN2	$56 \times 72 \times 1$	32	64	$5 \times 5$	$3 \times 3$	3,165,826
CNN3	$56 \times 72 \times 1$	32	64	$5 \times 5$	$5 \times 5$	2,756,226
CNN4	$56 \times 72 \times 1$	16	32	$5 \times 5$	$5 \times 5$	1,365,698
CNN5	$56 \times 72 \times 1$	8	16	$5 \times 5$	$5 \times 5$	680,034
CNN6	$56 \times 72 \times 1$	4	8	$5 \times 5$	$5 \times 5$	339,602
CNN7	$56 \times 72 \times 3$	32	64	$3 \times 3 \times 3$	$3 \times 3$	3,165,890
CNN8	$56 \times 72 \times 3$	32	64	$5 \times 5 \times 3$	$5 \times 5$	2,757,826




Fig. 7 Results obtained during the comparison of the supervised learning classifiers on 6 different evaluation tests (best seen in color). a)  $F_1$  score obtained in the training sets. b)  $F_1$  score obtained in the test sets

summarized in Tables 3 and 4, where C1 and C2 stand for Class 1 (*bucket class*) and Class 2 (*background class*) respectively. In these tables, the average values over the 6 tests are presented. In Table 3, the high scores obtained in most of the CNN configurations for the training set can be noticed. More relevant are the results presented in Table 4 for the test sets, where all the CNN configurations obtained a significantly higher  $F_1$  score than the one obtained by the L2R-LR and L-SVM classifiers. This fact can be confirmed in Fig. 7b, where in all the evaluation tests the CNN-based classifiers obtained a higher performance than the L2R-LR and L-SVM. As shown in Table 4, the number of feature maps was confirmed as being a crucial parameter for the CNN-based classifiers, where CNN6 obtained a significantly lower performance as compared to the rest of CNN configurations. Another important result extracted from Fig. 7b and Table 4 is that CNN configurations with a convolution filter of size  $5 \times 5$  provided higher performance than the analogous configurations with a filter size of  $3 \times 3$  (i.e. CNN3 vs CNN1, and CNN8 vs CNN7). However, these CNN configurations (CNN3 and CNN8) require the highest processing times per image and are quite demanding in

terms of the number of parameters, which will also lead to a higher memory consumption. All the results presented in Tables 3 and 4 have been obtained with a CPU Intel Core i7-6700HQ except the test time presented in Table 4, which has been measured with the UAV's onboard computer, and refers to the processing time per image, averaged over 128 images, taking into account the feature extraction plus the inference time of the classifier. Thus, in the case of the L2R-LR and the L-SVM, the test time encompasses the HOG feature extraction process plus the inference time of the classifier.

The execution of SAR missions using a fully-autonomous aerial robot implies running simultaneously multiple processes that implement the functionalities of each component of the system. Traditionally, UAVs are equipped with small and light computers with limited computational capabilities. In order to deal with these computational constraints, it is required to select the appropriate model in order to find a proper trade-off between performance and computational cost. For this purpose, in this work, we consider the selection of the most appropriate classifier based on two main variables: the performance of the

Table 3 Average training results obtained for the 6 evaluation tests of Fig. 7a

Classifier	Precision		Recall		F <sub>1</sub> score		Train time (s/epoch)
	C1	C2	C1	C2	C1	C2	
CNN1	0.998	1.00	1.00	0.998	0.999	0.999	30
CNN2	0.999	1.00	1.00	0.999	1.00	1.00	31
CNN3	1.00	1.00	1.00	1.00	1.00	1.00	44
CNN4	0.997	1.00	1.00	0.997	0.998	0.998	20
CNN5	0.996	1.00	1.00	0.996	0.998	0.998	10
CNN6	0.987	0.997	0.997	0.987	0.992	0.992	7
CNN7	0.996	1.00	1.00	0.996	0.998	0.998	33
CNN8	0.998	1.00	1.00	0.998	0.999	0.999	46
L2R-LR	0.992	0.992	0.992	0.992	0.992	0.992	–
L-SVM	0.989	0.991	0.991	0.989	0.990	0.990	–

**Table 4** Average test results obtained for the 6 evaluation tests of Fig. 7b

Classifier	Precision		Recall		F <sub>1</sub> score		Test time (ms/image)
	C1	C2	C1	C2	C1	C2	
CNN1	0.973	0.996	0.996	0.973	0.985	0.984	3.830
CNN2	0.978	0.993	0.993	0.977	0.985	0.985	3.956
CNN3	0.985	0.993	0.993	0.985	0.989	0.989	5.136
CNN4	0.973	0.998	0.998	0.972	0.985	0.985	3.027
CNN5	0.978	0.996	0.996	0.978	0.987	0.987	2.075
CNN6	0.963	0.989	0.989	0.962	0.976	0.975	1.921
CNN7	0.974	0.998	0.998	0.973	0.986	0.985	4.123
CNN8	0.980	0.996	0.996	0.980	0.988	0.988	5.830
L2R-LR	0.958	0.965	0.965	0.957	0.961	0.961	0.984
L-SVM	0.966	0.964	0.964	0.966	0.965	0.965	4.187

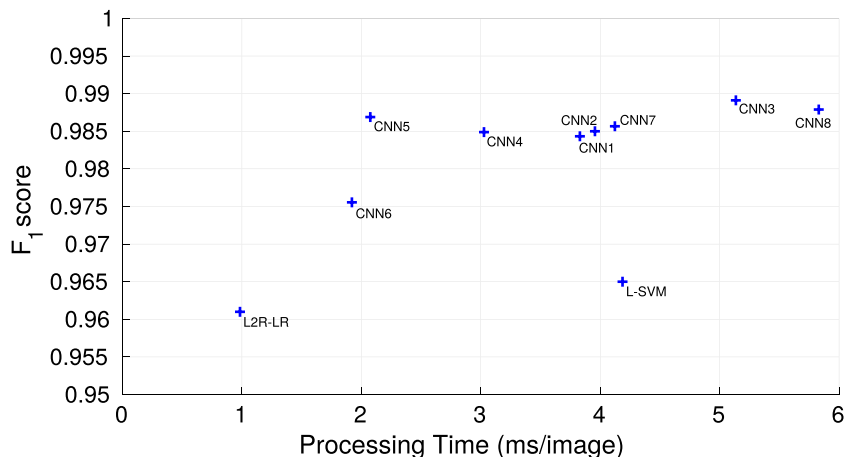
classifier and the processing time. Thus, we take into consideration the average  $F_1$  score and the test time presented in Table 4, which are plotted in Fig. 8 for a better visualization. Taking into account the results presented in Fig. 8 and considering that the best possible classifier is the one lying on point  $[0, 1]$  (i.e. minimizing the processing time while maximizing the  $F_1$  score), we have selected CNN5 as the most appropriate classifier for our purposes. This configuration provides the lower Euclidean distance (in normalized coordinates) to the desired point.

### 5.3 Training and Evaluation of the Reinforcement Learning agent for Image-Based Visual Servoing

#### 5.3.1 RL-IBVS Training Methodology


For training the agent in order to perform IBVS tasks, we use the RL-IBVS component (see Section 4.4.2) in training mode to command an AscTec Hummingbird quadrotor in the Gazebo environment presented in Fig. 9.

**Fig. 8** Average  $F_1$  score and processing time of the different supervised learning classifiers from the results presented in Table 4. The processing time has been measured on the Intel NUC6i5SYK onboard computer



In the RL-IBVS training mode, the *environment* is designed in an episodic RL setting, where the agent's experience is divided into a series of episodes, each one composed of several training steps. In each training step, the agent takes an action with added noise according to an Ornstein-Uhlenbeck distribution, and receives an *observation* and a *reward* from the environment (see Fig. 5). The current 4-dimensional observation is fed into the *actor* network, which generates a continuous 2-dimensional action in the range  $[-0.5, 0.5]$  m/s.

For stabilizing and accelerating the training process, the  $e_x, e_y$  variables from the state are measured taking into account the difference between the center of the ROI obtained by the projection into the image plane of the known 3D points of the target (see the cyan rectangle in Fig. 9) with respect to the reference defined in the image plane (see the blue rectangle in Fig. 9). During experimentation, we have found that this procedure is critical for allowing the convergence of the training process, as it removes large displacements of the detection of the target in the



**Fig. 9** RotorS Gazebo simulation environment used for training the RL-IBVS and image captured from the bottom camera of the UAV (best seen in color). The desired reference is depicted in blue color, while the detection of the target at the current time step is shown in green color. The cyan rectangle is obtained by projecting into the image plane the four corners of a square circumscribed to the circumference that forms the cylinder cover (see points  $p_0$  to  $p_3$  in Fig. 3)

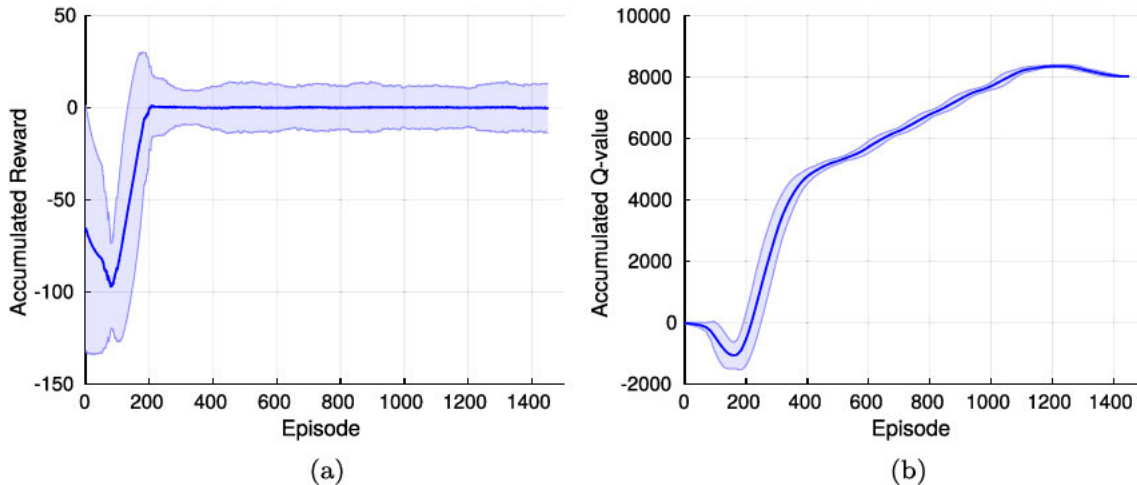
image plane due to sudden roll and pitch movements of the UAV.

At the beginning of each episode, the UAV is placed at a random position in a  $2\text{ m} \times 2\text{ m}$  area of the scenario at a constant altitude of 1.2 m. The episode is considered as finished when the maximum number of training steps per episode is reached or when the agent reaches a terminal state. In the specific case of the RL-IBVS presented in this work, the maximum number of steps per episode has been empirically set to 400 steps. A terminal state is reached when the target is out of the boundaries defined in the image captured by the bottom camera of the UAV (see the magenta rectangle in Fig. 9). When this situation occurs, the agent is penalized with a negative reward of -100 and a new episode

is started. In the rest of situations, the reward is computed using Eq. 7.

Using the previous setup, the training process of the *actor* and *critic* neural networks took approximately 7.5 hours, in which the Adam optimizer [24] was utilized with a base learning rate of  $10^{-4}$  for the *actor* and  $10^{-3}$  for the *critic*. As explained in Section 4.4.2, the architecture of the neural networks has been empirically obtained, using two hidden layers of 300 and 200 units. The output layer activation function of the *actor* is a *tanh* for providing a continuous linear velocity command, bounded to  $[-0.5, 0.5]$ . The rest of the hyperparameters are the same as the ones presented in [29]. The results obtained during the training process are illustrated in Fig. 10. In this figure, it is important to notice the different time instants in which the reward and the action-value function stabilize. The former reaches a stable value around zero from episode 200, while the action-value function continues increasing until episode 1250. Based on these results, the learned weights adopted for the neural networks of the DDPG (*actor* and *critic*) are the ones obtained at episode 1400, whose results have been additionally confirmed in a visual manner, and can be reviewed in the video demonstration.

It should be noted that although this training procedure is performed in simulation using the ground truth points of the target projected into the image plane, the realistic dynamics provided by the RotorS Gazebo simulator together with the appropriate policy learned by the actor network allow an almost direct transition to previously unseen simulated and real scenarios where the location of the target is unknown. In these scenarios, where no ground truth data relative to the position of the target is available, the stabilized detection of the target in the image plane is utilized for computing the state.



**Fig. 10** Training curves obtained during the training process of the RL-IBVS agent. Each curve is computed taking into account a moving average of 100 episodes. **a** Each point depicts the accumulated reward

at the specific episode. **b** Each point in the curve represents the accumulated action-value  $Q$  predicted by the *critic* network at the specific episode



### 5.3.2 RL-IBVS Evaluation

In this section, the experiments conducted in order to evaluate the proposed RL-IBVS algorithm using the trained agent are described in detail. For the sake of comparison, and in order to perform a thorough evaluation of the RL-IBVS system, the IBVS controller proposed in [11, 31] (referred to as *classic* IBVS) has been implemented using the Visual Servoing Platform (ViSP) library [32], and is utilized in this work as the baseline for evaluating our RL-IBVS. The *classic* IBVS is based on the well-known visual servo control law which establishes that  $\mathbf{v} = -\lambda \mathbf{L}_s^+ \mathbf{e}$ , where  $\mathbf{v}$  represents the spatial velocity of the camera,  $\mathbf{L}_s$  is the *interaction matrix*,  $\mathbf{e} = \mathbf{s} - \mathbf{s}^*$  is the error computed as the difference between the actual and desired features in the image plane, and  $\lambda$  is a positive control gain and has been empirically obtained as  $\lambda = 0.75$ . The estimation of the interaction matrix is obtained as proposed in [31], where  $\widehat{\mathbf{L}}_s^+ = \frac{1}{2}(\mathbf{L}_s + \mathbf{L}_{s^*})^+$ , being  $\mathbf{L}_{s^*}$  the interaction

matrix at the desired feature location, and  $(\cdot)^+$  represents the pseudoinverse operation as defined in [11].

In order to obtain a complete evaluation of the *classic* IBVS and the RL-IBVS, a simulated experiment using RotorS Gazebo with an AscTec Hummingbird quadrotor has been designed with the aim of following a moving target. In this experiment, the target performs a rhomboidal trajectory (see Figs. 11d and 12d), moving at a constant speed of 0.5 m/s along the edges of the rhomboid and stopping during several simulation steps in its corners. Using this setup, a complete evaluation of the IBVS approaches can be conducted as it involves the interaction with a static and a moving target in the same experiment.

The results obtained after the execution of the proposed experiment are shown in Figs. 11 and 12 for the *classic* IBVS and the RL-IBVS respectively. As can be seen in Figs. 11a and 12a, both controllers produced a smooth response in the error signal. A similar result can be observed in Figs. 11c and 12c for the commanded actions, with the

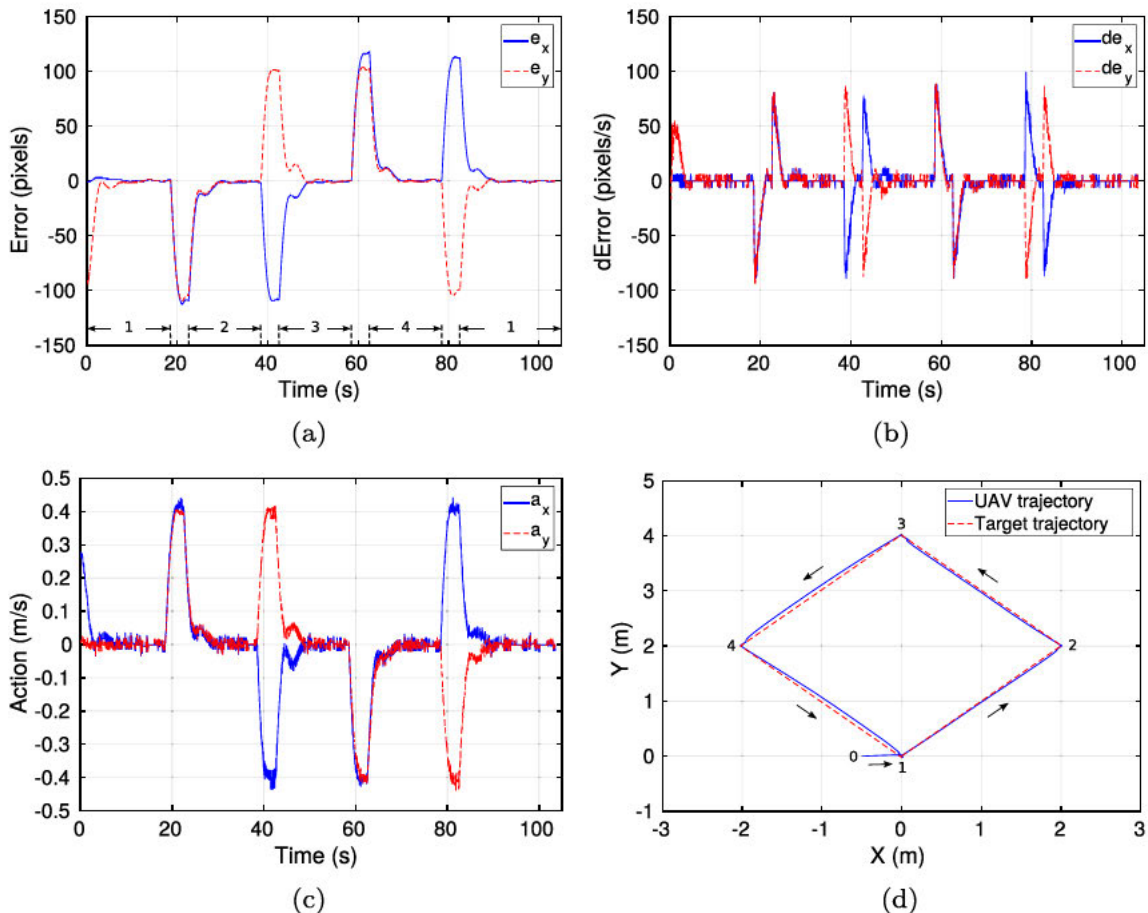


Fig. 11 Simulation results obtained for the classic IBVS approach while following a moving target executing a rhomboidal trajectory at a maximum speed of 0.5 m/s. The target stops for a few time steps on each corner of the rhomboid (1 to 4). a Error in the image with respect

to the reference. b Derivatives of the error in the image with respect to the reference. c Actions generated by the classic IBVS controller referred in the world frame of reference. d Trajectories followed by the UAV and the target

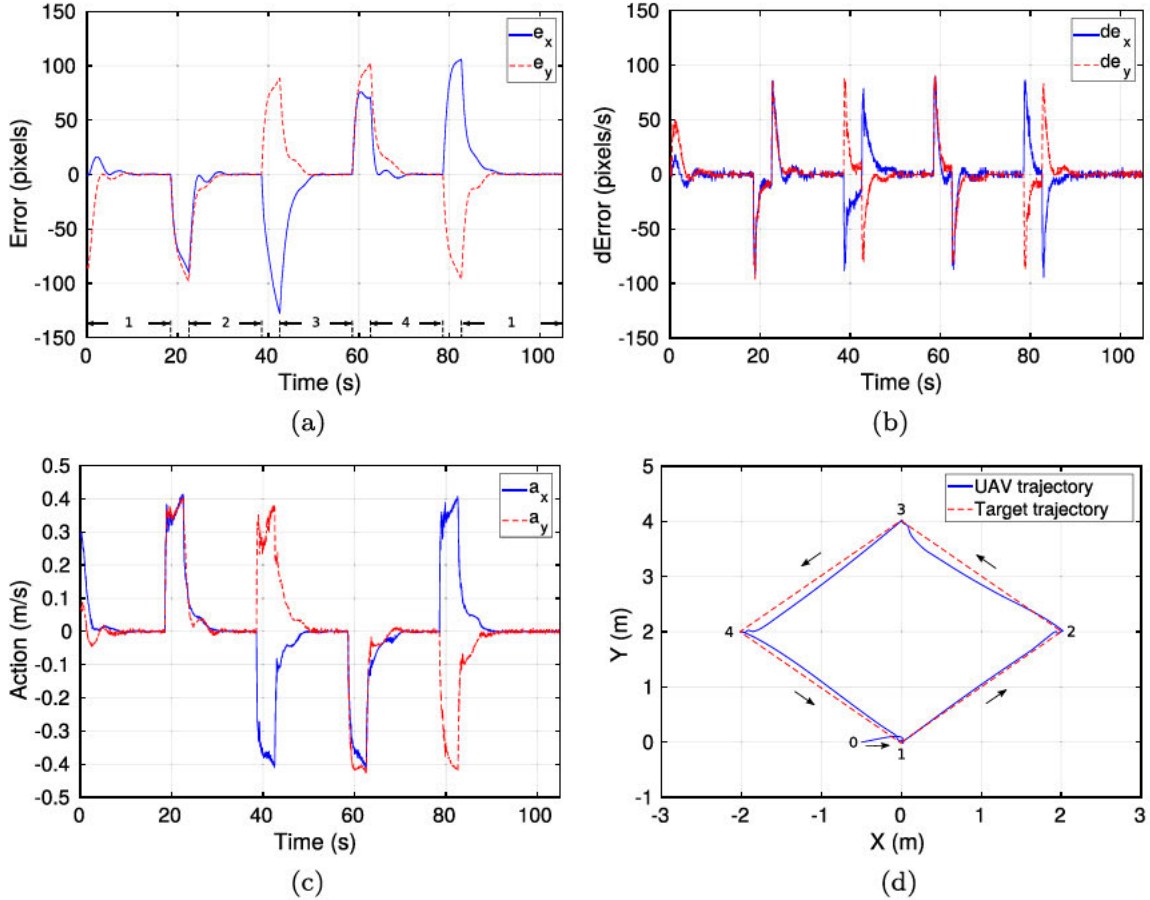


Fig. 12 Simulation results obtained for the proposed RL-IBVS approach while following a moving target executing a rhomboidal trajectory at a maximum speed of 0.5 m/s. The target stops for a few time steps on each corner of the rhomboid (1 to 4). a Error in the image with

respect to the reference. b Derivatives of the error in the image with respect to the reference. c Actions generated by the RL-IBVS controller referred in the world frame of reference. d Trajectories followed by the UAV and the target

RL-IBVS producing a slightly lower oscillatory response. The trajectories followed by the UAV and the moving target are depicted in Figs. 11d and 12d, where the classic IBVS produced a trajectory with a very small deviation with respect to the trajectory of the moving target. Finally, Table 5 summarizes the results obtained for both IBVS approaches in the case of the error signal. The metrics presented in this table are the mean and standard deviation of the error signal, the percentage overshoot (PO) and the steady-state error ( $e(ss)$ ) which we consider relevant for the rescue operations presented in this paper where the precision acquires an utmost importance. The last two metrics (PO and  $e(ss)$ ) are measured averaging the response of the controllers to the 5 steps produced when the UAV approaches the moving target stopped at each of the four

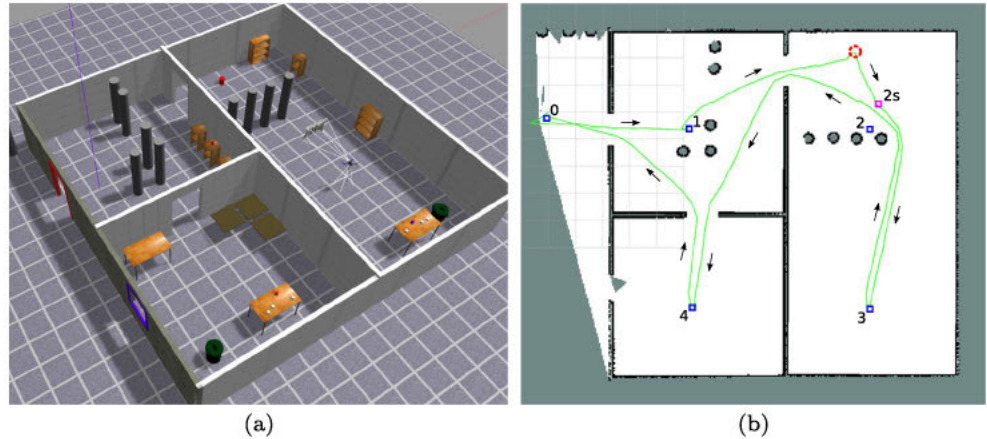
corners of the rhomboid (see steps on corners 1 to 4 in Figs. 11a and 12a). In Table 5 it can be noticed the appropriate behavior of both controllers, highlighting the very high precision of the RL-IBVS controller with an average error of less than one pixel and a steady-state error of less than 0.5 pixels in both  $x$  and  $y$  directions. Furthermore, the standard deviation of the error signal in the case of the RL-IBVS is about 6 pixels less than the *classic* IBVS, revealing the faster response of the former while following the target. According to these results, the RL-IBVS is adopted in order to conduct IBVS tasks in real flight experiments.

It should be remarked that the results obtained for the *classic* IBVS approach are the result of a long tuning process in which the corresponding gains of this controller

Table 5 Average errors measured in pixels and percentage overshoot obtained for the simulation flight experiment of Figs. 11 and 12

IBVS approach	$e_x$	$e_y$	$e_x(ss)$	$e_y(ss)$	$PO(e_x)$	$PO(e_y)$
Classic	$0.5 \pm 41.3$	$-1.6 \pm 39.3$	0.32	0.61	1.74%	1.35%
RL	$-0.4 \pm 34.6$	$-0.9 \pm 34.2$	0.22	0.12	3.46%	1.46%

**Fig. 13** Fully-autonomous simulation flight experiment for a SAR mission in a cluttered indoor scenario (best seen in color). The target is represented by a red bucket. **a** Gazebo simulation environment. **b** 2D occupancy grid map and trajectory followed by the UAV (green color). Blue squares depict the mission points (1 to 4) automatically generated by the GMP. The magenta square (2s) represents the safety mission point generated in substitution of mission point 2



were adjusted by a trial and error process. Conversely, the results presented in Fig. 12 for the RL-IBVS have been obtained by directly using the trained *actor* network without any further tuning.

#### 5.4 Search and Rescue

This section presents the experiments conducted and the results obtained during the execution of simulated and real flight experiments in SAR scenarios. A fully-autonomous SAR mission in a Gazebo simulated scenario (see Fig. 13) is presented in Section 5.4.1. Real flights showing the performance of the system in real SAR scenarios are introduced in Sections 5.4.2 and 5.4.3. For this purpose, two main experiments have been designed (see Fig. 14), one for evaluating the autonomous rescue of a moving target, and the second experiment which aims to evaluate the whole proposed system in a SAR mission inside a cluttered indoor scenario.

##### 5.4.1 Autonomous SAR in Cluttered Simulated Scenario

This section presents the simulation experiment that has been designed in order to evaluate the coordination between the components of the whole aerial robotic system and its capabilities to successfully accomplish SAR missions in a

cluttered indoor scenario. For the evaluation of the whole system, several simulated environments using Gazebo simulator (see Fig. 13a) have been designed based on the scenario proposed for the IMAV 2016 competition. Inspired by the use case of IMAV 2016, different layouts of indoors rooms have been created. In this experiment, the objective of the UAV is to explore the *a priori* unknown indoor scenario searching for a predefined target (red bucket). Once the target has been recognized, the UAV has to interact with the former to release an item inside of it for rescue purposes. Once the interaction is finished, the UAV has to continue exploring the scenario in order to locate more possible targets. As shown in Fig. 13b, the UAV is able to explore the entire indoor scenario, readapting the exploration mission when the target is recognized in order to perform the release maneuver. Another remarkable behavior that can be noticed in Fig. 13b is the safety point (see point 2s in Fig. 13b) generated by the AMP in order to avoid the mission point number 2 which is within the boundaries of an obstacle, taking into account the inflation radius of the obstacles.

##### 5.4.2 Autonomous Rescue of a Moving Target

This experiment has been designed for evaluating the proposed rescue strategy in real flight conditions (see

**Fig. 14** Indoor scenarios designed for real flight experiments. **a** Real flight experiment conducted for evaluating the autonomous rescue of a moving target. **b** Real flight experiment conducted for evaluating the complete proposed system while performing a SAR mission in a cluttered scenario


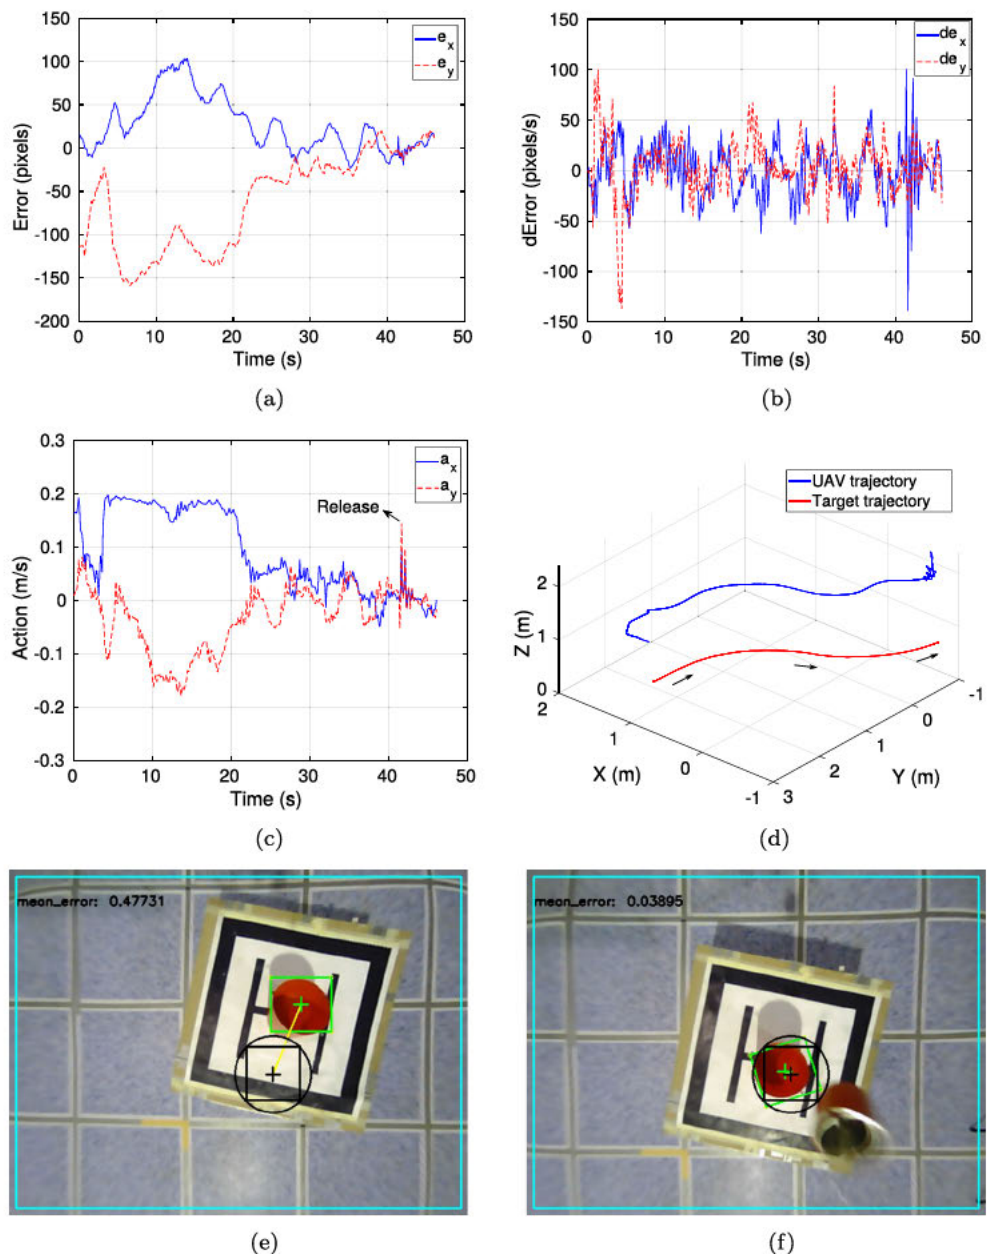


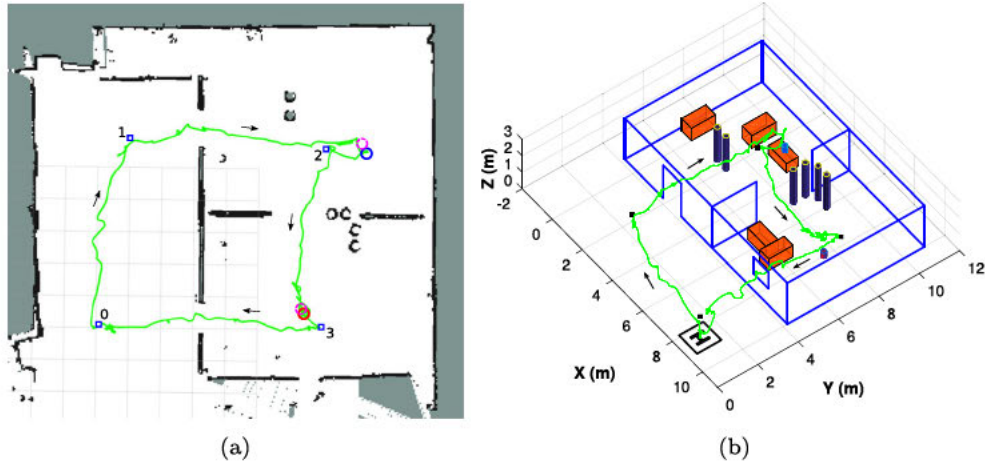
Fig. 14a), in which the RL-IBVS is the main component responsible for the rescue operation.

In order to obtain a detailed evaluation of the system, the experiment integrates a remotely controlled moving target which describes a random trajectory commanded by a human operator. The goal of the UAV is to autonomously track the moving target until a secure position is reached for performing the release maneuver of a preloaded item for rescue purposes. The results obtained during the execution of the experiment are presented in Fig. 15, which shows how the error in the image and the control actions tend to zero at the end of the path of the moving target, where it starts decreasing the speed until it is stopped. In this time

instant (see time instant 41.6 s in Fig. 15a), the error in pixels of the detected object in the image decreases below a predefined threshold and the *target locked* action for releasing the item is sent to the AMP, which commands its execution. In Fig. 15c, the decremental tendency of the control actions can be clearly seen as soon as the UAV is approaching the target. Around time instant 41.6 s, the release operation is commanded. The release of a 100 g item causes a perturbation on the dynamics of the UAV, successfully handled by our system. It is also very important to notice that, despite the turbulences generated while flying very close to the target (empty bucket), the UAV is able to maintain a stable flight, being the control actions under 0.04

**Fig. 15** Real flight experiment for evaluating the rescue strategy on a moving target. **a** Error in the image w.r.t the reference. **b** Derivatives of the error in the image w.r.t the reference. **c** Action commands generated by the RL-IBVS. **d** 3D paths followed by the UAV and the moving target. **e**, **f** Errors computed by the RL-IBVS component in two different instants of the mission. **f** When the target is locked, a release operation is commanded. An OptiTrack mocap system is used for recording the positions of the UAV and the moving target





**Fig. 16** Fully-autonomous real flight experiment for a SAR mission in a cluttered indoor scenario (best seen in color). The targets are represented by two colored buckets. The trajectory followed by the UAV during the mission is depicted in green color. a 2D occupancy grid map

m/s when the UAV is exactly above the target. Finally, the 3D trajectories followed by the UAV and the moving target are depicted in Fig. 15d, which shows the stable tracking trajectory followed by the UAV on top of the moving target. This trajectory is only perturbed in height when the release operation is performed. This fact can be clearly seen around point  $x : -0.82$  m,  $y : -0.84$  m, when the UAV moves slightly up.

#### 5.4.3 Autonomous SAR in Cluttered Indoor Scenario

In this section, a real flight SAR experiment for evaluating the whole system proposed in this work is presented. In the designed experiment the UAV is required to autonomously explore a cluttered indoor scenario with two possible entries: a  $1.2$  m  $\times$   $2.0$  m door or a  $1$  m  $\times$   $1$  m window. The indoor scenario is composed of several obstacles, mainly tables, chairs and columns, whose location is completely unknown (see Fig. 14b). Based on the use case of the IMAV 2016, the targets are represented by two colored buckets. For increasing the difficulty of the experiment, one of these targets is located on top of a table (blue bucket) while the other (red bucket) is positioned on the floor. The objective of the UAV is to autonomously explore the indoor environment, localize the targets, and interact with them by means of performing a release maneuver of the corresponding item. The items, both red and blue, are preloaded before takeoff and held in the hooks of the UAV. Once the interaction with the targets is finished, the UAV has to return to the initial takeoff point. Fig. 16 shows the 2D map generated (see Fig. 16a) and the 3D trajectory followed by the UAV (see Fig. 16b) when the SAR mission is completed. As can be observed in Fig. 16, the system is capable of adapting the current mission when the first


and trajectory followed by the UAV. Dotted pink circles represent the estimated target position provided by the Model-Based Object Recognizer. b 3D outline of the indoor scenario. Obstacles are represented by cuboids and columns. Squares depict the commanded mission points

target (blue bucket) is recognized while going from mission point 1 to 2. For this purpose, the mission planning system acts by queuing up the current exploration mission and adopting the rescue operation of the recognized target as the current mission to accomplish. Once the UAV finishes the rescue operation on the first target, the previous exploration mission is recovered in its corresponding state. The same behavior can be noticed when the UAV is commanded from mission point 2 to 3 and the target 2 (red bucket) is recognized.

The estimated positions of the targets computed by the Model-Based Object Recognizer component are depicted in Fig. 16a (see dotted circles in Fig. 16a). The ground truth positions of both targets were measured by a human operator, being  $[x : 7.84$  m,  $y : 7.5$  m,  $z : 0.0$  m] for *bucket A* (red bucket) and  $[x : 2.8$  m,  $y : 9.2$  m,  $z : 0.54$  m] for *bucket B* (blue bucket). The estimated positions were  $[x : 7.71$  m,  $y : 7.4$  m,  $z : 0.068$  m] for *bucket A* and  $[x : 2.54$  m,  $y : 9.14$  m,  $z : 0.596$  m] for *bucket B*. Thus, the average error in each coordinate for both targets is  $[x : 0.19$  m,  $y : 0.08$  m,  $z : 0.062$  m] what reveals the acceptable position estimation, taking into account the accumulated errors that might occur while measuring the ground truth positions, the ones introduced by the estimation of the UAV position, and the errors produced while computing the detection of the object in the image plane.

In Figs. 17 and 18 the recognition of the targets in different time instants during the execution of the real flight SAR mission is depicted. It is important to remark the capability of our system to automatically switch between different recognition modes. Based on this, once the object is considered as recognized using the front camera image, the system stops processing the images coming from that device and starts processing the images coming from the

**Fig. 17** Object recognition results for target 1 (best seen in color). a, b, c, d Recognition of the blue bucket (bucket B) from the front camera. e, f Detection of the bucket from the bottom camera. f) Detection of the bucket in the moment when the release command has just been triggered. The positions indicated in the figure refer to the computed relative pose of the target with respect to the front camera using a PnP algorithm. The output probabilities from the CNN classifier are indicated together with the object category



bottom camera. Subsequently, when the release maneuver for the rescue operation has finished, the recognition mode using the front camera is enabled again, stopping the processing with the bottom device.


## 6 Discussion

The development of a fully-autonomous aerial robot for SAR missions requires the implementation of versatile and efficient algorithms in a wide range of components ranging from control, mission planning, path planning, mapping, to object recognition, object following, etc. In order to successfully accomplish these missions, a precise

coordination between the different components is required. This coordination becomes more critical when it involves the interaction with the detected target, and even more in the presence of dynamic targets.

In this paper, we have focused our attention on this type of missions that have not been sufficiently addressed in the literature. In order to provide more flexible and versatile algorithms, here we extend our previous works [43, 44] by integrating learning-based algorithms for object recognition and object interaction. The experiments and results presented in Section 5, demonstrate the appropriate capabilities of our proposed system for the accomplishment of such high-level missions in a fully unsupervised manner (i.e. without human intervention).

**Fig. 18** Object recognition results for target 2 (best seen in color). a, b, c, d Recognition of the red bucket (bucket A) from the front camera. e, f Detection of the bucket from the bottom camera. f Detection of the bucket in the moment when the release command has just been triggered. The positions indicated in the figure refer to the computed relative pose of the target with respect to the front camera using a PnP algorithm. The output probabilities from the CNN classifier are indicated together with the object category



Regarding the object recognition task, several supervised learning classifiers have been evaluated in order to address the hard computational constraints imposed in the real-time operation of an aerial robotic platform. After a thorough evaluation, a 7-layer CNN model composed of 2 convolutional layers of 8 and 16 feature maps with convolution kernels of size  $5 \times 5$  showed the best balance in terms of accuracy ( $F_1$  score of 0.987) and processing time (2.07 ms/image) for the task of target/background classification presented in this work. This accuracy reflects the accurate detection of the object in the image plane (see Figs. 17 and 18), even in frames with a considerable amount of blur as shown in Figs. 17a, 17b, 18a and 18b. Remarkable results in terms of processing time were exhibited by the L2

Regularized Logistic Regression classifier which provided an average processing time per image of 0.98 ms/image. Despite providing an  $F_1$  score lower than other classifiers ( $F_1$  score of 0.961), we believe that this classifier can be considered as a suitable candidate in platforms with really hard computational constraints such as Micro Aerial Vehicles (MAVs).

With respect to the target interaction and following task for rescue operations, it is very important to remark that the proposed RL-IBVS has been trained only in simulation and with the dynamics of the UAV available in the RotorS Gazebo simulator (i.e. AscTec Hummingbird quadrotor). For the transition to real flights using our custom aerial quadrotor, no modifications have been made to the trained

---

model. This fact is even more relevant in the field of aerial robotics, where directly training a reinforcement learning algorithm in a real scenario is very complicated and to the authors' knowledge has not been yet addressed in the literature. Another important feature of the proposed RL-IBVS algorithm is its simplicity and versatility as the functionality of the RL-IBVS is independent of the object to be recognized, being the only required input the center of the ROI corresponding to the detected object in the image. The comparison results of the proposed RL-IBVS with respect to a classic IBVS approach shown in Figs. 11 and 12 and Table 5 demonstrate that the RL-IBVS can be utilized as an alternative to state-of-the-art IBVS approaches, which usually require a long tuning process of their parameters. The results presented in Table 5 reveal the outstanding precision of the RL-IBVS with steady-state errors lower than 0.5 pixels in both  $x$  and  $y$  directions. This precision is critical in the presented SAR missions where a high precision rescue operation for delivering preloaded items is required. Moreover, as shown in Fig. 10, the convergence in the training process of the RL-IBVS is achieved relatively fast (7.5 hours on a GPU Nvidia GeForce GTX 970) where the Q-value starts to settle around episode 1250.

This fast convergence was mainly obtained owing to the training strategy presented in this work, consisting in using as the detected ROI the known 3D points of the target projected into the image plane, which provides a more stable state of the target, removing the effects of sudden roll and pitch movements of the UAV. The main limitation of the proposed RL-based algorithm is its dependency on the stability of the detected object in the image, which is efficiently solved in this work at test time by an image stabilization technique. The inclusion of this stabilization technique in the training process of the agent will be further studied in future works.

Using the aforementioned learning-based capabilities, the proposed system has been evaluated in complex SAR real flight experiments involving the rescue operation on static and moving targets. The latter scenario was presented in Section 5.4.2, where the proposed system was able to efficiently deliver a preloaded 100 g item on a remotely controlled moving target. The efficient behavior of the RL-IBVS during the experiment is evidenced by the temporal evolution of the error signal shown in Fig. 15a, where the error decreases from more than 100 pixels in the first instants of the mission (e.g. time instant 14 s) to less than 7 pixels when the release operation is commanded (e.g. time instant 41.6 s). Furthermore, the stable behavior of the RL-IBVS can be derived from the low control actions (0.04 m/s) that the RL-IBVS is able to maintain when the UAV is exactly above the target, what permitted the precise release operation of the item.

A similar behavior of the RL-IBVS allowed the correct rescue operation on the static targets presented in the experiment of Section 5.4.3, which shows a fully-autonomous SAR mission in a cluttered indoor scenario. In this mission, the accurate detection of the object in the image plane represents a critical step for pose estimation purposes and the subsequent rescue operation on both targets. The pose estimation of the targets has been addressed in this paper by means of a PnP algorithm, which provides an acceptable pose estimation using monocular information. During the execution of the real flight SAR mission of Fig. 16, the average errors in the position estimation of both targets were [ $x$  : 0.19 m,  $y$  : 0.08 m,  $z$  : 0.062 m], highlighting the low errors obtained in the estimation of the  $y$  and  $z$  coordinates, both under 8 cm. The error in the  $x$  coordinate is slightly higher, which can be caused by sudden roll movements of the UAV while detecting the target. These errors are adequate for SAR missions and have low effect on our system as a fisheye-lens bottom-looking camera with a considerable field of view is utilized. This field of view allows the detection of the target within an area of 4 m  $\times$  4 m at the flight altitude of 1.2 m. Furthermore, the proposed RL-IBVS can handle quite big errors in the detection as shown in Fig. 12a.

Finally, it should be remarked that the success in the SAR missions presented in this work is mainly achieved due to the flexible mission planning system which allows the mission adaptation in real-time, switching from exploration to rescue mode and vice-versa when required. In addition, the learning-based techniques presented in this work can be easily adapted to other types of SAR missions as they are based on training models which can be retrained based on other requirements.

## 7 Conclusions and Future Work

In this paper, a fully-autonomous aerial robotic system for executing Search and Rescue (SAR) missions in cluttered indoor environments has been presented. The aerial robotic system developed in this work is based on the combination of a complete hardware configuration and a flexible system architecture which provide the appropriate capabilities for performing very high-level missions in a fully unsupervised manner. These capabilities include a dynamic mission planning system that allows mission re-planning in real-time, a reactive collision avoidance navigation system based on laser information, a complete multi-sensor fusion system for accurate pose estimation and altitude filtering, an accurate object recognizer component based on a Convolutional Neural Network (CNN) model, and a novel Image-Based Visual Servoing (IBVS) algorithm using deep



---

Reinforcement Learning (RL) for object interaction in rescue operations, among others.

Regarding the SAR problem, in this paper we have focused on object recognition and object interaction tasks, which we consider still as an open field of research in SAR missions and have been addressed in this work by means of learning-based techniques. For object recognition purposes, several supervised learning classifiers have been extensively evaluated for target/background segmentation. The final selected model consists of a 7-layer CNN which exhibits a good compromise between accuracy and computational cost, tackling the hard computational constraints of a real-time aerial robot. With respect to object interaction and following for rescue operations, a recent deep reinforcement learning algorithm, named Deep Deterministic Policy Gradients (DDPG), has been adapted in order to perform IBVS tasks. For training and evaluating the RL-IBVS algorithm, our own reinforcement learning framework has been developed. The proposed RL framework integrates a deep RL agent (e.g. DDPG) with a simulation environment for aerial robotic platforms (e.g. RotorS Gazebo simulator). After a thorough evaluation of the proposed RL-IBVS taking as baseline a classic IBVS controller, it has been demonstrated that reinforcement learning techniques can be efficiently trained and used for solving several tasks in SAR scenarios such as the object delivering maneuvers studied in this work.

The proposed system has been thoroughly evaluated by means of several simulated and real flight experiments. The RL-IBVS algorithm has been validated both on simulated and real scenarios with static and dynamic targets. Additionally, another set of experiments has been designed for validating the whole system on a complete SAR mission conducted in a cluttered indoor scenario, revealing its appropriate capabilities for the accomplishment of such high-level missions without the necessity of having a prior map of the scenario. One of the key capabilities of our system for the accomplishment of such missions is the versatile coordination between the mission planner and the perception components, providing high-level decision-making capabilities to the system, which is able to efficiently respond to different events that can occur during SAR missions (e.g. target recognized).

The development of the proposed system has motivated several future research directions that are summarized here. One of these directions is aimed towards the accomplishment of SAR missions using a swarm of aerial robots. We believe that an appropriate coordination of a team of UAVs can lead towards the execution of SAR missions in a more optimized manner. As it has been stated throughout this document, object interaction is one of our major concerns. For this reason, another future line of research will be focused on the exploration of

additional deep reinforcement learning techniques applied to object interaction tasks such as object release, object grasping, etc., which will increase the functionalities of the proposed system in SAR scenarios. Taking into account the outstanding results obtained in this document by using learning-based techniques, the extension of the capabilities of other components in the architecture (e.g. reactive navigation) using learning-based approaches will be considered.

**Acknowledgements** This work was supported by the Spanish Ministry of Science (Project DPI2014-60139-R). The LAL UPM and the MONCLOA Campus of International Excellence are also acknowledged for funding the predoctoral contract of the corresponding author.

## References

1. Abrahamsen, H.B.: A remotely piloted aircraft system in major incident management: concept and pilot, feasibility study. *BMC Emergency Medicine* **15**(1), 12 (2015)
2. Achtelik, M., Bachrach, A., He, R., Prentice, S., Roy, N.: Autonomous Navigation and Exploration of a Quadrotor Helicopter in Gps-Denied Indoor Environments. In: 1st Symposium on Indoor Flight, 2009. Citeseer (2009)
3. Bachrach, A., He, R., Roy, N.: Autonomous flight in unknown indoor environments. *International Journal of Micro Air Vehicles* **1**(4), 217–228 (2009)
4. Bavle, H., Sanchez-Lopez, J.L., Rodriguez-Ramos, A., Sampedro, C., Campoy, P.: A Flight Altitude Estimator for Multirotor Uavs in Dynamic and Unstructured Indoor Environments. In: 2017 International Conference on Unmanned Aircraft Systems (ICUAS), Pp. 1044–1051 (2017). <https://doi.org/10.1109/ICUAS.2017.7991467>
5. Bejiga, M.B., Zeggada, A., Nouffidj, A., Melgani, F.: A convolutional neural network approach for assisting avalanche search and rescue operations with uav imagery. *Remote Sens.* **9**(2), 100 (2017)
6. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym. arXiv:1606.01540 (2016)
7. Carrio, A., Sampedro, C., Rodriguez-Ramos, A., Campoy, P.: A review of deep learning methods and applications for unmanned aerial vehicles. *Journal of Sensors* 2017 (2017)
8. Caruana, R., Niculescu-Mizil, A.: An empirical comparison of supervised learning algorithms. In: Proceedings of the 23rd international conference on Machine learning, pp. 161–168. ACM (2006)
9. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* **2**, 27:1–27:27 (2011). Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
10. Chaumette, F.: Potential Problems of Stability and Convergence in Image-Based and Position-Based Visual Servoing. In: The Confluence of Vision and Control, pp. 66–78. Springer (1998)
11. Chaumette, F., Hutchinson, S.: Visual servo control. i. basic approaches. *IEEE Robotics Automation Magazine* **13**(4), 82–90 (2006). <https://doi.org/10.1109/MRA.2006.250573>
12. Chaumette, F., Malis, E.: 2 1/2 D Visual Servoing: a Possible Solution to Improve Image-Based and Position-Based Visual

- 
- Servoings. In: Proceedings of the IEEE International Conference on Robotics and Automation, 2000. ICRA'00, vol. 1, pp. 630–635 (2000)
13. Dalal, N., Triggs, B.: Histograms of Oriented Gradients for Human Detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005. CVPR 2005, vol. 1, pp. 886–893 (2005)
  14. Doherty, P., Rudol, P.: A Uav Search and Rescue Scenario with Human Body Detection and Geolocalization. In: Australian Conference on Artificial Intelligence, vol. 4830, pp. 1–13. Springer (2007)
  15. Dorigo, M., Colombetti, M.: Robot shaping: an experiment in behavior engineering. MIT press, Cambridge (1998)
  16. Erdos, D., Erdos, A., Watkins, S.E.: An experimental uav system for search and rescue challenge. IEEE Aerosp. Electron. Syst. Mag. **28**(5), 32–37 (2013)
  17. Erez, T., Smart, W.D.: What Does Shaping Mean for Computational Reinforcement Learning? In: 7Th IEEE International Conference on Development and Learning, 2008. ICDL 2008, pp. 215–219 (2008)
  18. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: a library for large linear classification. J. Mach. Learn. Res. **9**, 1871–1874 (2008)
  19. Fleck, M.: Usability of lightweight defibrillators for uav delivery. In: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 3056–3061. ACM (2016)
  20. Furrer, F., Burri, M., Achtelik, M., Siegwart, R.: Rotors-A Modular Gazebo Mav Simulator Framework. In: Robot Operating System (ROS), pp. 595–625. Springer (2016)
  21. Gatteschi, V., Lamberti, F., Paravati, G., Sanna, A., Demartini, C., Lisanti, A., Venezia, G.: New Frontiers of Delivery Services Using Drones: a Prototype System Exploiting a Quadcopter for Autonomous Drug Shipments. In: 2015 IEEE 39Th Annual, Computer Software and Applications Conference (COMPSAC), vol. 2, pp. 920–927 (2015)
  22. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 580–587 (2014)
  23. Grzonka, S., Grisetti, G., Burgard, W.: A fully autonomous indoor quadrotor. IEEE Trans. Robot. **28**(1), 90–100 (2012)
  24. Kingma, D., Ba, J.: Adam: A method for stochastic optimization arXiv:1412.6980 (2014)
  25. Kohlbrecher, S., Meyer, J., von Stryk, O., Klingauf, U.: A Flexible and Scalable Slam System with Full 3D Motion Estimation. In: Proceedings of the IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR) (2011)
  26. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet Classification with Deep Convolutional Neural Networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)
  27. Laud, A.D.: Theory and application of reward shaping in reinforcement learning. Tech rep (2004)
  28. Lee, A.X., Levine, S., Abbeel, P.: Learning visual servoing with deep features and fitted q-iteration arXiv:1703.11000 (2017)
  29. 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:1509.02971 (2015)
  30. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single Shot Multibox Detector. In: European Conference on Computer Vision, pp. 21–37. Springer (2016)
  31. Malis, E.: Improving Vision-Based Control Using Efficient Second-Order Minimization Techniques. In: Proceedings of the 2004 IEEE International Conference on Robotics and Automation, 2004. ICRA'04, vol. 2, pp. 1843–1848 (2004)
  32. Marchand, É., Spindler, F., Chaumette, F.: Visp for visual servoing: a generic software platform with a wide class of robot control skills. IEEE Robot. Autom. Mag. **12**(4), 40–52 (2005)
  33. Marder-Eppstein, E., Berger, E., Foote, T., Gerkey, B., Konolige, K.: The Office Marathon: Robust Navigation in an Indoor Office Environment. In: International Conference on Robotics and Automation (2010)
  34. Meier, L., Tanskanen, P., Fraundorfer, F., Pollefeys, M.: Pixhawk: a System for Autonomous Flight Using Onboard Computer Vision. In: 2011 IEEE International Conference on Robotics and Automation (ICRA), pp. 2992–2997 (2011)
  35. Moore, T., Stouch, D.: A generalized extended kalman filter implementation for the robot operating system. In: Proceedings of the 13th International Conference on Intelligent Autonomous Systems (IAS-13). Springer (2014)
  36. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: Proceedings of the 27th international conference on machine learning (ICML-10), pp. 807–814 (2010)
  37. Pestana, J., Mellado-Bataller, I., Sanchez-Lopez, J.L., Fu, C., Mondragón, I.F., Campoy, P.: A general purpose configurable controller for indoors and outdoors gps-denied navigation for multirotor unmanned aerial vehicles. J. Intell. Robot. Syst. **73**(1-4), 387–400 (2014)
  38. Polvara, R., Patacchiola, M., Sharma, S., Wan, J., Manning, A., Sutton, R., Cangelosi, A.: Autonomous quadrotor landing using deep reinforcement learning arXiv:1709.03339 (2017)
  39. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y.: Ros: An Open-Source Robot Operating System. In: ICRA Workshop on Open Source Software, vol. 3, pp. 5. Kobe, Japan (2009)
  40. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-Cnn: Towards Real-Time Object Detection with Region Proposal Networks. In: Advances in Neural Information Processing Systems, pp. 91–99 (2015)
  41. Rudol, P., Doherty, P.: Human Body Detection and Geolocalization for Uav Search and Rescue Missions Using Color and Thermal Imagery. In: 2008 IEEE Aerospace Conference, pp. 1–8. IEEE (2008)
  42. Sadeghi, F., Levine, S.: (cad)<sup>2</sup>rl: Real single-image flight without a single real image arXiv:1611.04201 (2016)
  43. Sampedro, C., Bavle, H., Rodríguez-Ramos, A., Carrio, A., Fernández, R.A.S., Sanchez-Lopez, J.L., Campoy, P.: A Fully-Autonomous Aerial Robotic Solution for the 2016 International Micro Air Vehicle Competition. In: 2017 International Conference on Unmanned Aircraft Systems (ICUAS), pp. 989–998. IEEE (2017)
  44. Sampedro, C., Bavle, H., Sanchez-Lopez, J.L., Fernandez, R.A.S., Rodriguez-Ramos, A., Molina, M., Campoy, P.: A Flexible and Dynamic Mission Planning Architecture for Uav Swarm Coordination. In: 2016 International Conference on Unmanned Aircraft Systems (ICUAS), pp. 355–363. IEEE (2016)
  45. Sanchez-Lopez, J.L., Molina, M., Bavle, H., Sampedro, C., Fernández, R.A.S., Campoy, P.: A multi-layered component-based approach for the development of aerial robotic systems: The aerostack framework. J. Intell. Robot. Syst. **88**(2–4), 683–709 (2017)
  46. Scherer, J., Yahyanejad, S., Hayat, S., Yanmaz, E., Andre, T., Khan, A., Vukadinovic, V., Bettstetter, C., Hellwagner, H., Rinner, B.: An autonomous multi-uav system for search and rescue. In: Proceedings of the First Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use, pp. 33–38. ACM (2015)
  47. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. **15**(1), 1929–1958 (2014)

- 
48. Sun, J., Li, B., Jiang, Y., Wen, C.Y.: A camera-based target detection and positioning uav system for search and rescue (sar) purposes. *Sensors* **16**(11), 1778 (2016)
  49. Tomic, T., Schmid, K., Lutz, P., Domele, A., Kassecker, M., Mair, E., Grixia, I.L., Ruess, F., Suppa, M., Burschka, D.: Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue. *IEEE Robot. Autom. Mag.* **19**(3), 46–56 (2012)
  50. Xiang, G., Hardy, A., Rajeh, M., Venuthurupalli, L.: Design of the Life-Ring Drone Delivery System for Rip Current Rescue. In: 2016 IEEE Systems and Information Engineering Design Symposium (SIEDS), pp. 181–186. IEEE (2016)
  51. Zhang, F., Leitner, J., Milford, M., Upercroft, B., Corke, P.: Towards vision-based deep reinforcement learning for robotic motion control arXiv:[1511.03791](https://arxiv.org/abs/1511.03791) (2015)
  52. Zhu, Y., Mottaghi, R., Kolve, E., Lim, J.J., Gupta, A., Fei-Fei, L., Farhadi, A.: Target-Driven Visual Navigation in Indoor Scenes Using Deep Reinforcement Learning. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 3357–3364 (2017)
  53. Zingg, S., Scaramuzza, D., Weiss, S., Siegwart, R.: Mav Navigation through Indoor Corridors Using Optical Flow. In: 2010 IEEE International Conference on Robotics and Automation (ICRA), pp. 3361–3368 (2010)

**Carlos Sampedro** received the BSc in Industrial Engineer (major in Industrial Electronics) in July 2011 obtaining the best marks degree award, and the Master’s degree in Automation and Robotics in July 2014, both from the Universidad Politecnica de Madrid (UPM), Madrid, Spain. He is currently working as a researcher and Ph.D. candidate in the Computer Vision and Aerial Robotics (CVAR) research group belonging to the Centre for Automation and Robotics and Universidad Politecnica de Madrid (UPM-CSIC). To this aim, he received a pre-doctoral grant from the Universidad Politecnica de Madrid in January 2017. In addition, he was a visiting researcher during three months (Sep. to Dec. 2015) at the Arizona State University (AZ, USA).

His research is focused on the application of learning-based techniques applied to aerial robotics, with special interest on object detection and recognition using machine learning and deep learning techniques, the development of Deep Reinforcement Learning algorithms to UAV control, and UAV autonomous navigation.

**Alejandro Rodriguez-Ramos** is a Telecommunication Engineer (major in Electronics and Micro-Electronics) graduated from Universidad Politecnica de Madrid (UPM). Currently, he is working as a researcher and Ph.D. candidate at the Computer Vision and Aerial Robotics (CVAR) group of the Centre for Automation and Robotics (UPM-CSIC). Previously, he worked for more than a year in the Aerospace sector, contributing to projects of the European Space Agency. His research interests include Deep Reinforcement Learning techniques applied to aerial robotics, Deep Learning, aerial robotics and image processing.

**Hriday Bavle** is a Ph.D. candidate at the Computer Vision and Aerial Robotics Group (CVAR), Universidad Politecnica de Madrid (UPM), Spain. He received his Bachelors in Aerospace and Masters in Avionics from Amity University, India. Aerial robotics being his core research field, his specialization includes localization and mapping techniques applied to UAVs in unknown indoor environments using several computer vision and sensor fusion techniques. He is one of the lead developers and testers of the Aerostack software framework developed in the CVAR group.

**Adrian Carrio** is a Ph.D. candidate in Automation and Robotics at Universidad Politecnica de Madrid (UPM) and a former researcher at ASU and MIT, where he has worked in the development of vision-based collision avoidance systems for UAVs. His research interests include machine learning for object recognition and autonomous UAV navigation.

**Paloma de la Puente** obtained her engineering degree in Automatic Control and Electronics in November 2007 and her Ph.D. in Robotics and Automation in December 2012, both from Universidad Politecnica de Madrid (UPM). She was a postdoctoral researcher at DISAMUPM and at ACIN Institute of Automation and Control-Vienna University of Technology. Her main research interests are related to navigation, mapping, spatial cognition and sensor data processing.

**Pascual Campoy** is a Full Professor on Automatics at the Universidad Politecnica de Madrid UPM (Spain) and visiting professor in TUDelft (The Netherlands), he has also been visiting professor at Tong Ji University (Shanghai-China) and Q.U.T. (Australia). He currently lectures on Control, Machine Learning and Computer Vision.

He is leading the Research Group on “Computer Vision and Aerial Robotics” at UPM within the Centre for Automation and Robotics (CAR), whose activities are aimed at increasing the autonomy of the Unmanned Aerial Vehicles (UAVs) by exploiting the powerful sensor of vision, using cutting-edge technologies in image processing, control and artificial intelligence.

He has been head director of over 40 R&D projects, including R&D European projects, national R&D projects and over 25 technological transfer projects directly contracted with the industry. He is author of over 200 international scientific publications and nine patents, three of them registered internationally. He is awarded several international prizes in UAV competitions: IMAV12, IMAV13, IARC14, IMAV16 and IMAV17.