2019 IMEKO TC-19 International Workshop on
Metrology for the Sea
Genoa, Italy, October 3-5, 2019

Visual Inspection of Vessels Cargo Holds: Use of a Micro-Aerial Vehicle as a Smart Assistant*

Alberto Ortiz[1], Francisco Bonnin-Pascual[1], Emilio Garcia-Fidalgo[1],
Joan Pep Company-Corcoles[1], Kai Yao[1]

[1]*Dep. Mathematics and Computer Science, Universitat de les Illes Balears (UIB), Spain*

Abstract – **This paper describes a *Micro-Aerial Vehicle* (MAV) intended for the visual inspection of cargo holds, whose development, among others, takes place within the context of the EU-funded H2020 project ROBINS, with the purpose of making ship inspections safer and more cost-efficient. To this end, the vehicle is equipped with specific sensors that are to permit teleporting the surveyor to the areas that need inspection, while the focus of the control software is on providing enhanced functionality and autonomy for the inspection processes. All this has been accomplished in the context of the supervised autonomy paradigm, by means of the definition of different autonomy levels and functionalities (including obstacle detection and collision prevention), and extensive use of behaviour-based high-level control, all intended for visual inspection as already mentioned. Automatic detection of defects is also addressed as part of ROBINS goals, through the adoption of deep learning approaches for enhanced performance. Results for some experiments conducted to assess the different functionalities are reported at the corresponding sections of the paper.**

I. INTRODUCTION

The movement of goods by vessels is today one of the most time- and cost-effective ways of transportation. Nowadays, the demand for maritime transport services is dealt by large-tonnage vessels specific for the kind of product to freight, namely oil tankers, bulk carriers, and general cargo or container ships, to name but a few. As any other installation or infrastructure, this type of vessels requires regular maintenance to avoid its deterioration due to a varied set of causes, ranging from design mistakes, use of sub-standard materials or procedures, structural overload or normal decaying of the metallic structures in the sea. Otherwise, accidents can result, with maybe catastrophic consequences for the crew (and passengers), environmental pollution or damage and/or total loss of the ship, its equipment and its cargo.

The inspection of those ship-board structures by humans is a time-consuming, expensive and commonly hazardous activity, which may cause damage to coating, creating the need to search for alternative solutions. Simultaneously, the need for broader and intensified data acquisition grows to provide a better basis for refined condition assessment. In line with the aforementioned, the EU-funded FP7 projects MINOAS (finished on 2012) and INCASS (finished on 2017) had among their main goals the development of robotic devices to assist and simplify the inspection of different vessels areas. As a follow-up, the EU-funded H2020 project ROBINS (started in 2018, `https://www.robins-project.eu/`) aims at filling the technology and regulatory gaps that today still represent a barrier to the adoption of Robotics and Autonomous Systems (RAS) in activities related to the inspection of ships, starting from understanding end-user's actual needs and expectations as well as analyzing how existing or near-future technology can meet them.

On the technological side, ROBINS intends to: (1) enhance the capabilities of RAS as for navigation and localization, access to and mobility within the environment, as well as improve their ability regarding sensing and probing, to increase the TRL attained within the frameworks of the previous projects; and (2) provide new tools for processing images and data collected by sensors, aiming at (a) the generation of 3D models that permit teleporting the surveyor to the areas of interest (by means of e.g. virtual tours) and (b) highlighting defects in collected images.

On the normative side, and concerning the regulatory aspects relevant to the use of RAS in ship inspection, ROBINS has the following objectives: (1) define criteria, testing procedures and metrics for the evaluation of RAS performance in terms of safety, functionality, dependability, security, data quality and economic viability; (2) design, implement and assess a Testing Facility (TF) where repeatable tests and measurements can be performed for the evaluation of the compliance between RAS and the requirements; and (3) provide a framework to assess the equivalence between the outcomes of RAS-assisted ship inspections and traditional inspection procedures. The possibility to replace, or at least reduce, field trials by means of equivalent test protocols carried out in the TF is one of the key concepts in ROBINS, and comes from the consideration that field trials on board ships are generally very

*This work is partially supported by projects ROBINS (EU-H2020, GA 779776), PGC2018-095709-B-C21 (MCIU/AEI/FEDER, UE), and PROCOE/4/2017 (Govern Balear, 50% P.O. FEDER 2014-2020 Illes Balears), and by scholarship BES-2015-071804 (MCIU/AEI).

Fig. 1. (left) Staging required during a vessel inspection. (right) Oil tanker in shipyard during construction.

difficult to be arranged, can be quite expensive and usually cannot be executed in controlled and reproducible testing conditions, what constitute severe limitations for the development of RAS platforms for ship inspection. Despite the aforementioned, field trials are not to be suppressed, but, to the contrary, are to provide feedback to revise the testing protocols and the TF itself, thus creating an iterative refining process that converges to a TF relevant for the assessment of RAS for ship inspection. (For additional details about the TF, the reader is referred to [1].)

Cargo holds and cargo tanks is one of the types of operational scenarios which is considered by the ROBINS project, as representative of the different environments where costs and risks connected to inspection activities are more significant: wide volumes with significant heights, characteristic of bulk carriers, container ships, tankers, and general cargos holds, which require costly access means to reach high points, while large unobstructed heights between levels imply severe consequences for the safety of surveyors. In these environments, inspections are actually carried out mainly by means of scaffolding, portable ladders, and cherry pickers (see Fig. 1). This paper describes recent progress achieved within the context of project ROBINS as for the development of an aerial platform specifically devised for the collection of inspection data from vessels' wide spaces, such as cargo holds, focusing on fitting the robot with enhanced autonomy functionalities for those particular areas. Besides, we also contribute with new visual methods for defect detection in images. After a first year since the project start in 2018, a first almost-complete version of the platform is available at both the hardware and software levels, together with software modules for processing the collected inspection data, e.g. for 3D reconstruction and defect detection. Advances in some of these different sides of the project are reported along the next sections of the paper.

II. AERIAL PLATFORM

Among others, the vertical structures that can be found in vessel holds are of prime importance. To make proper *repair/no repair decisions*, the surveyor must be provided with, among others, imagery detailed enough so as to enable the remote visual assessment of these structures. To this end, the platform can be either required to sweep the relevant metallic surfaces and grab pictures at a rate compatible with its speed, or else provide visual evidence of the state of a particular area suspected of being defective. Those images must as well be tagged with pose information, so that the areas suspected of being defective can be located on the vessel structure, or even for comparing images across inspections.

Therefore, the main requirements for the aerial platform stem directly from the very nature of the inspection process: the vehicle must be able to perform vertical, stationary and low speed flight, as well as permit indoor flight. These requirements rapidly discard fixed-wing aircrafts and focus the search on helicopter-type UAVs, naturally capable of manoeuvres such as hovering and *vertical take-off and landing* (VTOL). Additionally, the platform should not only rely on GPS data for positioning, because it could be required to operate indoors or in poor GPS reception areas (e.g. due to satellites being occluded by the vessel structures, multi-path effects, etc.)

Because of their fast deployment times and convenient size, a number of recent works have considered the use of multirotor-based Micro-Aerial Vehicles (MAVs) within the context of the inspection and monitoring of industrial facilities and assets, for data collection at remote or safety-compromised areas, difficult to reach by humans and ground vehicles, and with large areas to be covered as fast as possible [2, 3, 4]. The popularity these vehicles have gained in recent years has lead to the availability of a number of control and navigation solutions. They differ mainly in the sensors used to solve the navigation tasks, the amount of processing that is performed onboard/offboard, and the assumptions made about the environment. Apart from other devices, such as infrared and ultrasound sensors, laser scanners [5] and, lately, vision cameras [6] have become the preferred sensor modalities to undertake these tasks, mostly within Simultaneous Localization and Mapping (SLAM) frameworks and combined with Inertial Measuring Units (IMU).

A. Platform Overview

Our aerial platform is based on a multi-rotor design fitted with (1) a Flight Management Unit (FMU) for platform stabilization in roll, pitch and yaw, as well as thrust control, (2) a 3-axis IMU —which, according to today standards, is typically part of the FMU—, (3) a sensor suite able to supply vehicle 3D speed and height measurements, as well as distances to surrounding obstacles, (4) inspection sensors and (5) an embedded PC which avoids sending sensor data to a base station, but process them onboard and, thus, prevent communications latency inside critical control loops. Figure 2 shows a realization of this platform taking as a basis the Matrice 100 quadrotor by DJI.

To be more precise, apart from the FMU, the realization

Fig. 2. ROBINS aerial platform, fitted with a laser scanner, a height sensor, a camera set feturing LED-based illumination, and an embedded PC.

of Fig. 2 features: (a) the lightweight laser scanner Hokuyo UST-20LX, which provides 20 meters coverage for a 270° angular sector, which is used to estimate 2D speed as well as distances to the surrounding obstacles; (b) a downward-looking LIDAR-Lite v3 laser range finder used to supply height data for a maximum range of 40 meters; (c) two cameras to collect, from the vessel structures under inspection, RGB-D images on demand (Intel Realsense D435i camera) and video footage (Zenmuse X3 mounted on a gimbal); (d) two 7W DC20-24V pure white LEDs, supplying 2×700 lumen for a 140° beam angle; and (e) an Intel NUC7I7BNH embedded PC featuring an Intel Core i7-7567U 2×3.5GHz processor and 16 GB RAM.

Apart from other sensor suites capable of also supplying speed and height measurements, the previous configuration allows navigation under low-light conditions, as required in certain vessel compartments such as e.g. oil tanker cargo holds, for which a single manhole-sized entry point is typically available (see Fig. 3). Further, the LED-based system is intended to facilitate the capture of useful images despite the absence of ambient lighting. Thanks to a specific power system, they can be remotely dimmed from the embedded PC to adjust to the available illumination and the operating distance to the walls during flight.

B. Control Architecture

From a control viewpoint, the aerial platform implements a control architecture that follows the Supervised Autonomy (SA) paradigm [7]. This is a human-robot framework where the robot implements a number of autonomous functions, including self-preservation and other safety-related issues, which make simpler the intended operations for the user, so that the operator, which is allowed to be within the general platform control loop, can focus in accomplishing the task at hand. Within this framework, the communication between the robot and the user is performed via qualitative instructions and explanations: the user prescribes high-level instructions to the platform through an adequate Human Interaction Device (HID) while the robot provides instructive feedback.

Fig. 3. (left) Oil tanker manhole entry point. (right) Typical oil tanker cargo hold.

In our case, qualitative commands can be issued through both a Radio Controller (R/C) and a gamepad, as well as from the input side of the Graphical User Interface (GUI), while the robot feedback is handled from the output side of the GUI. While the platform is taking care of attitude stabilization, speed and height control, as well as prevents collisions, the user is only required to issue simple/qualitative commands related to the inspection operation, such as go up, go down, go left, go right, advance in a certain direction, etc.

The architecture comprises two separate agents: the MAV and the Base Station (BS). All the state estimation and control algorithms run over the computational resources of the MAV: as usual, the FMU runs low-level control tasks —attitude stabilization and direct motor control—, while the embedded PC executes, on top of the Robot Operating System (ROS) running over Linux Ubuntu, high-rate ROS nodes that (1) implement mid-level control tasks —height and speed controllers—, and (2) estimate velocity and height, as well as the distances to the closest obstacles surrounding the platform. Lower-rate top-level control also runs over this processor in the form of a number of different behaviours that implement the higher level autonomous functionalities of the platform. These behaviours combine the user desired speed command with the available sensor data, to obtain final and safe speed and height set-points that are sent to, respectively, the speed and height controllers.

The BS mainly runs the GUI used to supply the operator with information about the state of the platform as well as about the task in execution, e.g. images collected via the vision system attached to the platform. It runs ROS over Linux Ubuntu and is linked with the MAV via DJI Lightbridge and 5 GHz WiFi connections. The qualitative user commands and the robot feedback travel, respectively, forth and back through them.

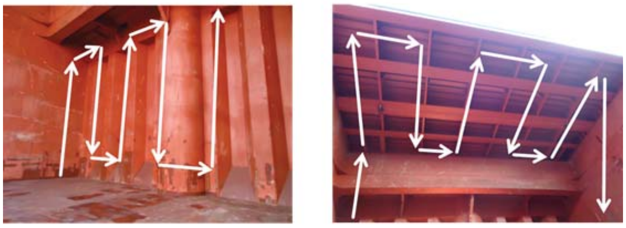

Fig. 4. Example of paths which are to be followed by the MAV to ensure an adequate level of coverage of the structure under inspection.

The MAV control software has been designed around open-source components and following modularity and software reutilization principles. In this way, adapting the framework for different platforms involving different payloads, or the selective activation of software modules, can be performed in a fast and reliable way.

C. Inspection Data Collection Capabilities

During flight, the aerial platform is intended to collect inspection data on demand or at a fixed rate, e.g. 10 fps, as well as log flight data. Regarding the latter, the vehicle pose, i.e. 3D position and 3D attitude, as well as the distance to inspected surfaces become of particular relevance here in order to be able to associate 3D position and scale information to the data collected and ultimately to the defects found after processing the images gathered. Besides, as an additional functionality during inspection missions, we aim at ensuring proper surface coverage in order to supply the surveyor with complete information about the surface under inspection, e.g. a bulkhead. This is also to contribute in an effective way to the 3D reconstruction of the area from the visual data collected (as well as other data modalities). To implement the aforementioned functionalities, the platform not only needs to estimate its speed but also its position, as well as it has to plan its motion to properly cover the areas of interest. Figure 4 illustrates the type of coverage which is expected, while Fig. 6 shows examples of sweeping flights and the 3D map resulting from one these flights.

D. Results from Field Trials

This section reports on a number of inspection missions performed during first field trials taking place onboard a 12.000 DWT Ro-Ro Cargo vessel while at drydock at the end of March 2019. Tests were performed at the Main Deck cargo hold (first day) and at the TankTop cargo hold (second day), both dry spaces of the vessel. It was a rather new vessel, built in 2012, which was in good condition, but served as a successful test bench for the different functionalities of the aerial platform. Thanks to the design of the base station and the vehicle, deployment times were about 5 minutes. Images from the field trials and plots for some of the flights are shown in Fig. 6.

III. DEFECT DETECTION

Steel surfaces can be affected by different kinds of defective situations, being coating breakdown/corrosion (CBC), in any of its different forms, the most common defect. As already mentioned, an early detection prevents vessel structures from suffering major damage which can ultimately compromise their integrity. ROBINS aim in this regard is to simplify visual inspection processes by means of tools for conveniently visualizing the inspection data collected, as well as by detecting and highlighting potential defects so as to draw the attention of the surveyor to relevant points of the hold under inspection.

To counteract complex lighting conditions and a great diversity of other conditions due to data obtained from possibly different robots, ROBINS approach for defect detection adopts Deep Convolutional Neural Networks (DCNN)-based methodologies as highly robust machine-learning approaches. As already known, recently DCNNs have already showed good results for object recognition in images [8], although what is most interesting is the fact that they have been shown highly promising for industrial inspection applications [9]. In contrast to manually designed image processing solutions, DCNNs automatically generate powerful features, i.e. *learn the representation*, from training data by means of hierarchical learning strategies with a minimum of human interaction or expert process knowledge.

In more detail, in this section we report on detection results for pixel-level CBC detectors based on deep learning semantic segmentation through a fully convolutional neural network (FCN) trained end-to-end. These networks learn and predict dense outputs from a whole-image-at-a-time by dense backpropagation/feedforward computation [10]. As part of the fine-tuning of the detector for the problem at hand, we have considered the FCN-8s architecture described by [11]. Due to their characteristics, our configuration allows the detector to find both large and small image areas affected by CBC, typically corresponding to, respectively, general corrosion and pitting. Table 1 and Fig. 5 show successful CBC detection results using the *Focal Loss* (FL) function [12] to train the detector. Performance is measured in terms of standard metrics for semantic segmentation solutions, such as *pixel accuracy* (PA), *mean accuracy* (MA), *mean region intersection over union* (mIU) and *frequency weighted IU* (fwIU), as well as through the traditional *precision* (P) and *recall* (R) values.

IV. CONCLUSIONS

A Micro-Aerial Vehicle to be used for vessel visual inspection of vessel cargo holds has been described. The MAV control approach is based on the SA paradigm,

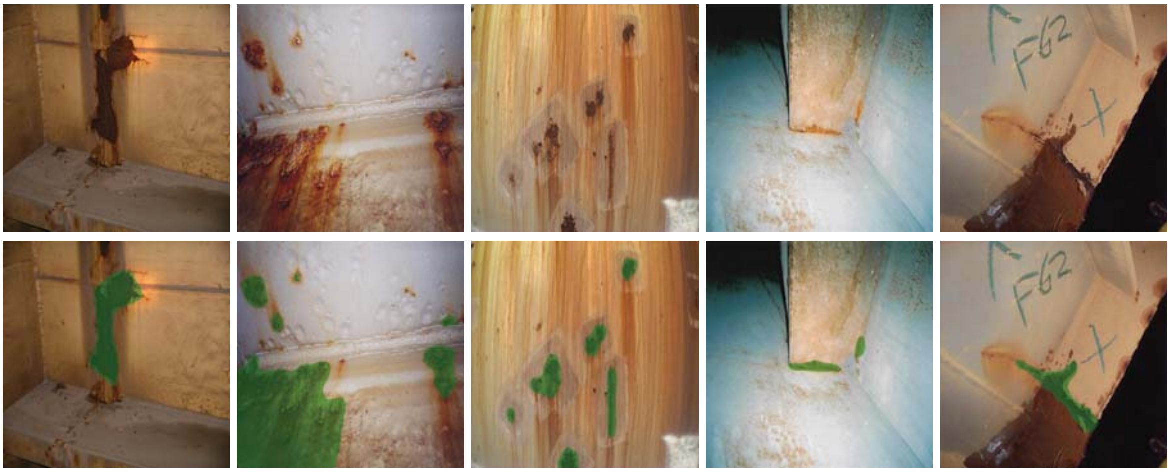

Fig. 5. Results of the corrosion detector: (1st row) original image, (2nd row) results superimposed in green.

Table 1. Detection performance of the corrosion detector.

	PA	MA	mIU	fwIU	P	R
FL	0.95	0.87	0.83	0.90	0.95	0.83

and hence the user is introduced in the platform control loop. For the specific problem of visual inspection, a behaviour-based control architecture tightly linked to the SA paradigm, including aspects such as vehicle state estimation, behaviour-based command generation and flight control, has been outlined. Results from first field trials, regarding control architecture validity and usefulness for visual inspection, have been reported. Furthermore, a deep learning-based CBC detector has been discussed and some experimental results shown.

REFERENCES

[1] L. Poggi, T. Gaggero, M. Gaiotti, E. Ravina, and C. M. Rizzo. Robotic inspection of ship structures: how to fill the gap between available technologies and current survey practices? In *Proc. IMEKO TC-9 International Workshop on Metrology for the Sea*, 2019.

[2] Tolga Ozaslan, Shaojie Shen, Yash Mulgaonkar, Nathan Michael, and Vijay Kumar. Inspection of penstocks and featureless tunnel-like environments using Micro UAVs. In *Proc. Intl. Conference on Field and Service Robotics*. 2015.

[3] Martin Stokkeland, Kristian Klausen, and Tor A. Johansen. Autonomous visual navigation of unmanned aerial vehicle for wind turbine inspection. In *Proc. Intl. Conference on Unmanned Aircraft Systems*, 2015.

[4] Alberto Ortiz, Francisco Bonnin-Pascual, Emilio Garcia-Fidalgo, and Joan P. Company-Corcoles. Vision-based corrosion detection assisted by a micro-aerial vehicle in a vessel inspection application. *Sensors*, 16(12), 2016.

[5] Ivan Dryanovski, Roberto G. Valenti, and Jizhong Xiao. An open-source navigation system for micro aerial vehicles. *Autonomous Robots*, 34(3):177–188, 2013.

[6] Chiara Troiani, Agostino Martinelli, Christian Laugier, and Davide Scaramuzza. Low computational-complexity algorithms for vision-aided inertial navigation of micro aerial vehicles. *Robotics and Autonomous Systems*, 69:80–97, 2015.

[7] G. Cheng and A. Zelinsky. Supervised Autonomy: A framework for human-robot systems development. *Autonomous Robots*, 10(3):251–266, 2001.

[8] D. Maturana and S. Scherer. VoxNet: A 3D convolutional neural network for real-time object recognition. In *Proc. IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, 2015.

[9] Daniel Weimer, Bernd Scholz-Reiter, and Moshe Shpitalni. Design of deep convolutional neural network architectures for automated feature extraction in industrial inspection. *CIRP Annals*, 65(1):417–420, 2016.

[10] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Intl. Conference on Computer Vision and Pattern Recognition*, 2015.

[11] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):640–651, apr 2017.

[12] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *Proc. IEEE Intl. Conference on Computer Vision*, oct 2017.

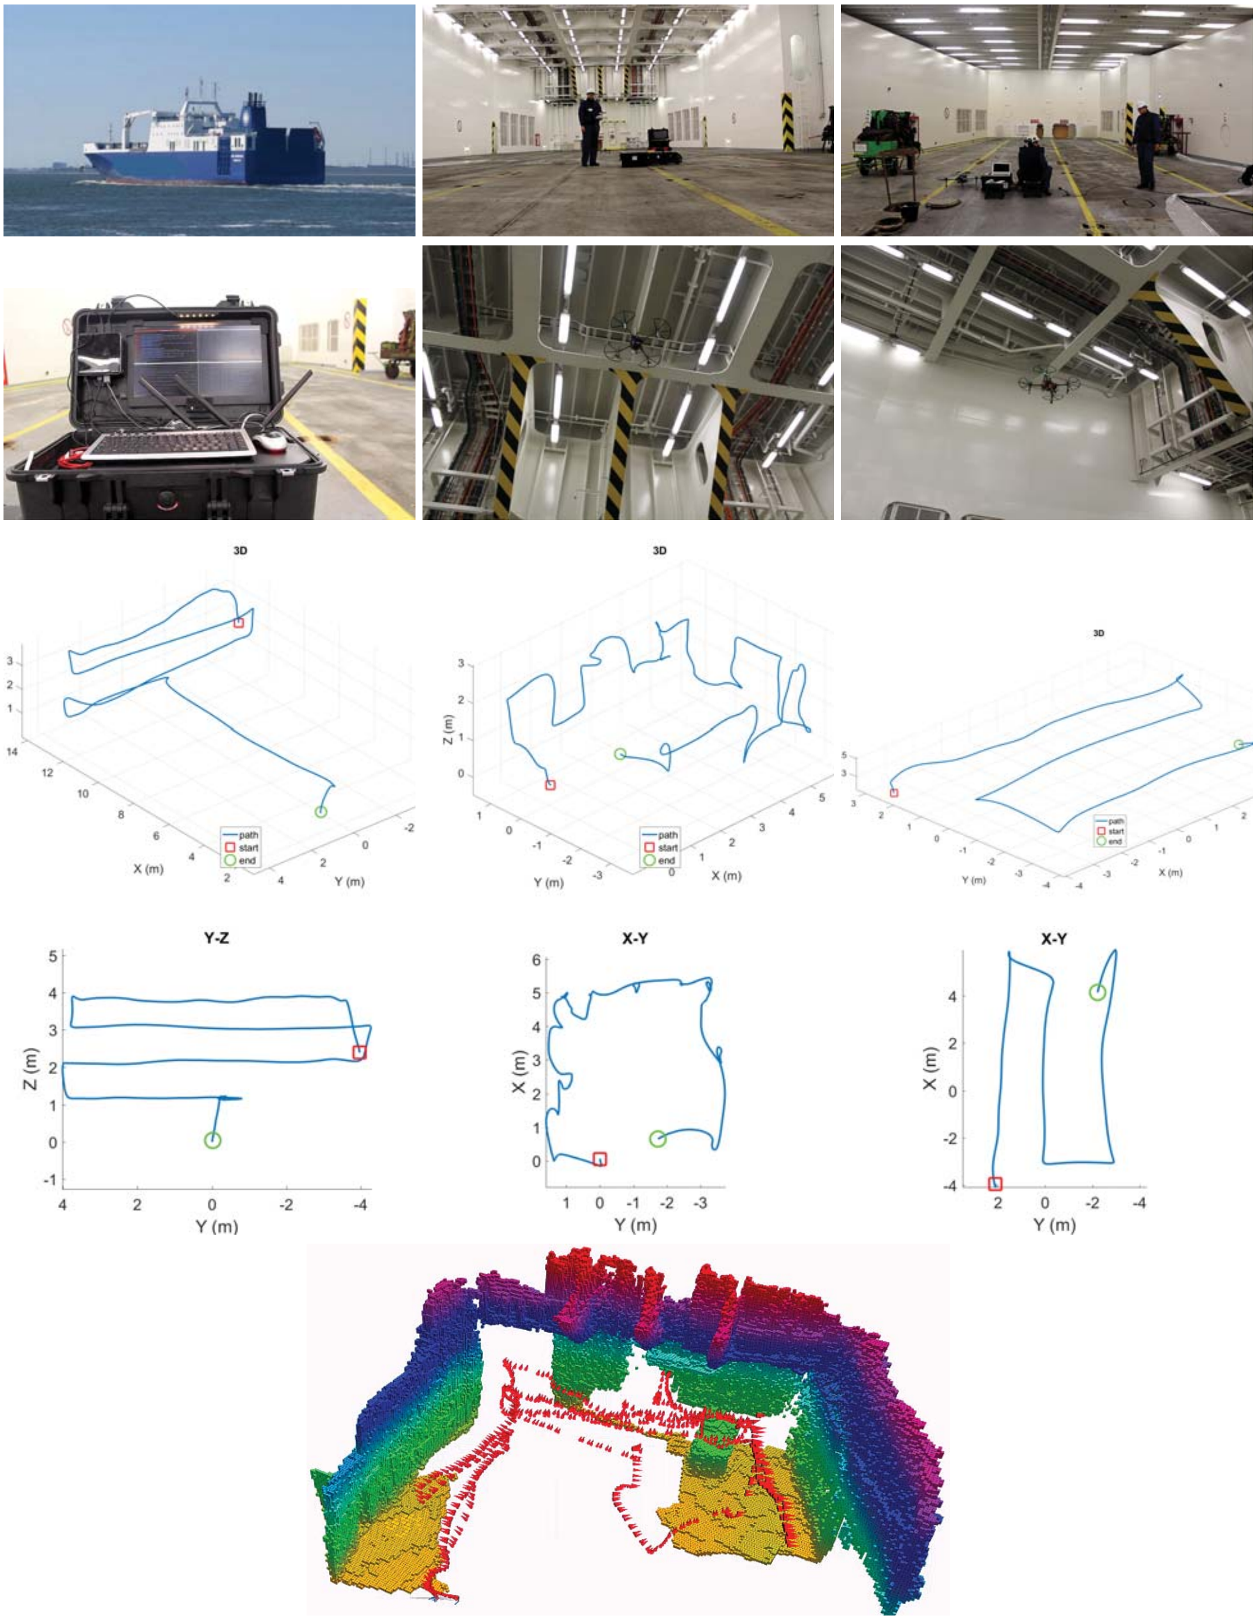

Fig. 6. Field trials: (1st row) Ro-Ro vessel, Main Deck cargo hold, and TankTop cargo hold; (2nd row) base station and images of the platform during inspection flights; (3rd/4th row) estimated paths for different sweeping flights in the cargo holds; (5th row) 3D map for one of the flights [path of the MAV shown in red].