Final prototype and validation for sewer inspection procedure

ARSI Deliverable D26.14

PDTI Urban Robotics - Sewer Inspection

December 20, 2018

Contents

1 Introduction **3**

2 MAV Platform **4**

2.1 Phase III platform overview 4

2.2 Propulsion system 5

2.3 Vibration absorption 6

2.4 Field tests 6

2.5 Platform improvements 7

3 Data processing **8**

3.1 3D reconstruction 8

3.2 Structural assessment 9

3.3 Distance map inspection 10

3.4 Obstacles detection 11

3.5 Defects 12

3.6 Tests and results 13

4 User interfaces **18**

4.1 Mission Planner 20

4.2 Operator Console 20

4.3 Data Analysis Interface 21

5 Conclusions and future work **25**

1 Introduction

In this document we present the final prototype for the ARSI Micro Air Vehicle (MAV) and data analysis system at the conclusion of the final phase of the project.

In section 2 we discuss the MAV platform, and in particular the firmware and propulsion system changes carried out since the July 2018 evaluation in Virrei Amat, Barcelona [1].

In section 3 we present the latest advances in the 3D reconstruction and defect detection modules, along with results from recent field tests in Barcelona. Then in section 4 we describe the new functionality added to the user interfaces, such as the Virtual Reality (VR) mode in the Data Analysis App.

We conclude this report by reflecting on the progress made in Phase III as well as the results from the final evaluation in December 2018.

References

[1] ARSI Consortium, *Deliverable D26.10 - Phase III tests and tests results*, 2018.

[2] ARSI Consortium, *Deliverable D26 - Changes and Improvements Based on Phase 2 Evaluation*, 2018.

2 MAV Platform

In this section we present the changes made to the MAV platform since the July evaluation in Barcelona. For a detailed description of previous MAV work in Phase III, please refer to deliverables [1] and [2].

After the July 2018 evaluation in Virrei Amat, it was concluded that the MAV flight was still too unstable for reliable sewer inspection and data acquisition. Therefore, the ARSI team identified two areas where changes could improve the MAV stability in sewer operations: the propulsion system and autopilot firmware.

In this section we describe those changes and their impact on flight stability during our various field tests in Barcelona.

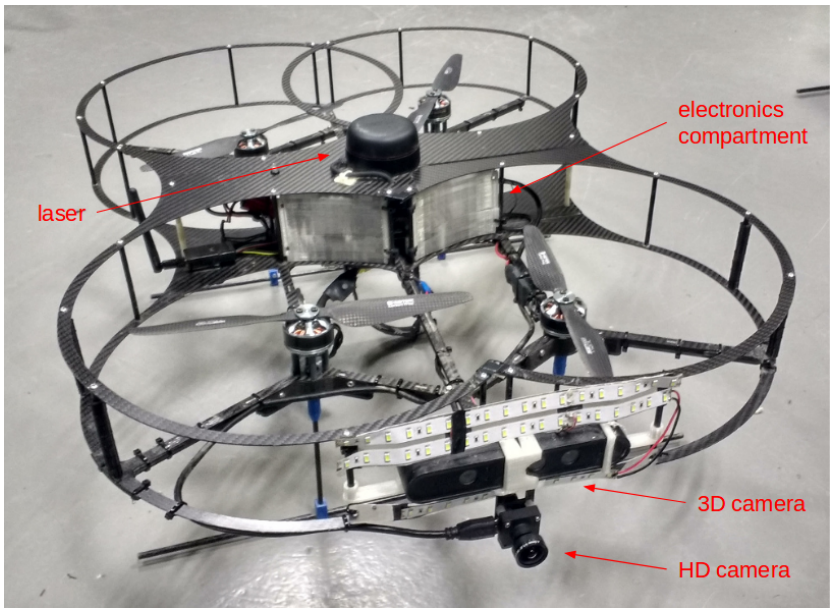

Figure 1: Final ARSI MAV platform

2.1 Phase III platform overview

The final MAV platform at the end of Phase III is shown in figure 1. It was entirely redesigned after Phase II, following the comments from the 2017 evaluation in Barcelona. The most notable changes includes:

- Introduction of a lightweight carbon-fiber frame providing full protection in case of contact or impact against the sewer walls;

- Integration of a 360 degree field-of-view laser allowing backwards flight to return autonomously to the point of origin;
- Mounting of all the critical electronics in an internal space protected from water and dust;
- Improved LED illumination;
- New propulsion system delivering a payload capacity of 1kg and a flight autonomy of approximately 14 minutes;
- Addition of a wide-angle HD camera at the front of the MAV, to provide high-resolution coverage of the entire tunnel section.

2.2 Propulsion system

The propulsion system of a Micro Air Vehicle (MAV) comprises the propellers, the motors and their Electronic Speed Controllers (ESCs). During Phase II of the project, our prototype used 10" (25cm) propellers with 4.5 pitch, and T-Motor MN3110 780kv (rpm*V) motors with 30A TBS ESCs running at 14.8V. The Phase II evaluation in October 2017 showed that this configuration delivered reactive and stable control in the sewers, but insufficient autonomy and payload capacity for commercial sewer inspection.

In Phase III we worked with our partner DroneTools to redesign the entire platform in order to meet the project requirements. We came up with a design using overlapping 14" (35cm) propellers with 5.5 pitch, using T-Motor Antigravity MN4006 380kv motors and Turnigy Flush 30A ESCs running at 22.2V.

The larger propellers in our new design meant that the MAV would move more air with each revolution, creating more thrust and therefore increasing the payload capacity. The design allows carrying heavier and higher-capacity batteries, therefore increasing the flight time, and includes more powerful but slower motors than those used in Phase II. While our new design performed well in our laboratory environment, even in our simulated sewer tunnels, our field tests in real sewers showed that it could not match the flight stability of our Phase II prototype, as was evident at the July evaluation in Virrei Amat, Barcelona.

After this evaluation, we decided to revisit the propulsion system to strike a compromise between the Phase II and III designs, to achieve acceptable autonomy and payload capacity whilst delivering stable flight control. We used smaller 11" (28cm) propellers with faster T-Motor MN3510 700kv motors and the same Turnigy Flush 30A ESCs. Using this configuration we our MAV is able to carry 1kg of payload for an estimated flight time of 14 minutes.

2.3 Vibration absorption

During our analysis after the July evaluation, we also investigated the whether vibrations might be contributing to our flight stability issues. Vibrations can have dramatic effects on MAVs especially when it comes to mounting the autopilot. Our Pixhawk, like most autopilots, carries accelerometers and gyroscopes that are used in the sensor fusion for state estimation. These sensors are extremely sensitive to the vibrations generated by the MAV motors.

In our case, a detailed analysis revealed that the level of vibration on our new platform design was elevated, and was likely to induce subtle control issues. We carried out a detailed comparison of various mount options, and we were able to significantly reduce the impact of vibrations on the Pixhawk sensors.

2.4 Field tests

Our new propulsion system showed good performance in the laboratory, as well as in field tests at the Mercat del Born, Av. Pearson and Forum evaluation areas in Barcelona. In all these areas we were able to perform stable flights in narrow tunnels (T111, T130, and NT120A sections, see figure 2).

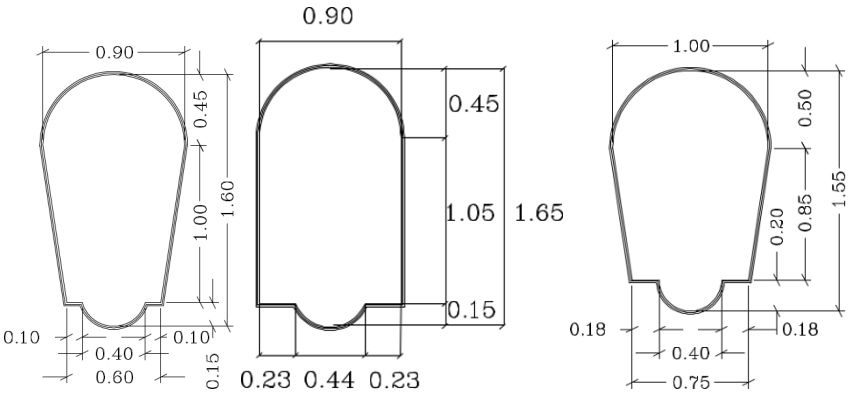

Figure 2: T111, T130 and NT120A section types

Flight videos are available for all these tests; note that the MAV flew autonomously, but we did have a safety pilot present to avoid accidents:

- Mercado del Borne (Dec 5th, 2018): video
- Forum (Dec 10th, 2018): video
- Av. Pearson (Dec 11th, 2018): video1, video2

Despite these good results during our field tests, the December evaluation in Forum, Barcelona, was unsatisfactory since the MAV exhibited serious flight instabilities. Please refer to the conclusion (section 5) for our plan to address this issue.

2.5 Platform improvements

Throughout Phase III we identified several limitations of the current MAV platform design and developed plans for an improved version with the following features:

- *Better structural design for the carbon fiber frame*

  A new reinforced carbon-fiber frame has been designed with beams crossing the base to increase the overall rigidity to the bottom structure, reduce vibrations in the frame, and improve flight stability as well as data capture.

- *Landing gear*

  Improved struts with a flexible PLA compound to absorb some of the impact during rough landings and help avoid tear and break.

- *Easier servicing of on-board electronics*

  The new design includes a service panel at the top for direct access to the on-board electronics, facilitating servicing, component replacing or cabling connecting/disconnecting.

- *Improved water and dust protection*

  The new design also includes additional protections around all the sensitive electronics to protect them from water and dust during flight.

- *MAV motors*

  The choice of MAV motors will be reviewed to reduce the risk hardware failure suffered in the recent Phase III field tests.

- *LED illumination*

  The current LED illumination was deemed satisfactory for inspection of the sewer walls and roof; however due to the typical topology of sewer tunnels, the central area of the image can be relatively dark and reduce the visibility for ground inspection.

  A new central long-range 3000 Lumens LED is being considered for better path illumination in the center of the sewer tunnel. This would improve motion estimation and data quality for ground inspection and defect detection.

3 Data processing

In this section we present the latest advances in the 3D reconstruction and detection modules developed in the last stage of phase III. The advances performed before the July 2018 evaluation were described in [1], while here we focus on the improvements implemented afterwards.

3.1 3D reconstruction

Before the July evaluation in Virrei Amat, several improvements were introduced in the 3D reconstruction module compared to Phase II. This included the automation of the reconstruction process and a new parameterization of the texture that took advantage of the GPU for a faster computation. After the field tests conducted on July 2018, we noticed that the accuracy of the reconstruction was still very dependent on the flight stability. In places where the odometry was not properly estimated or even lost, the obtained reconstruction suffered from these issues resulting in a bad reconstruction. To solve this we thus decided to introduce a more accurate estimation of the camera poses during the processing stage.

The new camera pose estimation is done using the RGB images provided by the drone, and performed offline during the processing stage. This estimation is based on a structure-from-motion approach that uses an accurate feature matching strategy and a global optimization of the camera poses. The obtained estimation is now more robust to the rapid movements of the platform and avoids the drift experienced when using the original odometry, as shown in Figure 3.

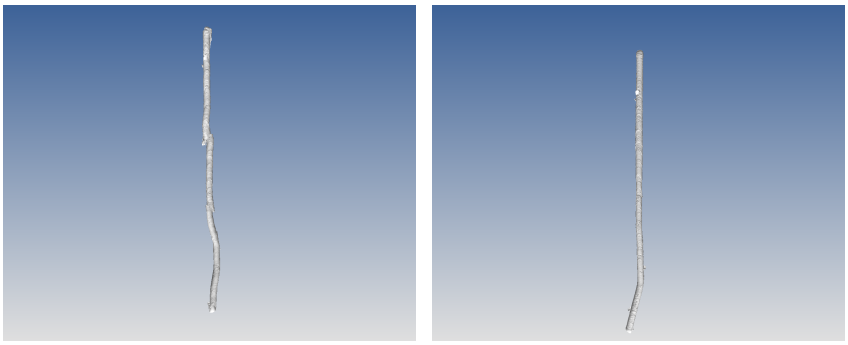

Figure 3: Comparison between original pose estimation (left) and the new one (right).

Besides this new camera estimation, we have incorporated other enhancements to further improve the accuracy of the reconstruction. In

particular, we noticed that the captured RGB and depth images could not be properly synchronized due to limitations on the camera driver. To handle this, we approximate this synchronization by interpolating depth maps between consecutive frames using the timestamps associated to each RGB and depth image. This results in a better alignment of the depth maps once projected in the 3D space. On the other hand, the estimation of the normals before the meshing step is now performed on each depth map independently instead of using the whole point cloud, which tends to provide a more accurate estimation. These two enhancements result in less artifacts and more smooth sewer models, as shown in Figure 4.

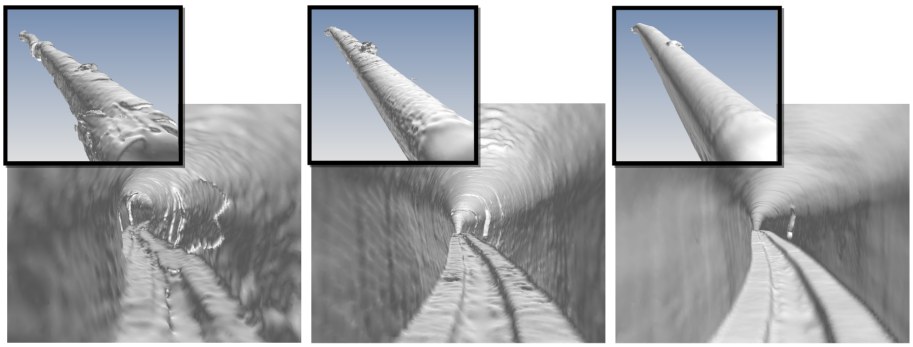

Figure 4: Comparison between original pose estimation (left), new one based on a global SfM approach (middle), and after improving depth synchronization and normal estimation (right).

In terms of texturing, the appearance delivered to the sewer models in Phase II suffered from blurring artifacts and a lower resolution compared to the original images. In order to solve this, we now include a projection step that projects each image into the 3D model and maps each view into a limited portion of the model. This projection is also directly performed into the texture image without resorting to the mesh vertices as done previously, which decouples the mesh resolution with its final appearance. All this results in improvements in terms of resolution and texture quality, as shown in Figure 5.

3.2 Structural assessment

The structural assessment of the inspected sewer takes as input the 3D reconstructed model and the expected (theoretical) structure of the sewer, as provided by the GIS database, in order to then compute differences between the two models. These differences are stored into a distance map that is subsequently passed to the inspection step. Com-

Figure 5: Comparison between old texturing approach (left) and the new one based on image projection (right).

pared to [1] no big changes have been done to this module, as the obtained results were already satisfactory. The main advances are related to the inspection of the distance map and the detection of defects, as detailed in the next subsections.

3.3 Distance map inspection

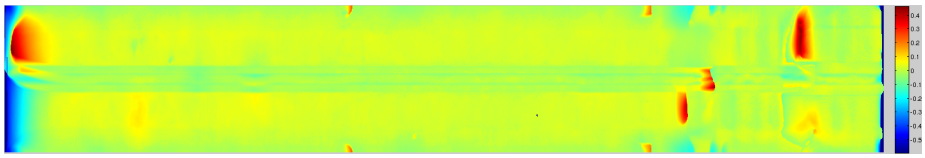

Figure 6: A texture image specifying the distance of each reconstructed point to the theoretical profile of the sewer.

Given the distance map specifying the positive and negative distances to the theoretical profile of the sewer (see, for instance, figure 6), detection of structural elements in the sewer is performed through image processing techniques. The texture image is an "unfolded" version of the sewer tunnel reconstruction, as if it were "cut open" along the ceiling. With respect to the set of elements that were detected in the previous version of this project (see deliverable D26.10), we have joint together different categories of defects in order to reduce false positive alarms. In particular, we report the following categories of incidents:

- **Unknown**: We start by removing those parts of the sewer reconstruction that are unidentifiable through the distance map. In

particular these are mainly reconstruction errors due to a bad collection of features, missmatching of the profile or others. It is important to detect such parts as their distance values live outside of the regular range. In practice, we look for outliers in the set of distance values and then for large compact regions of such values.

- **Manholes**: In the texture image (the 2D projection of the "unfolded" 3D reconstruction), manholes are located halfway in between the top and the bottom of the images. To locate manholes, the image is re-shaped as to have the manholes in the middle of the image, that is, as if the 3D reconstruction was mapped onto the 2D image cutting by the *cubeta*. On this new image, thresholding techniques are used to separate manhole candidates whose depth with respect to the profile is larger than a certain value (as manholes are, in fact, holes). Then different heuristics are applied to distinguish them: in particular, size, shape and position are evaluated to decide whether a candidate is in fact a manhole.

- **Structural elements**: This category includes all type of structural members such as junctions, gutters or connections. To detect them, the *cuvette* is located in the image so it can be deduced the position of the elements with respect to it. Secondly, a thresholding technique is applied to locate candidates and, finally, different heuristics are used to separate among other possible defects.

In all cases, whenever an element has been detected, that region is not analyzed again. That is, we do not report overlapping elements to avoid over-informing the user. We also report ***profile changes*** where the actual profile of the sewer has changed as detected in the registration step. Elements attached to a profile change are not reported since they might produce several false positives as the reconstruction close to a profile change is typically noisy.

3.4 Obstacles detection

To evaluate the serviceability of a particular section we look for obstacles. Obstacles are treated in two separated ways. In particular, we look for lacks of hydraulic capacity of the sewer and for large portions of the sewer where the texture image has a piece of sediment. This is different from the previous version of the project (see deliverable D26.10) in which obstacles or sediments where only sought in those regions where the sewer capacity was above a certain degree. In the current settings, we opt to identify both things at the same time in order to be more informative on how the platform detects a problem with the sewer serviceability.

- **Capacity lack**: In order to detect a section of the sewer where it has diminished its hydraulic capacity we perform a column-based analysis of the texture image. At each reconstructed profile, we integrate the negative differences (negative values mean protrusions with respect to the theoretical shape) and compare this with the actual area determined by the theoretical capacity of the section. A ratio between the computed capacity and its theoretical one determines the % of the sewer that is at service. If this situation occurs for a certain amount of profiles, this is, along a certain distance controlled by a parameter, we consider this situation a capacity lack of the sewer.
- **Sediments**: this category includes both sediments and general obstacles such as stones in the middle of the path. We proceed in a similar way as for the structural elements. We threshold the image (here we threshold the negative values) to have a set of candidates. The appearing blobs are selected based on their dimensions since small particles do not really constitute an obstacle and very big ones are most probably reconstruction errors, which are detected in a previous stage.

3.5 Defects

A new feature with respect to the previous version of the software tool is that we now report different types of defects. In particular, we report holes at the cuvette which may constitute a structural damage and for which the operator should be noticed. We have also developed a wall erosion detection module that does not use the distance map but the image itself taken from the platform. We detail both cases below.

Holes at cuvette

The procedure for detecting damages at the cuvette is rather similar to the one used for detecting structural elements. In particular structural elements are holes with a regular shape in a particular position of the sewer. Holes, though, may have any shape and reside within the limits of the cuvette. Thus, once the cuvette has been detected, we perform a similar analysis and look for holes, which are compact regions of the image after thresholding for a particular value.

Wall gravel

A radically different approach for detecting sewer damages than the one explained so far is to look directly into the RGB images that the platform

is gathering. As a testing bed, we have approach the problem of detecting eroded walls for which they might need re-plastering. Examples of this situation are shown in Fig. 7.

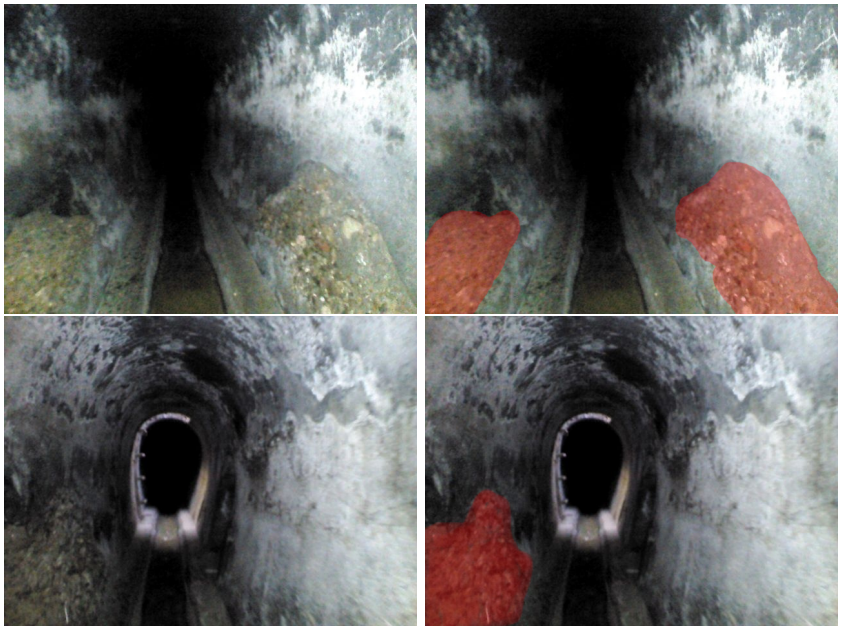

Figure 7: Examples of wall gravel. Left, the RGB image collected by the platform. Right, the region manually marked where wall gravels are present.

In order to detect wall gravels, we have trained a neural network for the problem of semantic segmentation. In particular, we want to have a probability value for each pixel in the image to belong to either the *gravel* category or not. To that end, we have manually labelled 3555 images with positive examples of the class under study and trained a fully convolutional neural network to predict whether a pixel belongs to this class or not.

3.6 Tests and results

Our new 3D reconstruction module and assessment pipeline has been tested on captured data obtained during the different field tests.

Figures 3 to 5 show a comparison between the old reconstruction pipeline and the new one for the Virrei Amat area, whose data was captured during the July evaluation.

Figures 8 to 10 show the reconstructions and distance maps ob-

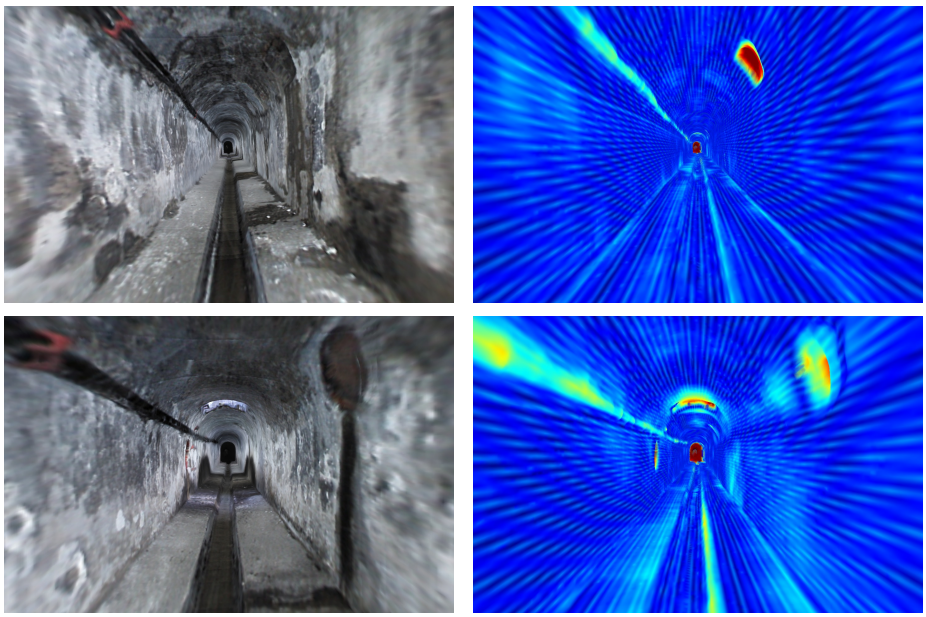

Figure 8: 3D reconstruction (left) and heat map (right) obtained for the Born area.

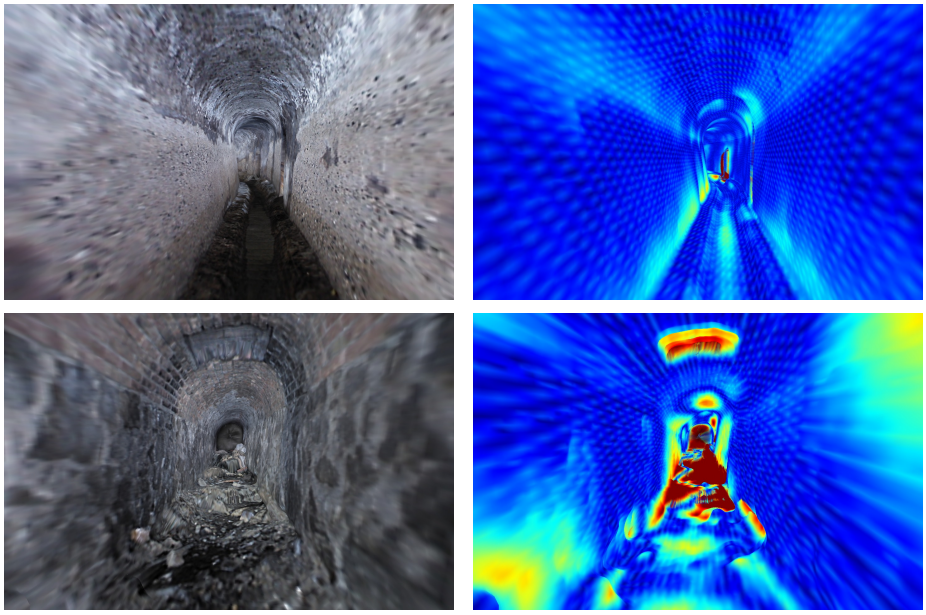

Figure 9: 3D reconstruction (left) and heat map (right) obtained for Av. Pearson (Pedralbes).

Figure 10: 3D reconstruction (left) and heat map (right) obtained for the Forum area.

tained after the field tests conducted at the end of Phase III. As described in 2.4, these belong to the Born area, Av. Pearson and Forum, respectively. The distance maps are here shown as heatmaps, where blue means low differences with regard to the theoretical profile up to red which means big differences. As can be seen, structural elements such as junctions, gutters or manholes are clearly captured within these maps, as well as other elements such as cables, obstacles or defects. The whole process (reconstruction + assessment) took around 10 to 15 minutes for each sewer model.

Incident type	Color
Unknown	Red
Profile change	Yellow
Manholes	Blue
Structural	Orange
Capacity lack	Pink
Obstacle or sediment	White
Holes at cuvette	Green

Table 1: Color legend for reported incidents

Figure 11 show the full maps/textures obtained for the different

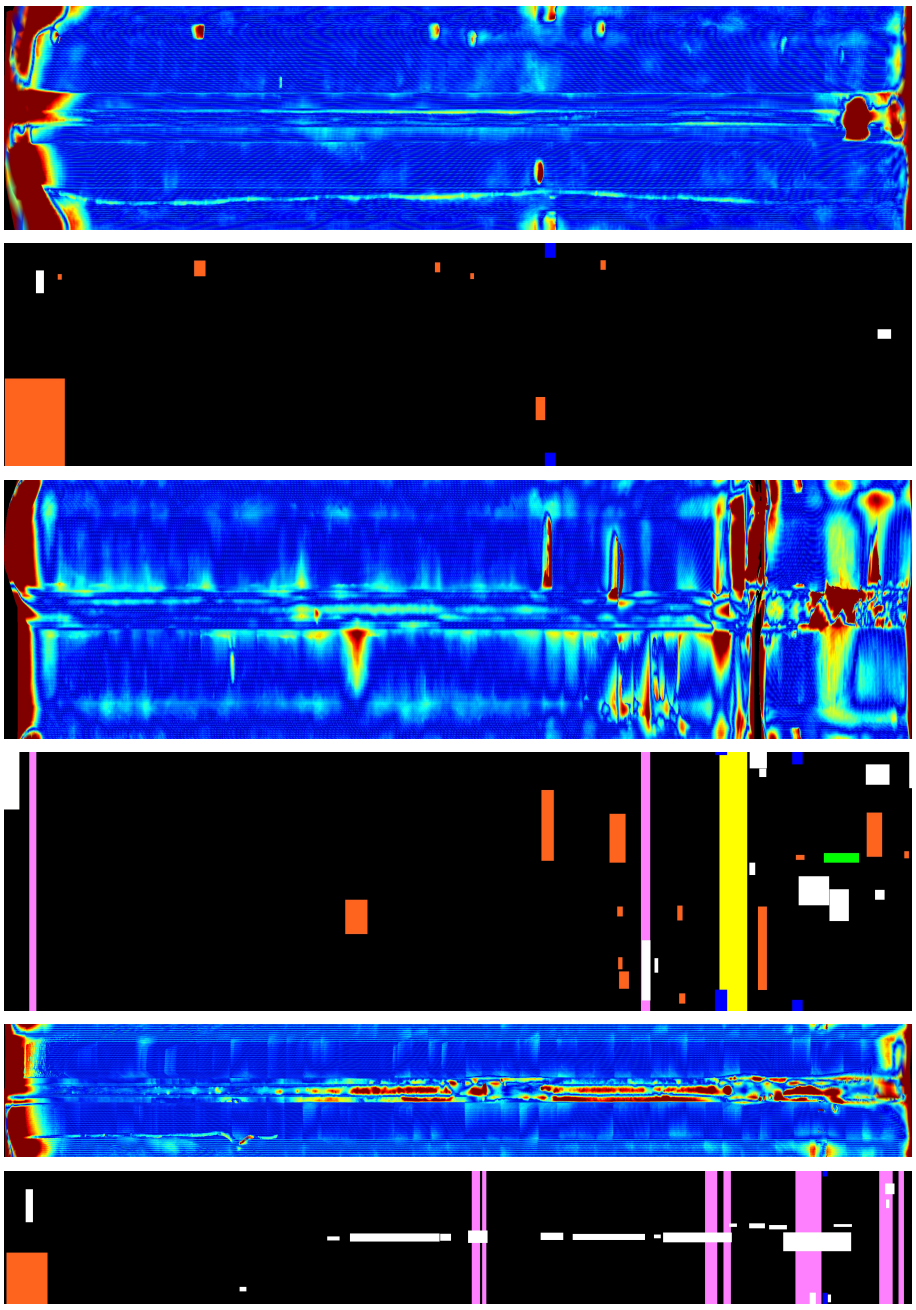

Figure 11: Heat maps and detection results obtained for the different sewer models: Born (top), Av. Pearson (middle), and Forum (bottom).

sewer areas and the corresponding detection of structural elements and obstacles computed afterwards. The color legend is as in Table 1. The Born's sewer area mainly contains structural elements such as junctions and gutters (in red) as well as a manhole (in blue). In Pearson's, a profile change is detected at the end of the sewer (in yellow). Several obstacles and holes are also detected at the end of the sewer due to the accumulation of debris and the lack of an exact theoretical profile. Finally, in the Forum sewer, the accumulated water results in a set of detected obstacles (in white) and capacity lacks (in pink).

Figure 12: Comparison between images captured by the MAV (left) and the obtained reconstruction of the sewer model (right).

Figure 12 finally compares the obtained reconstruction of a sewer with the captured images by the MAV. One can observe how the models obtained with the new reconstruction and texturing approach almost matches the quality of the image, leaving aside changes on the lighting which are corrected during the texturing stage.

Wall gravels

The resulting heatmaps after training the fully convolutional network for wall gravel segmentation are somewhat noisy (see Fig. 13).

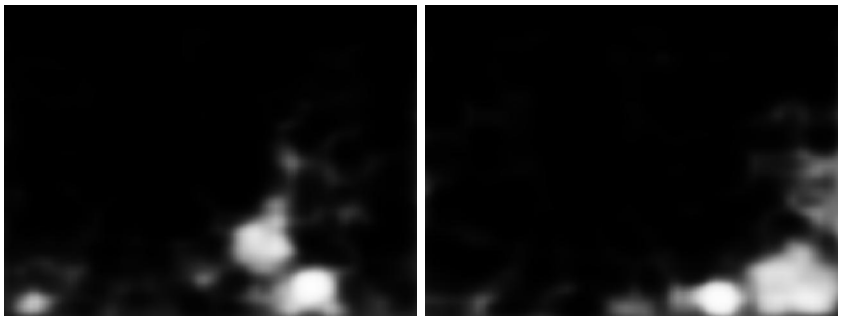

Figure 13: Examples of heatmap detection for wall gravel produced by the fully convolutional neural network.

However, as shown in Fig. 14, a simple processing for region compactness after probability thresholding and temporal consistency over multiple frames allows a successful detection of these regions. In particular, we keep only pixel probability values over a certain threshold (0.5 in our experiments) and ask for compact regions that are detected for at least 10 consecutive frames.

The training and testing data sets for learning the semantic segmentation network are based on four different recording of four different sewers. Three of them are used for training and the last one is used for validation purposes. The four sewers had similar conditions. We plan to collect more data with different conditions and re-train the neural network in order to be able to generalize better to unseen sewer conditions. In particular, we have tested other sewer with significantly different wall conditions —lighter background of the gravel, smaller and more separated stones— and the results suggest for a larger collection of images for training the model.

4 User interfaces

In the following sections we present the results of our effort to develop an intuitive and practical user interfaces allowing inspections brigades to plan and execute sewer inspections using the ARSI MAV, and to analyze the inspection data to generate informative reports for their clients.

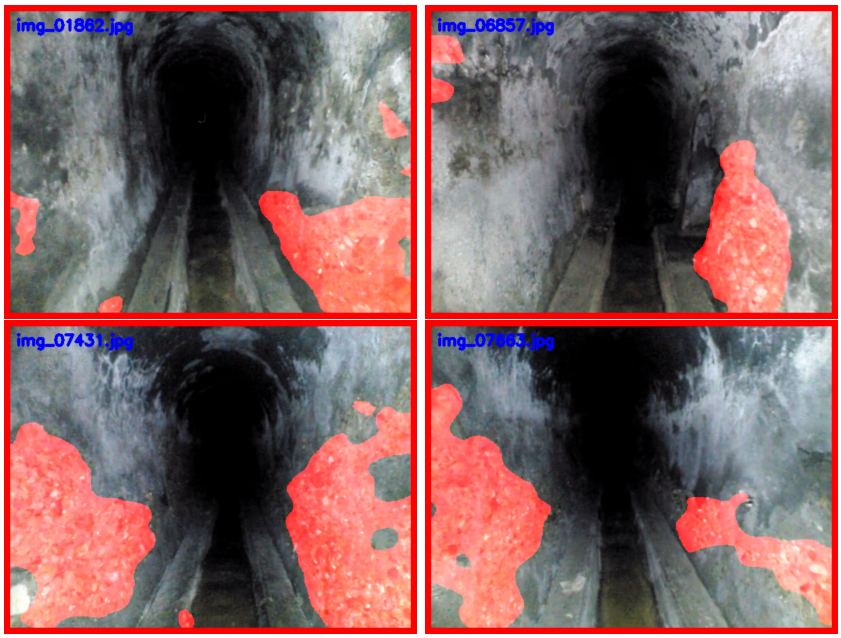

Figure 14: Wall gravel detection results after region compactness and temporal consistency.

4.1 Mission Planner

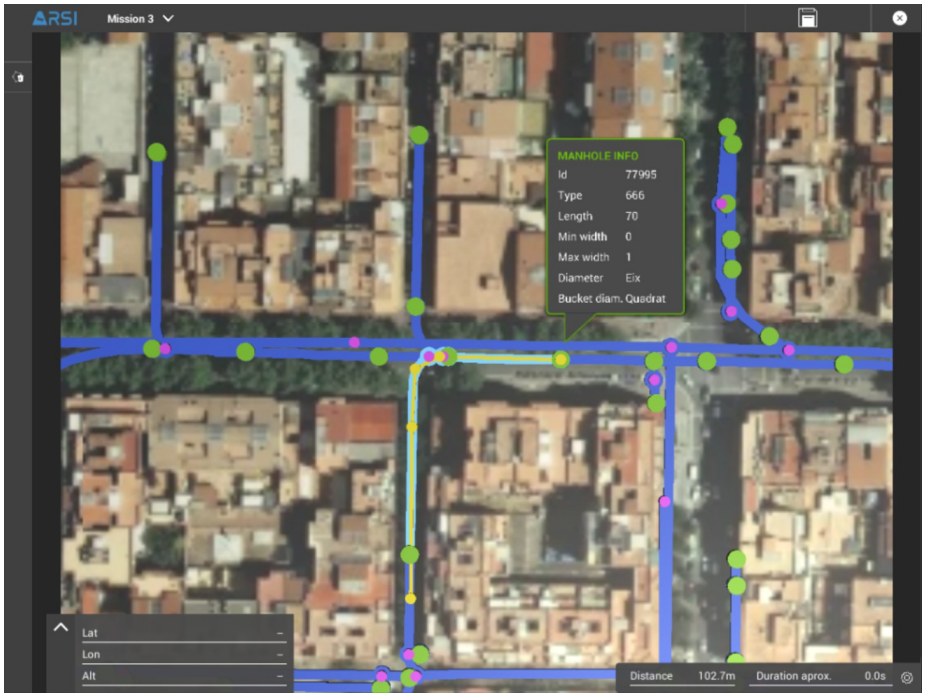

Figure 15: ARSI Mission Planner interface

A Mission Planner interface (see figure 15) was developed during phase III. It allows ARSI users to load GIS data and satellite imagery for the area selected by the client, and to plan a series of flights with the ARSI MAV in order to collect inspection data for the relevant sewer sections.

This tool was used extensively in numerous field tests in the sewers throughout Phase III, and demonstrated at the July 2018 evaluation in Virrei Amat. Please refer to [2] for a detailed description of the software capabilities.

4.2 Operator Console

The Operator Console is the user interface used by inspection brigades on the field to carry out sewer inspections. It was significantly upgraded in phase III: the interface used for the July evaluation featured GIS data display of the sewer networks along with satellite imagery (Google Maps or other map sources such as Mapbox or OpenStreetMaps). Operators can load ARSI missions into the Console, display them on the map and

query the GIS data (eg. manhole depth, type of sections, etc.) to plan operations before execution.

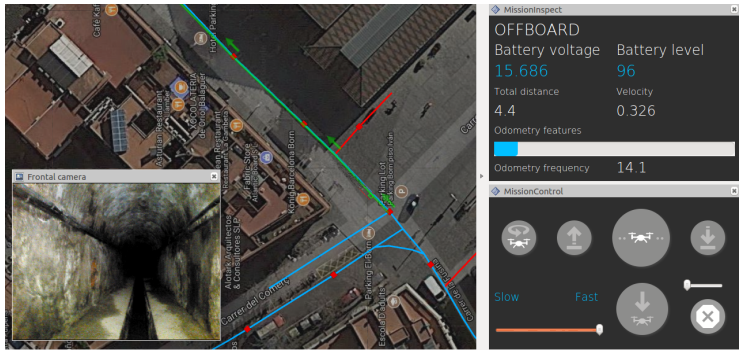

Figure 16: ARSI Operator Console with GIS and map display, live video feedback, and control panels

Before each inspection, the Operator Console connects to the MAV over WiFi to perform flight checks, notifying users of any sensor or system failure. New flight tests were added since the July evaluation, in particular safety checks to ensure that the Pixhawk autpilot is in the correct flight mode, as well as a validation of the ground ranger measurements while on the ground, to reduce the risk of rough landings.

The flight plan is then sent to the MAV for validation, and executed from our control panels by issuing simple high-level commands such as ”start mission”, ”pause”, ”land”, ”return to initial point”, etc. The Console provides live video feedback from the MAV and displays the mission trajectory on the map as well as sensor data. The Operator console was designed to be intuitive and requires only basic training, so that operators without any MAV piloting experience can use the ARSI system.

4.3 Data Analysis Interface

The ARSI Data Analysis tool is the user interface that provides an integrated visual representation for completed missions, combining the data collected by the drones and the post-processed information generated by the data processing components.

In the current version, in addition to the improvements achieved in the 3D reconstruction phase, this tool also uses new visualization techniques to increase the realism of the 3D model representation. The lighting system has been adjusted and a new visibility effect has been included to simulate in a more accurate way the drone lighting condi-

tions. Both lighting level and visibility level can be manually regulated by the operator from the Data Analysis tool.

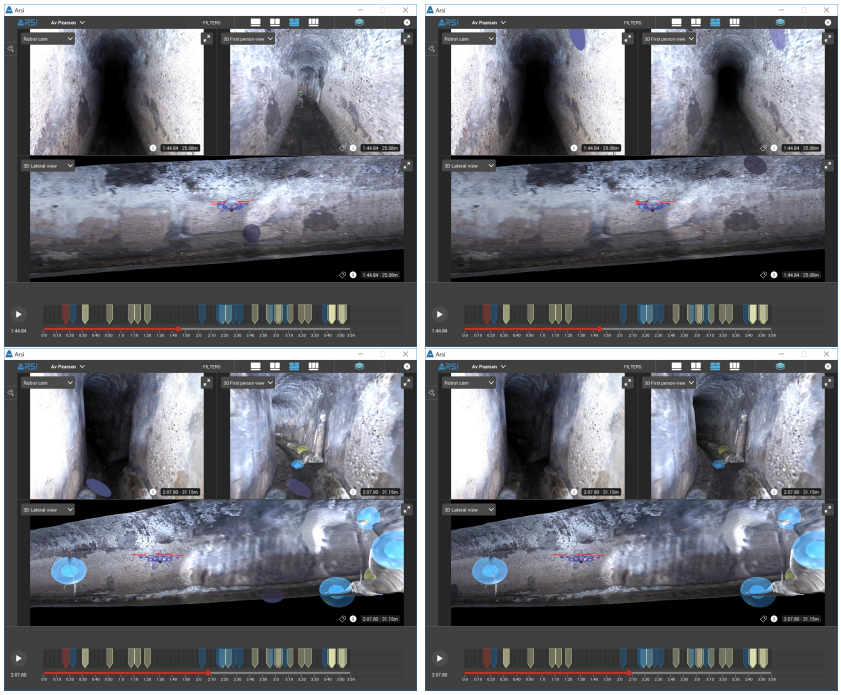

Figure 17: Drone lighting conditions simulation. Left, base reconstruction, Right, same reconstruction with the visibility effect enabled.

Also a new combined visualization mode has been added, which blends the original default mode with the heatmap mode included in the previous version. The combined visualization mode provides an intuitive way to recognize defects and potential conflictive areas, but maintaining the real image details captured by the drone cameras.

In addition to the improvements in the representation of defects that were included in the previous version, the tool now also gives the operator a more effective way to manually mark new incidences. Incidences can be reported at any point of the 3D reconstruction surface (pointing over the model), and not just using a timeline based approach.

On the other hand, adjustments have been applied in the rendering of the 3D model and in the visualization of video, in an effort to increase the system performance and to allow a better user experience. One of the most remarkable advances in this version is the integration of Virtual Reality devices. Now the ARSI Data Analysis tool allows immersive navigation through the 3D reconstruction using a Virtual Realty headset (Oculus Rift).

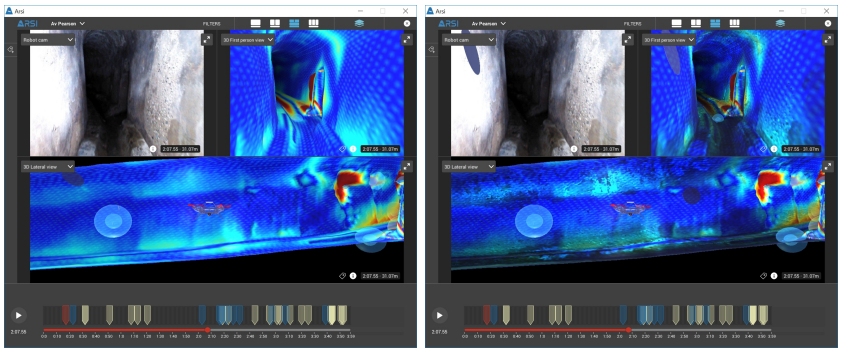

Figure 18: Tool visualization modes. Left, heatmap visualization mode, Right, combined visualization mode.

- The drone navigation can be controlled by using the VR touch controller. The tool allow bidirectional navigation with adaptable speed. The operator can also travel to a specific place just pointing it with the VR controller in a similar way to the desktop version.
- New defects can be manually marked in the virtual environment just pointing them with the VR controller and choosing the detected defect type using a VR user interface.
- Detected defects can be easily handled in order to change their current status (confirmation of defects, to mark them as false positives or just to remove them).
- Relevant visualization options can be controlled through the touch controllers (like show or hide the highlight of incidences or change the visualization mode)
- A new VR view is also available from the desktop user interface in order to obtain an approximated representation of what the operator is seeing through the VR headset.

Figure 19: ARSI Data Analysis tool working in immersive mode.

Figure 20: Integrated VR devices. Left, Oculus Rift headset, Right, Oculus Touch Controller.

Tests and test results

When a sewer inspection is completed and all the post-processed data has been generated, the resulting mission data structure can be imported into the ARSI Data Analysis Tool. Every carried out mission has

been used to test the ARSI Data Analysis tool, with the aim to check, in an iterative way, that the tool is working correctly after any change or improvement implemented in each part of entire ARSI workflow.

5 Conclusions and future work

In this report we detailed the work carried out in the second part of Phase III of the ARSI project, which was focused on three main tasks:

- Improving the MAV platform capabilities, in particular robustness and flight stability (section 2);
- Completing the development of the 3D reconstruction and defect detection algorithms, and validating them using datasets collected in various tests sites in Barcelona (section 3);
- Making inspection data analysis more intuitive, in particular through the introduction of an immersive mode in our Data Analysis App (section 4).

We believe that the significant improvements made to the 3D reconstruction, defect detection and data analysis components will boost the usability and marketability of the ARSI system as a practical and efficient tool for sewer inspection.

We put considerable effort into analyzing and addressing the flight stability issues observed at the July evaluation in Virrei Amat, Barcelona. This effort lead to important changes to the MAV propulsion system and autopilot firmware, as detailed in section 2. While these efforts produced good results during our field tests in various locations (Mercat del Born, Av. Pearson and Forum), we recognize that the MAV did not perform well at the final evaluation in December 2018.

Since our initial analysis did not reveal the cause of the issues seen in the evaluation, we are committed to carrying out a thorough investigation using all data available, in particular autopilot and onboard computer logs. We hope that our findings will allow us to identify the relevant software, firmware or hardware changes required to deliver stable and safe flights with our MAV platform, as well as high-quality data for sewer inspection.