Rescue Robots at Earthquake-Hit Mirandola, Italy: a Field Report

Geert-Jan M. Kruijff
DFKI GmbH
Saarbrücken, Germany

Viatcheslav Tretyakov, Thorsten Linder
Fraunhofer IAIS
Sankt Augustin, Germany

Fiora Pirri, Mario Gianni,
Panagiotis Papadakis, Matia Pizzoli, Arnab Sinha
University of Rome La Sapienza
Rome, Italy

Emanuele Pianese, Salvatore Corrao,
Fabrizio Priori, Sergio Febrini, Sandro Angeletti
Vigili del Fuoco
Rome, Italy

Abstract — In May 2012, two major earthquakes occurred in the Emilia-Romagna region, Northern Italy, followed by further aftershocks and earthquakes in June 2012. This sequence of earthquakes and shocks caused multiple casualties, and widespread damage to numerous historical buildings in the region. The Italian National Fire Corps deployed disaster response and recovery of people and buildings. In June 2012, they requested the aid of the EU-funded project NIFTi, to assess damage to historical buildings, and cultural artifacts located therein. To this end, NIFTi deployed a team of humans and robots (UGV, UAV) in the red-area of Mirandola, Emilia-Romagna, from Tuesday July 24 until Friday July 27, 2012. The team worked closely together with the members of the Italian National Fire Corps involved in the red area. This paper describes the deployment, and experience.

Keywords: *Robot-assisted USAR, human-robot interaction, human-robot teams, situation awareness*

I. INTRODUCTION

On May 20 2012, in the middle of the night, northern Italy was hit by an earthquake with epicenter in Finale Emilia, in the region of Emilia Romagna [12]. On May 29 at 09:00 AM local time, a 5.8 magnitude earthquake struck the already damaged area again. Overall, 246 seismic events with magnitudes between 3 and 6.1 occurred from May 20 until June 18 within a radius of 50km of the original epicenter, see Fig. 1 and [12], and affected some 900,000 people across six provinces, with a rich cultural heritage.

The National Fire Corps (CNVVF), the Italian Department of firefighters, public rescue and civil defense, has been in charge of disaster response and recovery for the area. The undertaking involves a national effort to ensure search and rescue of people, evacuation of several centers and recovery of valuable works of art, and to secure many of the damaged buildings and surveying them. For intervening in the main devastated churches and historical buildings, some of which were not accessible even by the CNVVF, and for art-works recovery, it was necessary to assess location and state of risks.

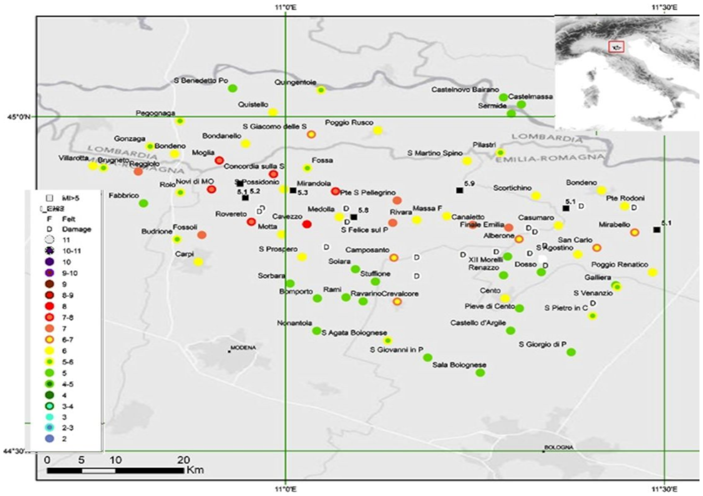

Fig. 1. Final Map of the intensities of the seismic sequence up to the 3rd of June, in the Emilia Romagna region; black squares denote quakes intensities $M > 5.0$. Source: *INGV*.

This is where NIFTi comes in. NIFTi is an EU-funded project, focusing on human-robot teams for exploring disaster sites in Urban Search & Rescue settings [3], `http://www.nifti.eu`. NIFTi adopts a user-centric approach to developing its models, from autonomous robot behavior (UGV, UAV) to human-robot collaboration, working together with several end user organizations. With the end users NIFTi sets up requirements, experiments with prototypes, and evaluates overall systems performance on a yearly basis. The CNVVF is one of the project partners involving both the *Instituto Superiore Antincendi* (ISA) and the *Scuola di Formazione Operativa* (SFO) in Montelibretti. Since the beginning of the project in early 2010, this meant that the CNVVF has built up a close working relationship with the research partners in NIFTi. Having this experience with the systems developed by the project, and knowing the potentials of the work done, the CNVVF has requested NIFTi to aid in structure damage assessment by deploying a team of humans, UGVs, and UAVs in one of the most damaged towns of the whole region, namely Mirandola, by entering the red area and intervening at the Duomo and at the *San Francesco* church; here, together with his dynasty, Pico della Mirandola is buried. Pico was author of *On the Dignity of Man* [2] considered the *Manifesto of the Renaissance*, and known for his prodigious memory.

This paper describes the experience gained in the deployment, in July 2012 and it is organized as follows. §II describes the sites of Mirandola red area where the team ran missions.

§III outlines the team structure, the robots, and the infrastructure used in the deployment. §IV describes various aspects of the experience gained, in human-robot team workflow, and UGV and UAV technology.

Robots have been deployed in real-life disasters before. Robin Murphy and her colleagues have been in the field with a wide variety of robots, aiding first responders across the world (cf. e.g. [8]–[10]), and Satoshi Tadokoro and his team recently conducted (and concluded) a long-term deployment at the Fukushima nuclear power station [7], [15], to cite the main contributions which, however took place in US and Japan. Within Europe, the only officially requested involvement of robots during a disaster we are aware of was at the Cologne city archive collapse in 2010 [6] – without running missions though, the robots were held on standby. The deployment described in this paper can thus (presumably) be seen as the first deployment of a large human-robot team in Europe, fielding multiple types of robots.

II. SITES

During the deployment, the team surveyed two sites in Mirandola red area: San Francesco church and the Cathedral (Duomo).

San Francesco church dates back to the thirteenth century, and is one of the first Franciscan churches in Italy. It houses the suspended arks (sarcophagi) of the Pico family, who ruled the Duchy for four centuries (1310-1711). Severely damaged during the earthquakes, only the façade and some of the walls are still standing. The lateral nave, where are the Pico's arks, is of particular cultural importance, and the Italian Cultural Heritage wanted to recover them, though the Gothic vaults were very dangerous. Indeed, the central nave and the eastern aisle are mostly destroyed, as the bell tower collapsed over the church roof. Fig. 2 (l.) shows the nave and a bit of the eastern aisle (in the back): A rubble heap by and large inaccessible to our UGV. The western aisle is still somewhat intact, but structurally highly unstable, see Fig. 2 (r.). The ceilings within most of the vaults are damaged, having holes and large unstable pieces of masonry about to come down. Throughout the aisle, there are rubble heaps with larger pieces of masonry fallen from the ceilings. At the end of the western aisle there is a heap of rubble due to a roof cave-in. The western aisle was accessible through the lateral coffered door. Size-wise, the western aisle was about 35 meters deep, and 6 meters wide.

Fig. 2. San Francesco nave (l.) and western aisle (r.).

Our mission targets for San Francesco church were to assess structural damage to pillars and the vaults, to identify passages to the altar so as to recover paintings and to record the state of the Pico's arks and the vaults over them. We ran missions at the church on Tuesday July 24 and Wednesday July 25, for a total of 5 UAV flights (27 minutes) and 2 UGV runs (1:05h).

The Duomo is a large cathedral finished in the 1470. Also here, the earthquake caused substantial damage. The façade of the cathedral, including a large clock, has largely fallen down – blocking access to the cathedral through its main entrances. The roof over the nave and the northern aisle has caved in, causing massive damage, as illustrated in Fig. 3 (l.). The bell tower of the Duomo is still standing, though structurally severely damaged. After consulting with the local CNVVF commander, and the cathedral's padre vicarious, we managed to get access to the Duomo through the vicarage. This provided direct entry to the southern aisle, still intact though with damage to pillars and to the cross vaults, and another entry in face of the chorus. The southern aisle, shown in Fig. 3 (r.), proved to be relatively easy to traverse for the UGV, much more difficult was to reach the S.S. Sacrament chapel at the end of the northern aisle, were the painting of Sante Peranda was held. The top of the nave was covered by a large rubble heap of masonry from the roof and supporting structures.

Fig. 3. The Duomo in Mirandola, nave (l.) and entrance to the S.S. Sacramento chapel behind the altar (r.).

For the Duomo, we again helped assessing structural damage, outside (bell tower) and inside to access the S.S. Sacramento chapel, to report on the state of the highly valuable paintings. Further missions were assessed to establish the state of the two wooden ancons covered in gold, although the one on the northern aisle could not be clearly assessed as the UGV had to stop at the exit of the S.S. Sacramento chapel towards the northern aisle, because of the very high heap of rubble. The chapel had been reached from the door accessing the back of the altar, though almost the whole vault was collapsed. We ran several missions at the cathedral on Thursday July 26, for a total of 4 UAV flights (15 minutes) and 3 UGV missions (about 1:20h).

III. DEPLOYMENT

The entire setup deployed in Mirandola has been developed within NIFTi, modulo the basic robot middleware (ROS). We deployed a subset of the available NIFTi functionalities, described in more detail in [3]. We focused on robust functionalities for robot control, video streaming from different omni-directional and monocular cameras, and laser-based 3D reconstruction of the environment, coupled to the NIFTi multimodal Operational Control Unit (OCU).

A. System & network infrastructure

As middleware we run the *Robot Operating System* (ROS) [13]. ROS is used for running processes on the robot, streaming data over WiFi to an operator control unit (OCU) and other visualization tools (RViz), and for logging purposes (`rosbag`'s). Off-board computers were used for processing 3D laser range data (point clouds), and for the OCU and visualization. We used a mixture of laptops, monitors, and a desktop computer.

We use a 2.4GHz WiFi network. We set up an antenna nearby the entrance to the actual deployment area, fixed to a tripod, and connected by ethernet cable to a router and DHCP server in the command post. The antenna is 50cm long, has 14dBi gain, and is extended with a Ubiquiti high power bullet enabling a transmission power of maximally 28dBm. Each robot (UGV and UAV alike) is also equipped with a bullet, and an omnidirectional rod antenna with a 9dBi gain. As we were mostly working in large open spaces, we did not experience substantial problems with network coverage.

Throughout the deployment, electricity was provided by a Honda 20i portable power generator, provided by the CNVVF. The power generator was capable of generating 230V ($\pm 1\%$) – provided it had enough fuel. Occasionally, fuel would run out causing a shut down of the desktop computer and the monitors, though fortunately never during a mission.[1]

All of these systems worked reliably, in outside working temperatures between 35 to 40 degrees centigrade, and dusty conditions. The gazebo-style roofing over the command post protected the staff, monitors, computers, and other equipment from direct sunlight.

B. UGV

We deployed two NIFTi UGV platforms in Mirandola: One as main system, and one as back-up should something go wrong. Fig. 4 shows the UGV platform used.

Fig. 4. NIFTi UGV with a rotating SICK-Laser (LMS100), a LadyBug3 omnicam, active flippers and active/passive bogeys, IMU, GPS, and a static mast mounting a PTU with a Kinect sensor.

The UGV platform has been developed within NIFTi, in close collaboration between research partners and end users. The platform provides a mixture of passive and active morphological adaptivity, to provide good mobility even in harsh terrain. Its bogeys are connected by a differential, allowing for passive adaptivity (following terrain contour) and active configuration (by blocking the differential). The four flippers can be independently controlled. The bogeys and flippers are constructed such that the platform has a ground clearance of over 15cm (which proved to be very useful in crossing "complex" rubble). The platform can traverse a wide variety of terrain, and per requirement climb up to 45 degrees; practice has shown that we can climb up 60 to 70 degrees inclines. Size-wise, the platform weighs in at about 25kg, and is airline compatible ($L + W + H < 158cm$). The entire body is IP53, it can drive through puddles (IP57), and is conceived for operating temperatures between -10 and 40 degrees centigrade.

The platform comes with a basic sensor suite consisting of a rotating laser (SICK LMS100) mounted in front of the robot, a LadbyBug3 omnidirectional camera mounted on top of the robot, as well as an IMU and a GPS sensor. In addition, we mounted a $25cm$-tall static mast on the battery compartment of the robot. On top of the mast was a pan-tilt unit with a Kinect camera. This provides a chase-style view of the robot, which is highly useful when navigating (tele-operating) the robot in tight or complex spaces – cf. also the recent experience with Quince reported in [15].

The platform has a quad-core on-board computer (mini-ITX type motherboard), and is powered by a battery pack with an operating time of between 2 and 4 hours. During the deployment we had the robot running all day long, only requiring a battery recharge in the evening.

C. UAV

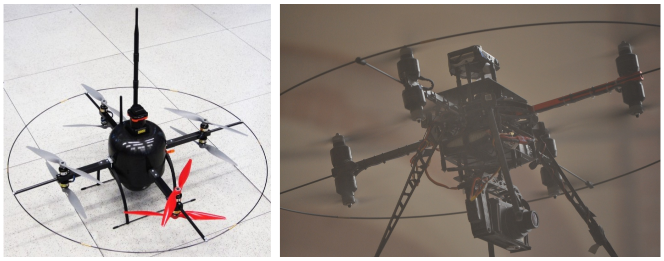

Fig. 5. Standard NIFTi UAV octocopter with a standard configuration (l.) and NIFTi UAV research octocopter (r.) with a mounted on top camera, and a camera in a tilt unit under the main body. Another configuration flown includes a PC and a Kinect-style sensor mounted on top of the research UAV.

Two different types of UAVs were prepared for the mission (see Fig. 5). The first platform has been developed within NIFTi according to requirements from end-user and research partners. It has a housing to protect from rain and dust, an interchangeable lower sensor compartment, and removable and easily exchangeable arms with motors. Disassembled it can fit into standard airline baggage. The UAV's standard configuration includes a Hokuyo UTM-30LX range-finder mounted on top, a sonar and a pressure sensor-based altimeter, an IMU module, a 3D magnetic compass, GPS module, and two cameras - one forward-looking (15 degrees tilted) and downward-looking. The on-board computer is a 1.6GHz Intel Atom-based PC. Operation time is between 10 and 15 minutes depending on operational mode.

The second vehicle is a research platform, based on a construction kit. It has eight high-power engines and can lift

---

[1] In retrospect, a sufficiently powerful UPS back-up would have been useful.

up to 2 kg additional payload. It has neither rain- nor dust protection, but it provides a high level of configurability. It can carry a high resolution camera on an IMU-stabilized tilt unit that can be controlled by the pilot, and it can be deployed with all the sensors described for the first platform. The camera is equipped with a high-power video signal transmitter (5.8 GHz, 1.5W). Instead of the Hokuyo laser scanner, a Kinect-like camera (ASUS Xtion Pro) can be mounted on top of the platform; see Fig. 6 for example output data. The vehicle is equipped with an Intel i7 2.6 GHz based single-board PC. All components can be easily removed or additional are added.

Fig. 6. 3D reconstruction of one of the Pico's arks in San Francesco church, using data from the NIFTi UAV Kinect sensor.

Because of the flexible configuration of the second platform, and the functionality it could thus make available, it was used as the primary UAV, with the first drone as backup. The UAV had to be navigated in a GPS-denied environment, turned on spots to acquire a better view, and the pilot could only remain on one spot next to the entrance into the area being surveyed. Crucial in this case was the UAV's ability to memorize its original orientation in space, and to translate the pilot's movement control commands relative to his position. This functionality was only tested on this UAV prior the mission.

D. Human-robot teaming

Before the actual deployment we set up an organizational- and communications structure for the team to be deployed. The organizational structure is based on previous experience [3], [8], and aims to set up chains of command which reflect *responsibility* (and ultimately, liability). The overall responsibility for each mission lays with the CNVVF, determining mission targets is resolved between NIFTi, the CNVVF and Cultural Heritage responsible, watching over missions targets.

Fig. 7 shows the structure, tied to *roles*. In practice, a single person can play multiple roles. The Mission Cmd (CNVVF) is in charge of the entire mission.

The UAV team largely deploys in the field. The UAV team consists of the UAV Operator, piloting the UAV in-field; a UAV Mission Spc, watching the UAV video streams and guiding the UAV Operator to mission targets; and a Safety Cmd (CNVVF) safeguarding the UAV team. During the deployment, the UAV Mission Spc mostly operated with the UAV Operator in-field, to provide the UAV Operator with an extra pair of eyes on the UAV. See Fig. 8. The UAV team assessed video material afterwards. The information gained from video material was provided directly to the CNVVF, and was also used in the briefings for follow-up UGV missions.

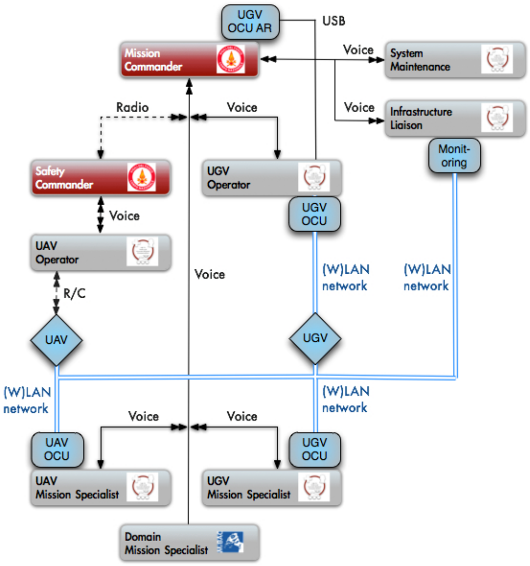

Fig. 7. Organizational structure for geographically distributed human-robot team, including a UGV and a UAV, and reflecting responsibility/charges in the chain of command.

Fig. 8. UAV team with an Operator (l.), Mission Spc (m.) and Safety Cmd (r.) operating near the Duomo, assessing damage to the Duomo bell tower.

Similarly, the UGV team consists of a UGV Operator, and a UGV Mission Spc. The UGV team is located "remotely" in a command post. Both teams are backed up by a System Spc and an Infrastructure Spc operating from the command post, who ensure the network- and system infrastructure remains alive. The intention was for the UGV and UAV Mission Spcs to collaborate with a domain expert, to establish mission targets before and possibly during the mission. During the deployment, however, mission targets were always determined beforehand, together with members of the CNVVF and others involved in recovery (vicarious and cultural heritage representative). The UGV team therefore "reduced" to a UGV Operator, Mission Spc, and a person doubling as Infrastructure/System Spc.

The UGV team uses the NIFTi OCU and RViz, to display video streams, and visualize incoming sensor information. The UGV is tele-operated using a gamepad (or alternatively, the OCU), and the PTU on the mast is controlled using a simple widget. The UAV team uses an R/C control to pilot the UAV, and can use an instance of the NIFTi OCU to watch streaming video from the UAV. Both teams, and the Mission Cmd, have the possibility to use Augmented Reality goggles to watch video streams. In the end, the UAV Mission Spc used this option to great effect. Given that the UGV team relies heavily

Fig. 9. UGV team in yellow helmets, with an Infrastructure/System Spc (l.), Mission Spc (m.) and Operator (r.) operating from the command post near San Francesco church.

on face-to-face communication, the goggles were impractical.

IV. EXPERIENCE

Below we describe first-hand experience with deploying complex human-robot teams in an earthquake-hit disaster area.

A. Human-robot teaming

Deploying robots in an Urban Search & Rescue mission is a *team effort*. This holds just as much for operating a UGV, as it does for flying a UAV. There is too much information to attend to, ranging from sensor information to system- and infrastructure-related monitoring, to be handled safely by a single person, cf. also [8]. The organizational structure shown in Fig. 7 reflects this, and as we already indicated in §III, this structure was pretty much implemented "as-is" in the field.

Nevertheless, both the UGV Operator and the UAV Operator suffered from cognitive overload. UGV missions typically lasted about half an hour, and were characterized by interleaving driving, and observing. (At the moment we cannot yet say whether the ratio reflects earlier experience as reported e.g. in [1], [5].) This interleaving made it possible for the UGV Operator to relax, momentarily – a luxury the UAV Operator did not have. The UAV did have some degree of autonomous flight control, but circumstances demanded that the UAV Operator continuously attended to the UAV.

This provides a first insight in, or rather perspective on, possible roles of "robot autonomy." In human-robot teams, humans and robots are (inherently) interdependent [4]. Robots can go where humans need to but cannot, whereas humans can aid robots in better understanding and operating in the environment. Both humans and robots are problem-holders – with the obvious "but" though that *the human users are the stake-holders*. Robot autonomy is ultimately to be in service of the human user, to reduce cognitive load (improved autonomous navigation, sensor data interpretation, collaborative decision making) and to improve the possibility for the human to collaborate with the robot as if "operating the world rather than the robot" [11]. We saw this over and again during the deployment, see also §IV-C and §IV-B: Autonomy is to make life easier for the human to understand the environment, (and not for the robot to bugger off on its own in "look ma no hands" mode).

The UAV serves as a good example here. The UAV Mission Spc used augmented reality eyewear (*Vuzix WRAP 920AR+*) to watch the video stream from the camera mounted in a tilt-unit under the UAV. This quickly led to a pseudo-immersive experience, and the desire to look left-and-right and have the UAV and/or the tilt-unit follow suit. More (and better) flight control autonomy, enabling the UAV to simply hover and turn on the spot, would have facilitated this.

Further insights concern the flow of information between the UAV team and the UGV team, in terms of tactical (team-level) situation awareness (tacSA) and mission planning. During the entire deployment, the UAV team and the UGV team never operated in the same area simultaneously. Partly, the reasons were technical (network) and environmental (dust). Another reason regarded the *use*, the workflow which emerged in using information from the different teams in establishing further missions. Based on in-field LOS observations of the area to be deployed in, and a first set of recon missions by the UAV team, we would establish a first sketch of the environment. Most importantly, we would identify important landmarks to navigate by, establishing explicit names for them (e.g. "column 4"), and determining targets for future missions. These targets typically included areas and objects to be observed, and how these observations were to be made. Targets were discussed together with members of the CNVVF.

Follow-up missions then helped detail out tacSA and revise mission targets. Since tacSA was built up from operational SA coming from the different teams, we occasionally found mismatches in expectations which then required further missions; (as was to be expected, cf. [14]). For example, video from initial UAV recon missions in San Francesco church gave the impression that the top of nave would be reachable from the western aisle, either from between the fourth and fifth columns, or the opening behind that. This would then make it possible for the UGV to drive close to the altar, and provide close-up video. As it turned out at the end of the second UGV mission, what seemed accessible terrain from the viewpoint of the UAV, was not so in UGV-reality. The UGV did manage to take video of the altar, but an additional mission was then planned for the UAV to fly in over the main nave and record video from that viewpoint.

The UAV and the UGV thus supported each other, but indirectly so. It did result in the required tacSA for the team, and the other stake-holders. At the same time, it also opened new questions as for how to optimally transfer data from one mission to the next, to make the tacSA consolidated so far available online. Before the deployment, we had developed a basic viewer for post-mission analysis. During a mission, a Mission Spc could take snapshots in an OCU, annotate them with a description. Snapshots were stored with the text annotation and robot position information. For post-mission analysis, the viewer could then load snapshots and a 2D map, mark the snapshots on the map, and enable the user to browse snapshots. We did use some of this functionality, particularly to get high-definition snapshots of cultural artifacts, but what was missing was the possibility to correlate *geo-referenced video* from one mission, and show this during another mission in a *context-aware fashion*, i.e. show previously recorded video of the environment in which the robot in the current mission is located. This is a form of information fusion to provide continuous tacSA across different missions within a single

Fig. 10. 3D reconstruction based on NIFTi UAV data

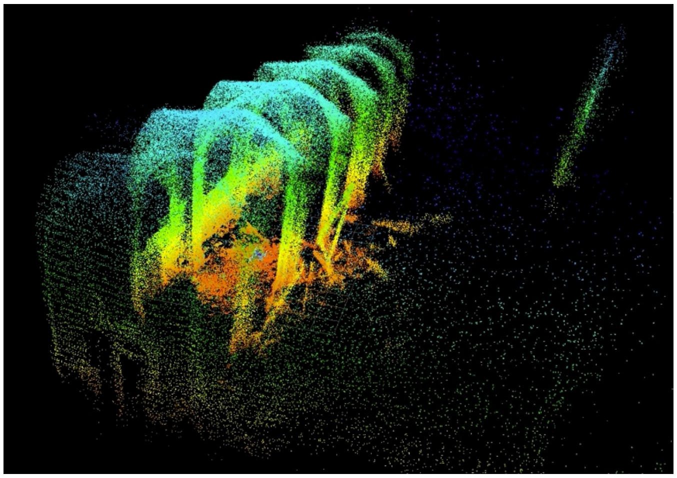

Fig. 11. 3D map constructed by the NIFTi UGV.

area. We made similar observations about map information. The UAV could be deployed to gather a 3D reconstruction of the environment. This map would not need to be so detailed as to enable the UGV to localize itself in it. All the map would need to make possible is a form of *forward mapping/scouting* for the UGV team to determine the optimal path amongst different alternatives. While operating in a harsh environment like the ones in Mirandola we would have greatly benefited from such functionality, as it could have saved time, or have indicated paths where none were obvious (like a traversal from the western aisle to the nave in San Francesco church). See 10 and 11: Coupling the UAV 3D information to the dense 3D metrical representation for the UGV could improve situation awareness for the Operator as well as the robot.

In summary, we observed several issues regarding the operations of a geographically distributed human-robot team, with team members operating both in-field and at a remote command post. As the UGV and UAV teams operated asynchronously, maintaining and transferring tactical situation awareness between missions was an issue to the extent that system automatization could help (in the future) to make aspects of operational situation awareness from one team available to the next in an operational context-aware fashion.

B. UGV

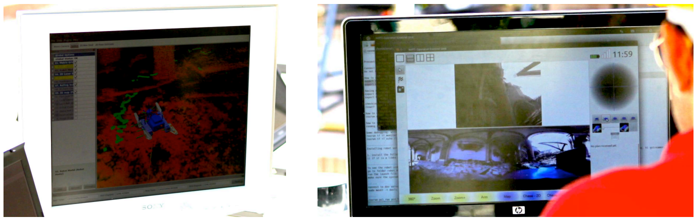

Fig. 12. 3D map with robot model visualized in RViz (l.) and OCU Visualization with Pan-tilt-camera feedback on top and LadyBug3 feedback on bottom (r.)

UGV piloting does not require keeping the LOS on the robot, due to the sophisticated visualization facilities offered by the OCU and RViz, providing the remote operator with situation awareness, as mentioned in §III. The UGV exploration is carried out by alternating two phases. The first phase corresponds to *navigation* and it is functional to terrain understanding: the pilot retrieves information on how to change the robot morphology in order to choose the best route and overcome or avoid obstacles. The second phase is *observation*. This is functional to the collection of data: the Kinect sensor is properly oriented in order to perform data acquisition towards the locations of interest. The UGV Mission controlled the alternation of these two phases and communicated to the UGV Operator the items of interest to inspect and whether to stop the navigation in order to acquire data in details.

Four different flipper configurations have been pre-defined within the OCU to support the navigation phase (e.g. for flat terrain or to approach an obstacle) that are manually adjusted by the UGV Operator according to the information provided by the camera feedbacks, the rotating laser and the robot 3D model within the computed 3D map (Figure 13). The latter was eventually deemed as the main source of information in order to suitably adjust the configuration of the UGV rear flippers. The 3D robot model provides an immediate visual feedback of the current flippers configuration with respect to the terrain, which is displayed on the 3D map. Adequately setting the UGV rear flippers is mandatory to prevent the UGV from tipping-over whenever overcoming rubble or climbing stairs was the case. Continuously switching among the camera feedback and the 3D visualization of the UGV in the map increased the UGV Operator's cognitive load, which increased the time required to accomplish the mission, indeed.

During the *observation* phase, the UGV Operator was guided by the mission specialist to acquire data on the items of interest, by changing the orientation of the pan-tilt camera. This had a negative effect on the awareness of the UGV operator regarding the relation between the current coordinate frame of to the camera orientation and the robot pose. To alleviate this problem, a set of pan-tilt configurations were pre-set in order to allow the UGV Operator to quickly recover the views related to the pre-defined reference poses.

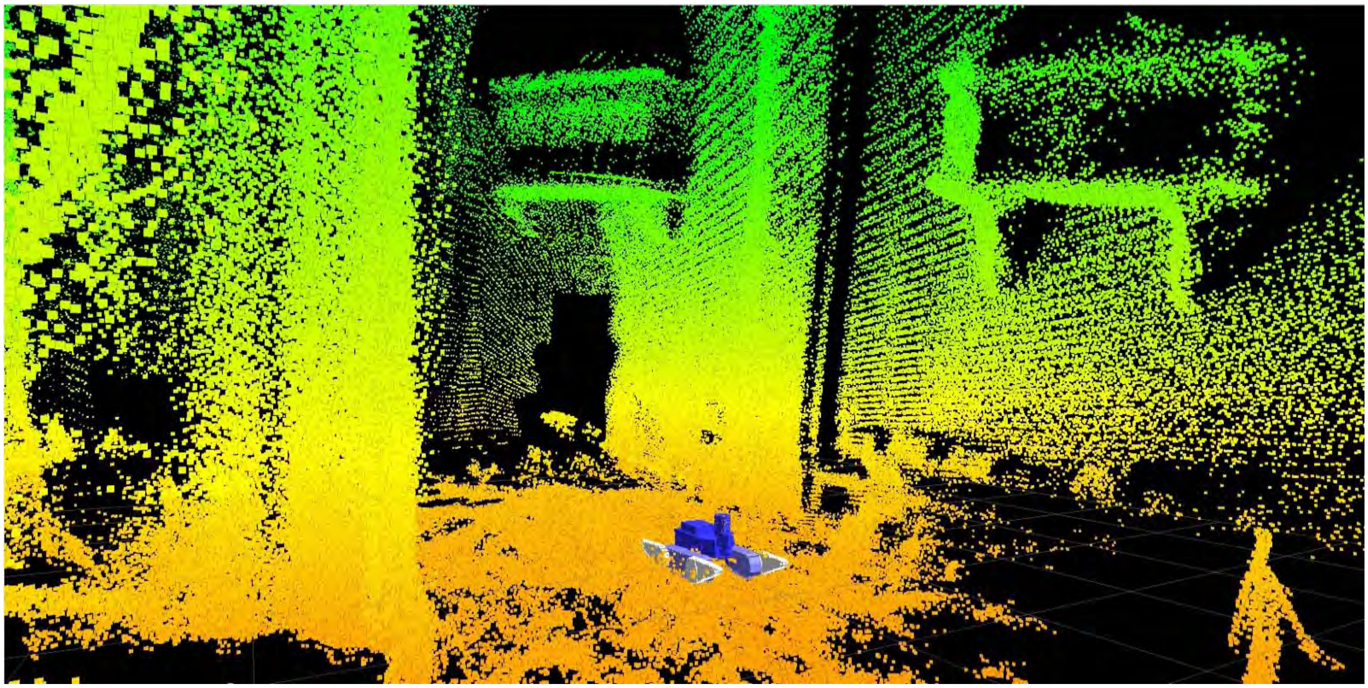

Fig. 13. The UGV model displayed in the current 3D map, illustrating 2 of the Pico's arks.

Within the UGV sensor suite, the pan-tilt camera could not be used to inspect the back of the UGV. When this was required, the stream coming from the LadyBug3 together with the RViz 3D visualization constituted the only source of information about the robot status. Another issue with the pan-tilt-mounted Kinect concerned the dynamic adaptation in response to changes in the lighting conditions. In the interior of San Francesco church the lighting conditions were such that even within a shadow, the Kinect camera provided a qualitatively good output. In the case of the Duomo church, the situation was quite different as the Operator has to drive the UGV inside the church through a corridor with very constrained lighting. In that situation, a staircase had to be traversed and the most valuable information came from the LadyBug3 camera and the 3D scanner.

The above mentioned situation at the Duomo highlighted another important issue, namely the fact that the presence of the static mast increased the risk of tipping over. This was particularly harmful when the UGV was operated to overcome the staircases and the large rubble heap occupying the way to the western aisle.

Nonetheless, according to the UGV Operator, the pan-tilt camera proved to be the highest priority perception instrument and most immediate choice among the sensor feedbacks and visualizations. Such a configuration-adaptive camera turned to be the most valuable sensor in the *navigation* phase because of the easiness in the interpretation of its output but it was also used in the *observation* phase to collect data regarding the surrounding of the UGV.

C. UAV

Piloting a UAV in a confined space is a challenging task by itself, even if the pilot has LOS to the UAV. In addition, it was clear that if the UAV looses control or crashes into an obstacle, we will not be able to retrieve it. All these factors negatively impacted the UAV Operator's stress levels: No mistake was allowed.

The first deployment was in San Francesco church. It was required to provide images of structural damage to arcs and ceilings inside the western gallery of the church as well as to capture the condition of coffins and art, to observe the altar, and the end of the eastern gallery. The only access to the church was provided via a door of the western gallery with a less than 1 by 1 meter surface free of rubble for take-off and landing or via the caved in roof.

Before the actual deployment all UAV functionalities were tested, including GPS-aided hovering next to the church, flying along two galleries of a seminary that was intact, testing turning maneuvers, and testing the range of the video transmitter. After being satisfied with the performance, missions were undertaken.

Each deployment had a clearly identified task, which with a high level of stress and cognitive load helped the pilot to stay concentrated. One example of such a task is to reach the end of the western gallery and look to the right to see if a UGV could traverse there. During the first two missions the Mission Spc was sitting in a distant location from the Operator and viewing the video stream from the on-board camera in goggles (see above). Since the Mission Spc can only observe the area from the video and cannot intervene, and two successful missions were carried out without accident, it was decided that the Mission Spc would better assist the Operator, standing next to him. This was especially needed because of the amount of dust produced by propellers. Without high-contrast dark lenses it is hard to see the UAV. The video from the camera provides a better view but it is impossible to control a flying vehicle in such a confined space with the camera view only. Even having a direct LOS, it is hard to understand how close you are to obstacles or walls, especially if the drone is far away. Our eyes are loosing clear depth perception with long distances.

It is important to mention that even though the above mentioned functionality of translating control commands independent on the UAV orientation was of a great help, it did not function properly. The problem was in that the orientation calculation relied strongly on the on-board magnetic compass. In an open space this worked perfectly, but in the church there were numerous metal bars that influenced the compass thus causing sudden turns of the drone to several degrees. This dramatically increased the cognitive load of the Operator, forcing him to concentrate more on maintaining the UAV's position. Even if each mission lasted in average no longer than 5 minutes, the stress level of the Operator at the end was so high that it could easily provoke making control mistakes. This was clearly seen by the landing maneuvers which later on were no longer performed smoothly, and caused hitting the entrance door several times, breaking several propellers (3 in total).

Another important mission to mention from the point of view of the obtained experience is the structural inspection of the bell tower of the Duomo. The total flight lasted slightly less than 6 minutes while the first 4 minutes appeared to be almost unusable. The reason was simple - the pilot could not see what he was filming. The task was to concentrate on a defined part of the tower in which CNVVF were particularly interested. The pilot had a plan to reach a desired altitude and distance to the tower, fix the UAV by means of GPS, and then manipulate the camera and the copter orientation. After reaching the required altitude the pilot could not achieve a stable hovering of the UAV due to appeared wind. The pilot takes a decision to operate fully manually and approach the point of interest. The task seemed to be accomplished but shortly before landing the pilot noticed that the camera was facing upwards. That means a wrong point of interest was filmed and had to make a second approach.

Summarizing the experience of piloting the UAV for such a mission following conclusions can be made:

- Only functionality that has been tested, and therefore can be trusted, should be used.
- A high level of stress in combination with cognitive load can provoke pilot mistakes.
- A minimum level of autonomy is required even for pure tele-operation, as it provides the possibility for the UAV to autonomously maintain its position when no movement commands are received.
- A better situation awareness for the pilot is of high importance, particularly to provide information about close obstacles, and better depth perception.
- Observation camera control should be independent from the pilot, unless the vehicle is fully autonomous.

V. CONCLUSION

This paper reports on the deployment of the NIFTi human-robot team at the core of sites that were heavily hit by earthquakes at the northern Italy from May until June 2012. Among the distinguishing contributions of this mission, we highlighted the challenges as they were encountered within a USAR scenario of highest realism and the means by which the NIFTi human-robot team in collaboration with the CNVVF managed to jointly and effectively address. Apart from that primary goal, the NIFTi team further managed to assist in the inspection of damaged cultural heritage art-work and altogether, despite the diversity of risks, completed the mission safe and sound for both human and hardware resources. Both the deployment and the experience are reported in the paper, with the aim of disseminating the lessons we learned from the described mission.

ACKNOWLEDGMENT

The research presented in this paper is funded by EU-FP7 *NIFTI* project.

REFERENCES

[1] J.L. Burke, R.R. Murphy, M. Coovert, and D. Riddle. Moonlight in Miami: An ethnographic study of human-robot interaction in USAR. *Human Computer Interaction*, 19((1–2)):85–116, 2004. 5

[2] Giovanni Pico della Mirandola. De hominis dignitate, 1486. 1

[3] G.J.M. Kruijff *et al.* Experience in system design for human-robot teaming in urban search & rescue. In *Proceedings of Field and Service Robotics (FSR) 2012*, Matsushima/Sendai, Japan, 2012. 1, 2, 4

[4] M. Johnson, J.M. Bradshaw, P.J. Feltovich, R.R. Hoffman, C. Jonker, B. van Riemsdijk, and M. Sierhuis. Beyond cooperative robotics: The central role of interdependence in coactive design. *IEEE Intelligent Systems*, pages 81–88, May/June 2011. 5

[5] B. Larochelle, G.J.M. Kruijff, N. Smets, T. Mioch, and P. Groenewegen. Establishing human situation awareness using a multi-modal operator control unit in an urban search & rescue human-robot team. In *Proceedings of the 20th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2011. 5

[6] T. Linder, V. Tretyakov, S. Blumenthal, P. Molitor, D. Holz, R. Murphy, S. Tadokoro, and H. Surmann. Rescue robots at the collapse of the municipal archive of cologne city: A field report. In *Safety Security and Rescue Robotics (SSRR), 2010 IEEE International Workshop on*, pages 1 –6, july 2010. 2

[7] F. Matsuno, N. Sato, K. Kon, H. Igarashi, T. Kimura, and R. Murphy. Utilization of robot systems in disaster sites of the great eastern japan earthquake. In *Proceedings of Field & Service Robotics (FSR) 2012*, Matsushima/Sendai, Japan, 2012. 2

[8] R.R. Murphy and J.L. Burke. The safe human-robot ratio. In M.J. Barnes and F. Jentsch, editors, *Human-Robot Interactions in Future Military Operations*, Human Factors in Defence, pages 31–49. Ashgate, 2010. 2, 4, 5

[9] R.R. Murphy, K. Pratt, and J.L. Burke. Crew roles and operational protocols for rotary-wing micro-UAVs in close urban environments. In *Proceedings of the ACM/IEEE Conference on Human-Robot Interaction (HRI'08)*, Amsterdam, The Netherlands, 2008. 2

[10] R.R. Murphy, E. Steimle, C. Griffin, C. Cullins, M. Hall, and K. Pratt. Cooperative use of unmanned sea surface and micro aerial vehicle at Hurricane Wilma. *Journal of Field Robotics*, 25(3):164–180, 2008. 2

[11] D. R. Olsen and M. A. Goodrich. Metrics for evaluating human-robot interactions. In *Proceedings of PERMIS 2003*, 2003. 5

[12] QUEST. Macro-seismic report of 20(ml5.9) and 29th of may (ml5.8 and 5.3) in the emilia valley. Technical report, `http://terremoti.ingv.it/it/ultimi-eventi/842-terremoti-in-pianura-padana-emiliana.html`, 2012. 1

[13] M. Quigley, K. Conley, B.P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A.Y. Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009. 2

[14] P.M. Salmon, N.A. Stanton, G.H. Walker, and D.P. Jenkins. *Distributed Situation Awareness: Theory, Measurement, and Application to Teamwork*. Human Factors in Defence. Ashgate, 2009. 5

[15] T. Yoshida, K. Nagatani, S. Tadokoro, T. Nishimura, and E. Koyanagi. Improvements to the rescue robot quince -toward future indoor surveillance missions in the fukushima daiichi nuclear power plant. In *Proceedings of Field & Service Robotics (FSR) 2012*, Matsushima/Sendai, Japan, 2012. 2, 3