Noname manuscript No.
(will be inserted by the editor)

Safety, Security, and Rescue Missions with an Unmanned Aerial Vehicle (UAV)

Aerial Mosaicking and Autonomous Flight at the 2009 European Land Robots Trials (ELROB) and the 2010 Response Robot Evaluation Exercises (RREE)

Andreas Birk · Burkhard Wiggerich · Heiko Bülow · Max Pfingsthorn · Sören Schwertfeger

Received: date / Accepted: date

Abstract The usage of an Unmanned Aerial Vehicle (UAV) in different realistic safety, security, and rescue exercises is presented. First, results from two safety and security missions at the 2009 European Land Robot Trials (ELROB) are presented. An Unmanned Aerial Vehicle (UAV) in form of an Airrobot AR100-B is used in a reconnaissance and in a camp security scenario. The UAV is capable of autonomous waypoint navigation using onboard GPS processing. A digital video stream from the vehicle is used to create photo maps - also known as mosaicking - in real time at the operator station. This mapping is done using an enhanced version of Fourier Mellin based registration, which turns out to be very fast and robust. Furthermore, results from a rescue oriented scenario at the 2010 Response Robot Evaluation Exercises (RREE) at Disaster City, Texas are presented. The registration for the aerial mosaicking is supplemented by an uncertainty metric and embedded into Simultaneous Localization and Mapping (SLAM), which further enhances the photo maps as main mission deliveries.

Keywords Unmanned Aerial Vehicle (UAV) · Autonomy · Registration · Mosaicking · Visual SLAM

1 Introduction

Aerial vehicles are obviously well suited to give a bird's eye view over an incident site. There offer hence interesting options as robotics tools for assisting in safety, se-

A. Birk, H. Bülow, M. Pfingsthorn, and S. Schwertfeger
Jacobs University Bremen
Campus Ring 1
28759 Bremen, Germany
E-mail: a.birk@jacobs-university.de
http://robotics.jacobs-university.de

B. Wiggerich
AirRobot GmbH & Co. KG
Werler Strasse 4-8
59755 Arnsberg, Germany
E-mail: info@airrobot.de
http://www.airrobot.de

curity, and rescue missions. We present here results from using an UAV in two field tests where different realistic response scenarios were investigated, namely the 2009 European Land Robot Trials (ELROB-2009) and the 2010 Response Robot Evaluation Exercises (RREE-2010).

ELROB-2009 was the fourth event in an annual series of evaluation events for Safety, Security, and Rescue Robotics (SSRR) (ELROB 2009). ELROB-2009 took place from 14.-18. June 2009 in Oulu, Finland. The related results presented in this article are based on a joined participation of the German company Airrobot with the Jacobs Robotics Group. The team used an Airrobot AR100-B Unmanned Aerial Vehicle (UAV) in two SSRR missions, namely reconnaissance and camp security, i.e., two safety and security related missions. The Airrobot AR100-B is a quadrocopter (figure 1), which is capable of onboard processing of GPS signals and of stored way points, hence allowing fully autonomous operations. It is also equipped with a high quality digital video link to the operator station.

A third application scenario was of interest at the 2010 Response Robot Evaluation Exercise (RREE-2010) at Disaster City in College Station, Texas (TEEX 2008). RREE is an annual event organized by the Intelligent Systems Division (ISD) of the National Institute of Standards and Technology (NIST). At this event, the task for the aerial vehicle was to gather overview of a rubble pile, i.e., the mission is more related to rescue than safety and security. A slightly different Airrobot model, namely an AR100, was used at RREE-2010.

Fig. 1 The Airrobot AR100-B taking off for a mission at ELROB 2009.

Data from an onboard video link is used in all missions for photo mapping - also known as mosaicking - by registering the video frames with each other and hence providing a larger overview. These photo maps serve two main purposes. First, they help the operator in the mission: they stabilize the video stream, which is quite shaky, and they provide much larger overview than a video. Second, they are an important mission deliverable. Even though there are many Geo Information Systems (GIS) available including widespread popular ones like Google Maps (Google 2009) and Microsoft Bing Maps (Microsoft 2009), they have severe limitations. First of all, there are still sur-

prisingly many regions that are not or only poorly covered by certain systems. Google Maps has for example only an extremely low resolution coverage of Oulu and especially the competition site (figure 2). Also, GIS may provide inaccurate information (figure 3) and it is even in the best case less precise than what can be achieved with a low flying UAV. But the most important motivation for real time photo mapping is that it provides an actual overview of the mission site at the time of the mission, in contrast to the older data stored in a GIS.

Fig. 2 The ELROB 2009 competition site is only poorly covered in Google maps; the above image shows the best available resolution of the area.

Naively, the generation of a photo map is trivial as long as the exact pose of the vehicle is know. But GPS localization is very imprecise, the 3D orientation of the UAV is hard to determine, and so on. An alternative is to use registration of subsequent frames from a camera on the vehicle. Related work in the areas of visual odometry and visual Simultaneous Localization and Mapping (SLAM) predominantly uses feature based approaches based for example on the scale invariant feature transform (SIFT) (Lowe 2004, Lowe 1999) or the Kanade-Lucas-Tomasi Feature Tracker (KLT) (Shi & Tomasi 1994). This holds for the area of Visual SLAM (VSLAM) in general (Davison, Reid, Molton & Stasse 2007)(Eustice, Singh & Leonard 2006)(Angeli, Filliat, Doncieux & Meyer 2006)(Davison 2003)(Castellanos, Neira & Tardos 2001)(Castellanos, Martinez, Neira & Tardos 1998)(Broida, Chandrashekar & Chellappa 1990)(Grandjean & Robert De Saint Vincent 1989) and in particular for work dealing with UAV and photo mapping (Steder, Grisetti, Stachniss & Burgard 2008)(Angeli et al. 2006).

In our approach, not only local descriptors but the whole information within the images is used, which seems to be very beneficial in the context of photo mapping as described in more detail in (Buelow & Birk 2009). An variant of the Fourier Mellin Invariant (FMI) transform for image representation and processing (Chen, Defrise & Deconinck 1994)(Reddy & Chatterji 1996) is used with two significant modifications in

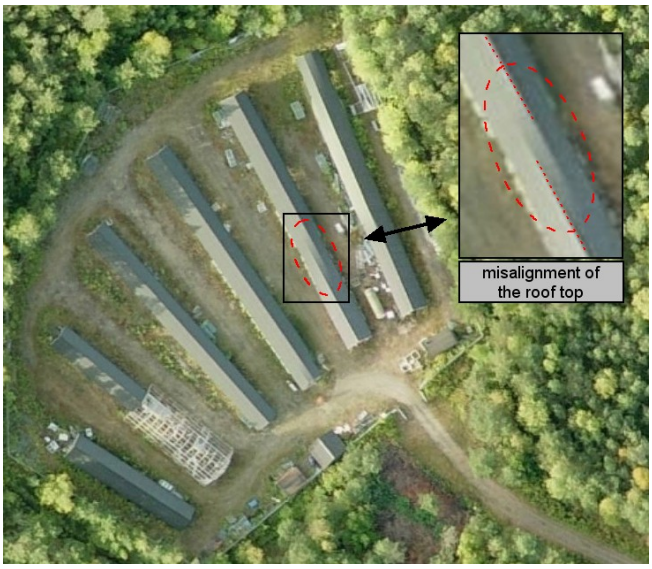

Fig. 3 Even professional map services are far from perfect. The above image shows an obvious bug in the coverage of the ELROB reconnaissance site by Microsoft Bing Maps.

comparison to previous work. First, a logarithmic representation of the spectral magnitude of the FMI descriptor is used. Second, a filter on the frequency where the shift is supposed to appear is applied. An analysis of the benefits of this improved FMI (iFMI) is provided in (Buelow & Birk 2009). The iFMI provides 2D translations, rotations, and scaling, i.e., registrations when the UAV is moving horizontally, rotating its yaw, and changing its altitude. In addition, it requires in theory a strictly down looking camera and a flat world hypothesis is assumed. But as shown in the experiments, these constraints do not have to be strictly fulfilled to achieve meaningful results.

We furthermore developed an uncertainty metric for iFMI registrations (Pfingsthorn, Schwertfeger, Buelow & Birk 2010), which is up to our knowledge the first time that this was done for a spectral registration method. This allows to embedded the iFMI registration in SLAM, which further improves the mosaicking results as also shown in this article.

2 Improved Fourier Mellin Mapping

2.1 Improved Fourier Mellin Image Registration

The basis for the Improved Fourier Mellin (iFMI) (Buelow & Birk 2009)(Buelow, Birk & Unnithan 2009) image registration method used in this work is a Phase-Only Matched Filter (POMF). In this correlation approach, the signal shift is calculated based on the phase difference of the image spectra. If the POMF is applied to two periodically shifted signals, the result is an ideal Dirac pulse indicating the signal shift. This resulting Dirac pulse deteriorates with changing signal content of both signals. As long as the POMF yields a clear detectable maximum this method can be used for matching two signals.

The POMF can be used to detect both translation in the image content as well as rotation and scaling. Translation and rotation/scaling is decoupled because the 2D spectrum rotates and scales in the same way as the image in the time domain. However, the spectrum magnitude is invariant in translation. The spectrum magnitudes are resampled in polar coordinates using a logarithmic radial axis. A translation is this log-polar space signifies a rotation along the angular axis and a scaling along the radial axis. The logarithmic resampling is a way to compute the Mellin transform with a Fourier transform, hence the name (Chen et al. 1994)(Reddy & Chatterji 1996). It is thus possible to use the 2D spectrum magnitudes of the images to register rotation and scaling parameters, and the image content itself to register translation.

Algorithm 1: The Phase Only Matched Filter, later used as procedure $\Delta = \mathrm{POMF}(A, B)$

Data: Matrices A and B
Result: A matrix Δ containing a discretized Dirac pulse
$F_A = \mathrm{FFT}_{2D}(A)$;
$F_B = \mathrm{FFT}_{2D}(B)$;
$D = \mathrm{PHASE}(F_A) - \mathrm{PHASE}(F_B)$;
$\Delta = \mathrm{IFFT}_{2D}(\exp(iD))$;
return Δ;

Algorithm 2: iFMI Image Registration Algorithm

Data: Grayscale Images A, B
Result: Transformation Parameters T: Translation x, y, Rotation θ, Scaling γ
`/* Calculate spectra */`
$S_A = \mathrm{FFT}_{2D}(A)$;
$S_B = \mathrm{FFT}_{2D}(B)$;
`/* resample magnitudes into polar coordinates with logarithmic radius */`
$P_A = \mathrm{LOGPOLAR}(|S_A|)$;
$P_B = \mathrm{LOGPOLAR}(|S_B|)$;
`/* apply POMF */`
$\Delta_{rot/scale} = \mathrm{POMF}(P_A, P_B)$;
`/* find dirac peak within scaling/rotation parameter space */`
$\{\theta^*, \gamma^*\} = \arg\max_{\theta,\gamma} \Delta_{rot/scale}(\theta, \gamma)$;
`/* rotate and scale second image into reference frame */`
$B' = \mathrm{ROTATE\_SCALE}(B, \theta^*, \gamma^*)$;
`/* apply POMF once more */`
$\Delta_{trans} = \mathrm{POMF}(A, B')$;
`/* find dirak peak within translation parameter space */`
$\{x^*, y^*\} = \arg\max_{x,y} \Delta_{trans}(x, y)$;
return $T(x^*, y^*, \theta^*, \gamma^*)$;

The steps are used in the Fourier Mellin based Mapping in a straightforward way. Algorithm 1 shows the calculation of the Phase Only Matched Filter, and algorithm 2 details the iFMI registration method. The image registration is done incrementally, keeping track of the global transformation parameters from a main reference frame to the current image.

A first reference image I_0 is acquired or provided to define the reference frame. For each subsequent image I_n the image registration is performed with the preceding image I_{n-1} resulting in a transformation estimate T_{n-1}^{n}. Two subsequent transformations T_n^m and T_m^k can be combined as follows:

$$
\begin{aligned}
T_n^k &= T_n^m \oplus T_m^k \\
&\quad \text{with specific parameters:} \\
x_n^k &= x_n^m + \gamma_n^m \left(\cos(\theta_n^m) x_m^k - \sin(\theta_n^m) y_m^k \right) \\
y_n^k &= y_n^m + \gamma_n^m \left(\sin(\theta_n^m) x_m^k + \cos(\theta_n^m) y_m^k \right) \\
\theta_n^k &= \theta_n^m + \theta_m^k \\
\gamma_n^k &= \gamma_n^m \gamma_m^k
\end{aligned}
$$

At each iteration, the new transformation estimate T_{n-1}^{n} is combined with the previous global estimate T_0^{n-1} to form the new global estimate T_0^n for the image I_n. This new global transformation is used to project image I_n into the global image map.

2.2 Improvements for more registration stability

In order to make the registration more stable the search area within the FMI descriptor the search area for scaling and rotation is restricted to reasonable parameters between two image frames. As a second measure the FMI descriptor is in addition to a standard spectral window processed by a logarithmic function in order to suppress low frequencies which often lead to erroneous registration results. Furthermore the domain after the inverse Fourier transform is processed by a FIR interpolation filter to suppress interference from aliasing image content and lift the energy from a correct registration peak which can be distributed around several frequency cells. A more detailed discussion of the approach can be found in (Buelow & Birk 2009).

2.3 Uncertainty Estimations for SLAM

Please note that the iFMI described above is a pure registration method, i.e., it determines a spatial transform between images. We have recently extended iFMI in own work (Pfingsthorn et al. 2010) with an uncertainty analysis, which is up to our knowledge done for the first time for a spectral registration method. A covariance matrix is extracted from the result of a Phase-Only Matched Filter, which is interpreted as a probability mass function.

Algorithm 3: Uncertainty Analysis of iFMI Registration Result as compute by algorithm 2

Data: Parameter Estimates x, y, θ, γ, POMF Results $\Delta_{rot/scale}, \Delta_{trans}$, and Interpolation Neighborhood N

Result: Covariance Matrices $C_{rot/scale}$ and C_{trans}, for parameters θ, γ and x, y, respectively

`/* normalize POMF Result histograms */`

$$\Delta_{rot/scale} = \Delta_{rot/scale} \cdot \left(\sum_i \sum_j \Delta_{rot/scale}(i,j) \right)^{-1};$$

$$\Delta_{trans} = \Delta_{trans} \cdot \left(\sum_i \sum_j \Delta_{trans}(i,j) \right)^{-1};$$

`/* fit covariance to histograms around estimates */`

$$C_{rot/scale} = \sum_{i=\theta-N}^{\theta+N} \sum_{j=\gamma-N}^{\gamma+N} \Delta_{rot/scale}(i,j) \cdot (i-\theta, j-\gamma)^T (i-\theta, j-\gamma);$$

$$C_{trans} = \sum_{i=x-N}^{x+N} \sum_{j=y-N}^{y+N} \Delta_{trans}(i,j) \cdot (i-x, j-y)^T (i-x, j-y);$$

return $C_{rot/scale}, C_{trans}$;

The method is embedded in a pose graph implementation for Simultaneous Localization and Mapping (SLAM) (Olson, Leonard & Teller 2006, Grisetti, Stachniss, Grzonka & Burgard 2007). The covariances matrices computed in algorithm 3 are combined into one block diagonal matrix for the parameters x, y, θ. Scaling (γ) is encoded into a height parameter z, as it is directly proportional to it in the SLAM setting. Within the implementation, the covariances are reordered and scaled accordingly.

Each registration result gives rise to an edge in a pose graph, while vertices represent poses at which images were acquired. This representation is especially useful as it allows for easy manipulation at a later time. More integrated representations, such as the image mosaics shown in the next section, are easily reconstructable from the graph. Loops in the pose graph, i.e. substantial overlaps between non-consecutive image frames, are detected automatically by comparing current pose estimates of two nearby vertices. Once all loops are found, the excellent TORO pose graph global optimization library (Grisetti et al. 2007) is used to reduce the overall map error.

We demonstrate the benefits of iFMI-based SLAM in (Pfingsthorn et al. 2010) through results in the underwater domain, i.e., mosaicking of sea beds. The results presented here from ELROB-2009 are based on registration only, which already gives very useful results as can be seen later on in the related sections 3.2 and 3.3. However, the field tests from RREE-2010 show that the usage of SLAM further improves the map quality. The according results are presented in section 3.4.

3 Results

3.1 Safety and Security Missions at ELROB-2009

As discussed in a bit more detail later on, the UAV performed significant amounts of time in autonomous mode (table 1) at ELROB-2009 despite challenging weather

scenario	autonomy (mm:ss)	overall (mm:ss)
reconnaissance	11:00	22:00
camp security	20:00	57:30

Table 1 The Autonomous Flight Times of the UAV. Please note that the overall mission times include the set up and the time for the provision of the mission deliverables like maps, photos of the intruders, etc. to the organizers.

	reconnaissance			camp security	
fig.	6	7	8	11	12
#frames	31	52	87	337	784

Table 2 Number of frames in the example maps

conditions. The main limitation in the photo mapping are error frames in the video (figure 4), which could not be registered. When an error frame occurred, the generation of a new photo map was started. The photo mapping took place on a Laptop with a Core-2 Duo 2.0 GHz CPU running Ubuntu 9.04 Linux.

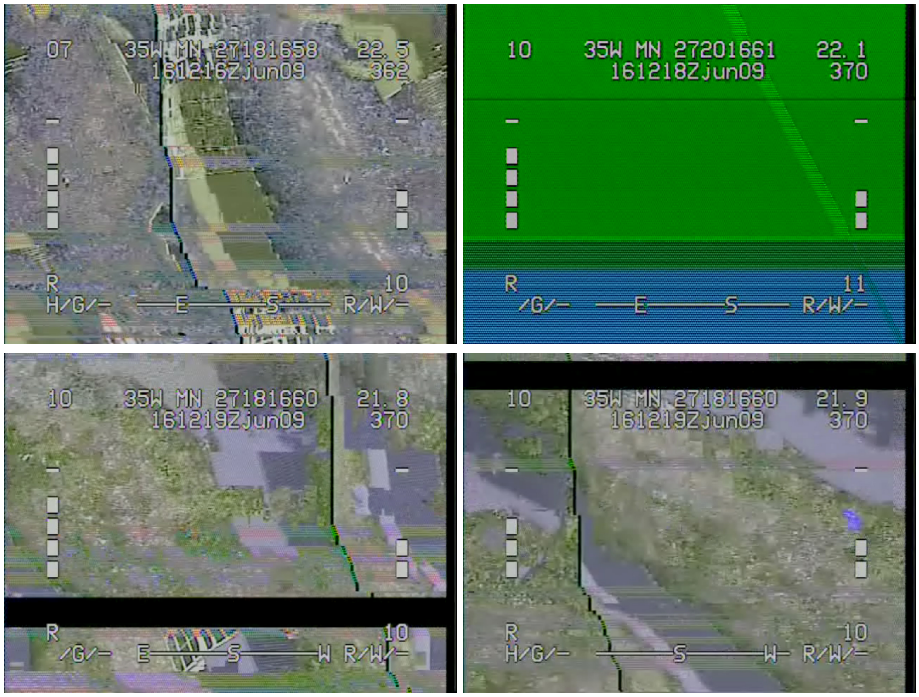

Fig. 4 The digital video link from the UAV had to cover up to 300 m through forest under humid weather conditions. There were occasional video errors despite using a high quality, military grade RF connection. The above images show some typical examples.

3.2 Reconnaissance at ELROB-2009

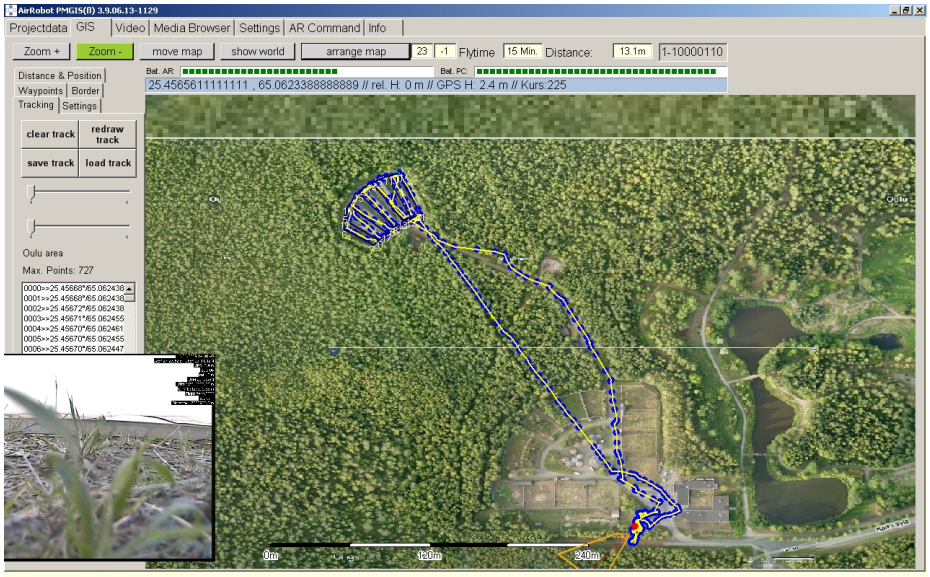

Fig. 5 The waypoints in the reconnaissance path

The mission task for the reconnaissance mission was to detect hazardous spots in an Nuclear-Biological-Chemical (NBC) incident scenario. The hazardous spots were marked by Emergency Response Intervention Cards or short ERI-cards. ERI-cards are rectangular plates of about 40 cm × 30 cm, with black number codes for different hazardous materials on orange background.

The main site for the reconnaissance operations was about 360 m away from the operator's station. The site itself spans about 60 m by 60 m (figure 5). The UAV used GPS way points to autonomously reach the site and to explore it.

The UAV operated fully autonomously for 11 minutes in total out of 22 minutes mission time (table 1). Please note that the mission time includes set up and the provision of mission deliverables to the organizers. There were only about 4 minutes of manual intervention during the actual flight time of about 15 minutes. The manual operations included start and landing as well as the upload of a second set of preplanned GPS way points to pursue a different search strategy. The vehicle was capable of autonomous operation under difficult wind conditions; wind gusts in the air exceeded 12 m/sec. The UAV did not detect any of the ERI-cards. But also 5 out of the 6 other teams - all of them using land robots - did not find any ERI-cards either. Only 1 team out of the 7 teams managed to find 2 ERI-cards.

The mission scenario involved significant communication challenges for the video transmission. The maximum distance was about 400 m with significant amounts of vegetation, namely forest in humid weather conditions, between the vehicle and the operator's station. There were hence quite often error frames in the video data stream. As mentioned, we are currently working on an automatic detection of error frames to

Fig. 6 An example of a photo map in the reconnaissance scenario (31 frames).

prevent a disruption of the map generation. Typical error-free video sequences consisted of about 75 frames. Figures 6 and 7 show two examples with 31, respectively 52 frames. It has to be noted that the approach also gives reasonable results when the camera is tilted though it should fail in theory under such conditions. Figure 8 shows an according example based on 87 frames. The maps were generated in real time at the operator's station with about 50 msec processing time for registration and visualization, i.e., 20 Hertz update rate.

Fig. 7 A second example of a photo map in the reconnaissance scenario (52 frames).

3.3 Camp Security at ELROB-2009

The mission task for the camp security was to detect intruders. The operator's station was located about 30 m from the camp site, which spans about 90 m by 90 m (figure 9). The UAV mainly autonomously hovered over the site and intruders were "chased" in manual mode by the operator when being detected. The overall mission time was 60 minutes including set up and the provision of mission deliverables in the end.

Fig. 8 An example of a photo map of the way of the UAV to the reconnaissance site (87 frames). Please note that there is a significant tilt of the camera and hence skew in the images. Though the algorithm is in theory not capable of registering under these conditions, the result is nevertheless quite usable of a rough orientation.

In total, 6 intruders out of 8 possible ones were detected and "caught" by the UAV. Two examples of detected intruders are shown in figure 10. Only a 2nd team out of the 7 teams in total - all other 6 teams using land robots - was also able to catch 6 intruders. One other team caught 2, three teams caught 1, and one team caught 0 intruders. The 2 intruders that were missed entered the camp while the UAV was for a 5 minute stop at the operator's station for swapping the batteries - as mentioned before, the maximum operation time of the UAV is about 25 to 30 minutes, i.e., batteries had to be changed once during this mission.

As shown in table 1, the UAV operated fully autonomously for 20 minutes in total out of 57.5 minutes flight time, i.e., for about 34.8% of the time. The wind conditions were also quite challenging during this mission, the wind speed on ground was measured with up to 7.2 m/sec, wind gusts in the air exceeded 12 m/sec.

The mission scenario involved shorter communication ranges, less than 150 m maximum, and fewer vegetation, especially few trees in the line of sight. But the "camp" - being a part of the former Oulu Zoo - contains a significant amount of metal struc-

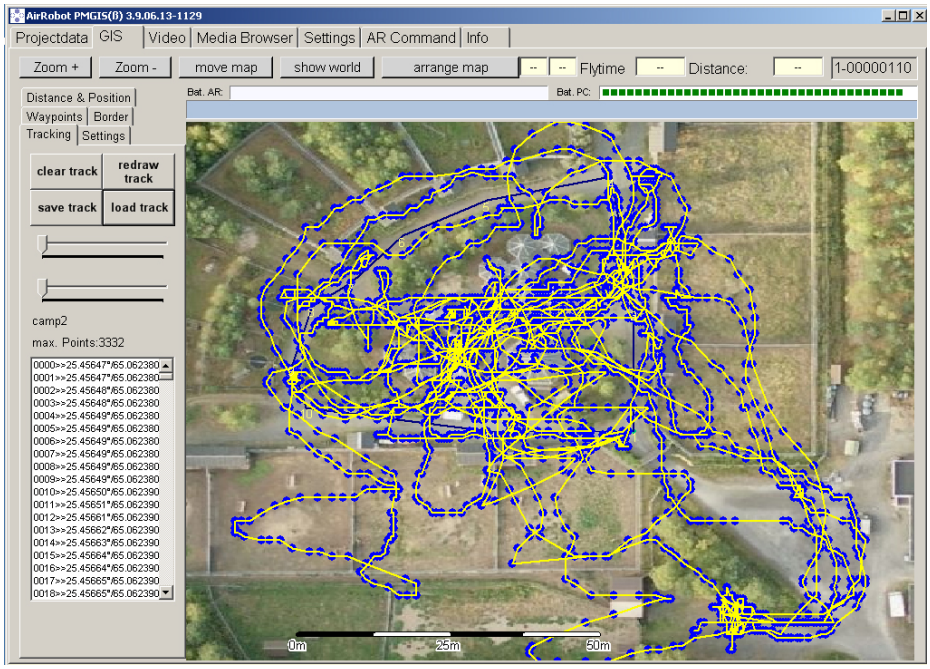

Fig. 9 The actual path of the UAV in the camp security scenario based on GPS track points.

tures, especially fences and cages with wire meshes. Nevertheless, the video link was much less error prone than for the reconnaissance mission. The video sequences between error frames were hence significantly longer. On average, video sequences with 650 frames were turned into photo maps. Figures 11 and 12 show two examples with 337, respectively 782 frames. The maps were generated in real time at the operator's station with about 50 msec processing time for registration and visualization, i.e., 20 Hertz update rate.

3.4 Mapping a Rubble Pile at RREE-2010

The ELROB missions required - quite typically for safety and security scenarios - a significant distance between the operator and the vehicle and hence quite some amounts of autonomous operation. The RREE mission presented in this section is in contrast quite typical for a rescue mission. The operator of the vehicle can get relatively close to the incident site and manual flight in eyes-on-mode is possible.

The RREE scenario presented here is a very good example for a possible usage of UAV in rescue missions, namely the generation of an overview map of a rubble pile. Rubble piles are potentially very dangerous to physically access, both for responders as well as for victims trapped inside of the pile. All operations have therefore to be carefully planned and a bird's eye overview generated from aerial video is a very good basis for this.

The most interesting aspect of the rubble pile map shown in figure 13 is in the context of this article that it has been generated - in contrast to the ELROB-2009

Fig. 10 Two examples of intruders detected during the camp security mission.

Fig. 11 An example of a photo map of the camp generated from 337 video frames in real time.

Fig. 12 A second example of a photo map of the camp generated from 784 video frames in real time.

mosaics - with full SLAM and not only with sequential registrations. Concretely, 630 video frames where taken. Figure 14 shows on the left the path of the robot and the map generated by sequential registrations. On the right in the same figure, the pose-graph is shown that consists of both the edges of the path and of edges that have been generated by loop closures, i.e., the detection of non-sequential image pairs that can be registered with each other.

The loop closures lead to a significant increase in the number of edges, namely from the 629 sequential ones that represent the path of the robot to a graph with 12,364 edges. All edges are supplemented by uncertainty information, i.e., a covariance matrix that is extracted from the result of the phase-only matched filters in the iFMI registrations. The pose-graph is relaxed using the open source TORO optimizer (Grisetti et al. 2007). As shown in figure 15, the SLAM improves the map quality.

The computation of the uncertainty information hardly influences the runtime of the registration at all, i.e., it can still be done in real-time with 20-30 Hz. The optimization of the pose-graph is quite fast, too. It takes 2.12 seconds in the example

Fig. 13 A map of a rubble pile at Disaster City generated at RREE-2010 from 630 video frames via iFMI-SLAM.

with the 12,364 edges. Please note that previously computed transformations on the images can be reused, hence speeding up the registrations. So while the UAV is flying, the iFMI already generates a reasonable photo map. Once the mission is finished or a higher quality overview is required, a loop closing and optimization process can be triggered and an improved map is generated after a few seconds via SLAM.

4 Conclusions

The usage of an Unmanned Aerial Vehicle (UAV) in different realistic safety, security, and rescue exercises is presented. First, results from two safety and security missions at the 2009 European Land Robot Trials (ELROB) are presented. An Unmanned Aerial Vehicle (UAV) in form of an Airrobot AR100-B is used in a reconnaissance and in a camp security scenario. The UAV is capable of autonomous waypoint navigation using onboard GPS processing. A digital video stream from the vehicle is used to create photo maps - also known as mosaicking - in real time at the operator station. This mapping is

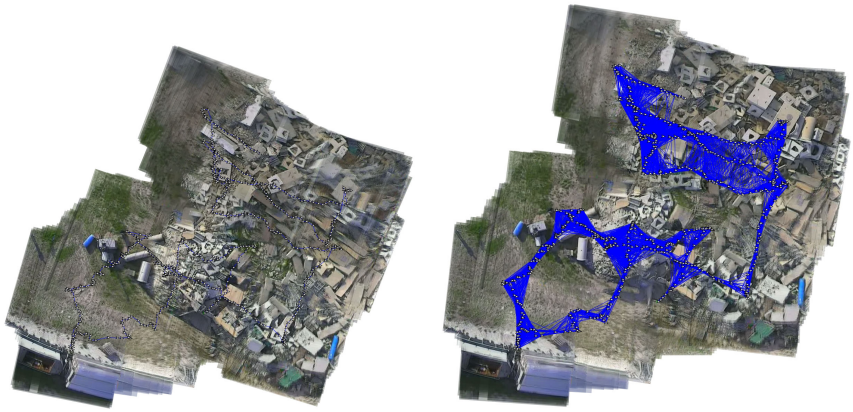

(a) sequential only (b) including loop closures

Fig. 14 The path of the UAV over the rubble pile as pose-graph, once without (left) and once with loop closures edges for SLAM (right). The little triangles show the 630 places where a photo was taken.

(a) sequential only (b) SLAM

Fig. 15 A detail of the rubble pile map, once generated with sequential registration only (left) and once embedded in posegraph SLAM (right).

done using an enhanced version of Fourier Mellin based registration, which turns out to be very fast and robust. Furthermore, results from a rescue oriented scenario at the 2010 Response Robot Evaluation Exercises (RREE) at Disaster City, Texas are presented. The registration for the aerial mosaicking is supplemented by an uncertainty metric and embedded into Simultaneous Localization and Mapping (SLAM), which further enhances the photo maps as main mission deliveries.

Results dealing with the usage of an Unmanned Aerial Vehicle (UAV) in different safety, security, and rescue exercises were presented. First, two SSRR missions with an Unmanned Aerial Vehicle (UAV) at the 2009 European Land Robot Trials (ELROB) were presented. An Airrobot AR100-B was used in a reconnaissance and a camp security scenario. The UAV is capable of GPS based way point navigation, which was used for significant amounts of autonomous operations despite very challenging weather conditions with gusts of wind with top speeds of more than 12 m/sec. A digital video link was used for photo mapping in real time. An improved version of Fourier Mellin based registration with two core modifications is used for this purpose. First, a logarithmic representation of the spectral magnitude of the FMI descriptor is used. Second, a filter on the frequency where the shift is supposed to appear is applied. The resulting photo maps provide high resolution, up to date information about the mission site.

Second, results from a rescue scenario at the 2010 Response Robot Evaluation Exercises (RREE) at Disaster City, Texas were presented. Concretely, a manually flown UAV was used to generate an overview map of a rubble pile. The iFMI registration for the aerial mosaicking is supplemented in this mission by an uncertainty metric. this allows to embed the registration into a pose-graph for Simultaneous Localization and Mapping (SLAM). The SLAM further enhances the photo map.

References

Angeli, A., Filliat, D., Doncieux, S. & Meyer, J. A. (2006). 2d simultaneous localization and mapping for micro aerial vehicles, *Proceedings of the European Micro Aerial Vehicles (EMAV 2006) conference*.

Broida, T. J., Chandrashekar, S. & Chellappa, R. (1990). Recursive estimation of 3d motion from a monocular image sequence, *IEEE Transactions on Aerospace and Electronics Systems* **26**: 639–656.

Buelow, H. & Birk, A. (2009). Fast and Robust Photomapping with an Unmanned Aerial Vehicle (UAV), *International Conference on Intelligent Robots and Systems (IROS)*, IEEE Press.

Buelow, H., Birk, A. & Unnithan, V. (2009). Online Generation of an Underwater Photo Map with Improved Fourier Mellin based Registration, *IEEE OCEANS*, IEEE Press.

Castellanos, J., Martinez, J., Neira, J. & Tardos, J. (1998). Simultaneous map building and localization for mobile robots: a multisensor fusion approach, *in* J. Martinez (ed.), *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, Vol. 2, pp. 1244–1249 vol.2.

Castellanos, J., Neira, J. & Tardos, J. (2001). Multisensor fusion for simultaneous localization and map building, *Robotics and Automation, IEEE Transactions on* **17**(6): 908–914.

Chen, Q.-S., Defrise, M. & Deconinck, F. (1994). Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition, *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **16**(12): 1156–1168.

Davison, A. J. (2003). Real-time simultaneous localisation and mapping with a single camera, *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pp. 1403–1410.

Davison, A. J., Reid, I. D., Molton, N. D. & Stasse, O. (2007). Monoslam: Real-time single camera slam, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**: 1052–1067.

ELROB (2009). European Land-Robot Trial (ELROB), `www.elrob2009.org`.

Eustice, R. M., Singh, H. & Leonard, J. J. (2006). Exactly Sparse Delayed-State Filters for View-Based SLAM, *Robotics, IEEE Transactions on [see also Robotics and Automation, IEEE Transactions on]* **22**(6): 1100–1114.

Google (2009). Google Maps, `http://maps.google.com`.

Grandjean, P. & Robert De Saint Vincent, A. (1989). 3-D modeling of indoor scenes by fusion of noisy range and stereo data, *in* A. Robert De Saint Vincent (ed.), *Robotics and Automation, 1989. Proceedings., 1989 IEEE International Conference on*, pp. 681–687 vol.2.

Grisetti, G., Stachniss, C., Grzonka, S. & Burgard, W. (2007). A Tree Parameterization for Efficiently Computing Maximum Likelihood Maps using Gradient Descent, *Robotics: Science and Systems (RSS)*, Atlanta, GA, USA.

Lowe, D. G. (1999). Object Recognition from Local Scale-Invariant Features, *Proceedings of International Conference on Computer Vision*, pp. 1150–1157.

Lowe, D. G. (2004). Distinctive Image Features from Scale-Invariant Keypoints, *International Journal of Computer Vision* **60**(2): 91–110.

Microsoft (2009). Bing Maps, `http://www.bing.com/maps/`.

Olson, E., Leonard, J. & Teller, S. (2006). Fast iterative alignment of pose graphs with poor initial estimates, *in* J. Leonard (ed.), *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pp. 2262–2269.

Pfingsthorn, M., Schwertfeger, S., Buelow, H. & Birk, A. (2010). Maximum Likelihood Mapping with Spectral Image Registration, *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE Press.

Reddy, B. & Chatterji, B. (1996). An FFT-based technique for translation, rotation, and scale-invariant image registration, *Image Processing, IEEE Transactions on* **5**(8): 1266–1271.

Shi, J. & Tomasi, C. (1994). Good features to track, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR94)*.

Steder, B., Grisetti, G., Stachniss, C. & Burgard, W. (2008). Visual SLAM for Flying Vehicles, *Robotics, IEEE Transactions on* **24**(5): 1088–1093.

TEEX (2008). Disaster City, `http://www.teex.com/teex.cfm?pageid=USARprog\&area=USAR\&templateid=1117`.