This is a repository copy of *Autonomous and scalable control for remote inspection with multiple aerial vehicles*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/109395/

Version: Accepted Version

Autonomous and Scalable Control for Remote Inspection with Multiple Aerial Vehicles

Ruaridh A. Clark, Giuliano Punzo, Charles N. MacLeod, Gordon Dobie,
Rahul Summan, Gary Bolton, Stephen G. Pierce and Malcolm Macdonald

University of Strathclyde, Glasgow, UK

Abstract

A novel approach to the autonomous generation of trajectories for multiple aerial vehicles is presented, whereby an artificial kinematic field provides autonomous control in a distributed and highly scalable manner. The kinematic field is generated relative to a central target and is modified when a vehicle is in close proximity of another to avoid collisions. This control scheme is then applied to the mock visual inspection of a nuclear intermediate level waste storage drum. The inspection is completed using two commercially available quadcopters, in a laboratory environment, with the acquired visual inspection data processed and photogrammetrically meshed to generate a three-dimensional surface-meshed model of the drum. This paper contributes to the field of multi-agent coverage path planning for structural inspection and provides experimental validation of the control and inspection results.

Keywords: Automatic Optical Inspection, Photogrammetry, Swarm, Unmanned Aerial Vehicles.

1. Introduction

This paper considers the challenge of using autonomous unmanned aerial vehicles (UAVs) to visually inspect structural assets such as tanks, flair stacks, chimnes, and wind turbines. Results are presented for a mock inspection of a nuclear intermediate level waste (ILW) storage drum chosen for its convenient size, which enabled laboratory based research, and relevance to industrial partners (National Nuclear Laboratory). Visual structural inspection, using UAVs, has seen interest in the literature through its application to the inside and outside of buildings, for example [1, 2, 3], industrial facilities

[4, 5, 6], landscapes/agriculture and archaeological sites [7, 8, 9]. The current state-of-the-art for UAV motion control has not progressed far beyond the use of waypoints in conjunction with position tracking, such as Global Positioning System (GPS), to autonomously travel to desired points in a set sequence. Such approaches require the calculation of many possible trajectories, followed by collision risk assessment and optimisation routines to reduce the options before a final selection, as in [10, 11]. The authors instead propose a computationally light and highly scalable, collision-free, control scheme for multiple UAVs that could be implemented in a distributed manner.

This scheme is applied to automated structural inspection, where coverage is achieved by minimising the un-scanned areas. Techniques that aim to achieve complete area coverage over a desired region are performing Coverage Path Planning (CPP), which refers to any procedure used for determining that a path passes over all points of an area or volume while avoiding obstacles [12]. CPP algorithms can be classified as heuristic or complete depending on if they probably or definitely guarantee complete coverage of the free space [12]. CPP for aerial vehicles has primarily received research attention in areas such as surveillance [13], agriculture [14] and disaster/emergency management [15]. When considering quantifiable structural inspection, the authors are aware of little research on coverage path planning, especially when considering aerial vehicles. This paper presents a simple but novel approach to automated CPP for multiple aerial vehicles for remote inspection. This approach enables inspection with minimal manual intervention but also features inherent scaling flexibility, which could enable a swarm of vehicles to work in parallel to reduce the time required for inspection by following, vehicle and target, collision free paths.

The approach, detailed in this paper, relies on an analytically defined vector field (often referred to as a Lyapunov vector field) that is used to create a guidance law. A popular way of defining such a field is through the gradient of a scalar function, called an artificial potential function (APF). This was first used in a control architecture by Khatib [16], before Spears et al. applied this concept to the pattern formation problem in multi-agent systems [17, 18]. This control technique has the advantage of being analytically verifiable and has been successfully demonstrated for real world applications [19, 20].

To enable the desired inspection trajectories the vector field has to be defined without considering an associated potential function, for this reason it shall be referred to here as an Artificial Kinematic Field (AKF). AKFs have been the subject of UAV research for both single vehicle control and

coordinated formations [21, 22, 23]. AKF convergence to a desired behaviour, just like with APF, can be proved by means of Lyapunov stability theory. In particular, the problem of circling a target has been considered with fixed wing UAV in simulations and in an open field by Lawrence [21] and Frew [22], where the circling of a moving target was also undertaken. The task in this paper is to adapt these concepts to multiple aerial vehicles, where the main challenge is in preventing vehicle collisions.

A popular way to perform collision avoidance in multi-agent systems is through mutual repulsive potential, see for example [19, 23, 24, 25]. This way each agent alters the kinematic field by producing a short-range repulsive action on the other agents. This is an efficient but crude mechanism for performing collision avoidance as the trajectories generated only consider collision prevention and, hence, can be detrimental to the overall objective. Therefore, a less disruptive approach will be presented that blends well with the global kinematic field.

1.1. Three-Dimensional Model Construction

The inspection can be completed using the aforementioned control methods but the collected data needs to be processed before reviewing the structural state of the target. An intuitive method for displaying visual inspection data is to map it into a three-dimensional (3D) space to allow a reviewer to explore an object to check for faults and failures. Photogrammetry analysis can enable the creation of a 3D model from the inspection images, providing an estimation of the geometry as well as a reviewable environment.

State-of-the-art image based reconstruction systems derive markers or features from the texture information present in the images and by triangulating markers, creating point clouds corresponding to the geometry of the scene. This method is attractive when dealing with remote environments and has been previously investigated in the literature [26, 27, 28, 29]. In particular for robotic inspection [30] but also for archaeological work [31, 32], where the reconstruction tools were the same as those adopted in this work. The system detailed in this paper uses this reconstruction method to enable a remote inspection that then provides data review in a similar manner to an in-person inspection.

2. Guidance and Control

To provide autonomous control for multiple vehicles, a kinematic field is produced that is a function of the vehicle's position with respect to a central target. It provides a highly nonlinear guidance law which is mapped to the control action through a linear controller. This arrangement generates smooth trajectories for the vehicles, which is desirable for many tasks including autonomous visual inspection.

The control architecture is illustrated in Fig. 1. The desktop computer is provided with the vehicle's own position and the relative positions of the other vehicles in the field by using a Vicon MX motion capture system [33], which is discussed in more detail in Section 3. Based on this information, the local kinematic field of each vehicle is computed to produce the desired velocities in the horizontal plane of the external reference frame. These are passed to a linear controller that provides the pitch and roll angles to the on-board controller with desired yaw and vertical speeds also supplied. This in turn commands the motors to execute the requested manoeuvre. The control scheme is designed to be highly scalable if the vehicles have sufficient on-board processing capabilities and position tracking. For the vehicles used in the following work (Parrot AR.Drone 2.0) the on-board processing is limited and hence the control is handled in a distributed manner but by a central processor. The specifics of this scheme are detailed in the following paragraphs.
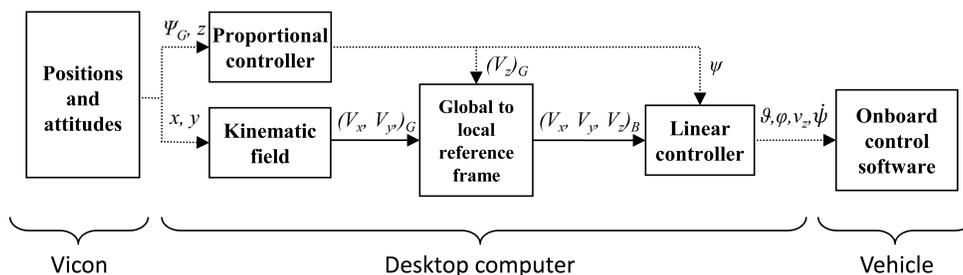

Figure 1: Control architecture scheme for one vehicle (wireless data transfer indicated by dotted line).

4

2.1. Kinematic Field Definition

The kinematic field is defined in the horizontal plane, as the target is assumed to have central symmetry characteristics, with the origin of the global reference frame at the centre of the target. The fundamental structure of the field is a modified version of the Hopf bifurcation function used in [23]. This function is described for an X-Y plane, where x and y are distances from the centre of the field, as

$$v_x^d = c_1(y + \mu x) - Rx\sqrt{x^2 + y^2} \tag{1}$$

$$v_y^d = c_1(-x + \mu y) - Ry\sqrt{x^2 + y^2} \tag{2}$$

where, R defines the radius of an ideal circular trajectory in the horizontal plane enclosing a central target, c_1 is a constant, v_x^d and v_y^d are desired velocities in the x and y directions respectively and μ is a dimensionless scalar parameter that, taken positive, guarantees the emergence of a limit cycle in the field.

In particular, the choice of

$$\mu = \frac{R^2}{c_1} \tag{3}$$

guarantees a circular trajectory of radius R around the centre. This can be verified by transforming eqs. (1) and (2) into polar coordinates and checking that the radial velocity is always null at a distance R from the centre. It can also be easily verified that, along the circular trajectory, the tangential velocity is constant. The kinematic field is completed by a function that provides a more robust control system by strengthening the control action close to the target to avoid collisions, while effectively leaving the characteristics of the field produced with eqs. (1) and (2) unaltered. This function is a radial field in the form $1/(1 + \sqrt{x^2 + y^2})$, which increases the repulsion from the centre while decreasing the attraction at large distances, thus making approaching manoeuvres smoother and preventing overshoots in the direction of the target. The resulting field is described by

$$v_x^d = \frac{c_2}{c_U + \sqrt{x^2 + y^2}}[c_1(y + \mu x) - Rx\sqrt{x^2 + y^2}] \tag{4}$$

$$v_y^d = \frac{c_2}{c_U + \sqrt{x^2 + y^2}}[c_1(-x + \mu y) - Ry\sqrt{x^2 + y^2}] \tag{5}$$

where c_2 is a constant used to scale the whole expression as appropriate to fit its output within the control architecture and c_U is a unitary constant with dimensions [m]. The values for the constants used in this work were found empirically and are tailored to the Parrot AR.Drone 2.0, but could remain the same for implementations of this system with more vehicles. They are $c_1 = 3 \times 10^5$ m^2, $c_2 = 1.5 \times 10^{-7}$ m^{-1}s^{-1} and $R = 1200$ mm. In Fig. 2 a view from above this field is represented with arrows and seven streamlines with different starting points joining in the limit cycle. The constant c_2 performs a scaling function of the vectors in the field, but if negative then the vehicles will be repelled from the desired radius. In the next section, the influence of c_1 shall be explored by showing its function in enabling collision avoidance.

2.2. Vehicle Collision Avoidance

The novel approach to vehicle collision avoidance, presented here, modifies a vehicle's kinematic field when it approaches another vehicle, reducing the magnitude of the field's rotating component. In order to be effective only the trailing vehicle is inhibited. Identification of this vehicle is achieved by considering the scalar product of the relative position vector with the desired velocity vector. In reference to Fig. 3 a binary variable, h, is defined on the basis of the scalar product with

$$V_{des1} \cdot P_{2-1} \geq 0 \rightarrow h = 1 \tag{6}$$

$$V_{des1} \cdot P_{2-1} < 0 \rightarrow h = 0 \tag{7}$$

where P_{2-1} is the position vector of vehicle 2 with respect to vehicle 1 in the global reference frame and V_{des1} is the desired velocity vector, as defined by the kinematic field, for vehicle 1. This enables the kinematic field to be modified asymmetrically, i.e. only the trailing vehicle, where $h = 1$, is affected.

The desired velocity of vehicle 1, as calculated in eqs. (4) and (5), is filtered to create the asymmetrically modified kinematic field. This is achieved by replacing the constant c_1 with the following function

$$c_1^* = H(|P_{2-1}|)c_1 \tag{8}$$

where $H(|P_{2-1}|)$ is dimensionless and scales the rotational component of the field as a function of P_{2-1}; v_x^{d*} and v_y^{d*} are the desired x and y velocity vectors for vehicle 1 within the modified field. This change does not affect the radial
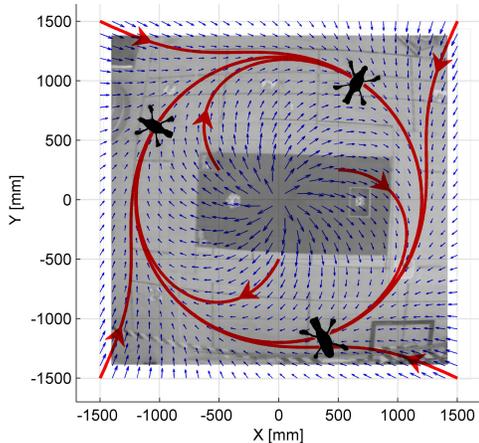
Figure 2: Image of three vehicle lab test overlayed with kinematic field where arrow plots and streamlines highlight how a vehicle would be affected ($R = 1200mm$).
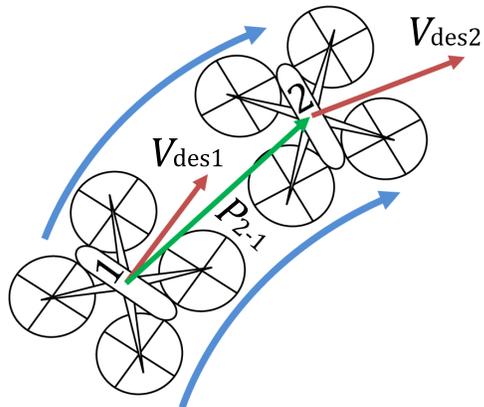

Figure 3: Scheme for the avoidance manoeuvre based on the direction of travel.

velocity at distance R from the target centre, which remains null, with the calculation of μ in eq. (3) updated to

$$\mu = \frac{R^2}{c_1^*}. \tag{9}$$

This asymmetrically modified field only occurs when vehicles are within close proximity. The $H(|P_{2-1}|)$ term that governs this proximity enables a threshold distance between two vehicles to be defined, whereby passing this point results in a switch of direction for the rotational component of the global kinematic field affecting that drone, as depicted in Fig. 4. The modified field, therefore, enables station keeping, relative to the leading vehicle, at the defined distance from the target until the leading vehicle moves on. The function used is in the form

$$H(|P_{2-1}|) = \frac{|P_{2-1}| - \rho}{||P_{2-1}| - \rho|} + \frac{\rho - |P_{2-1}|}{|\rho - |P_{2-1}||} e^{-\frac{(|P_{2-1}| - \rho)^2}{c_s}} \tag{10}$$

where ρ defines the threshold distance between vehicles and c_s is an empirically sought scaling factor that influences the gradient of the function ($c_s = 15 \times 10^4$ m^2). The value of c_s may require adjustment, for example to ensure collision avoidance while operating in a more turbulent environment where the vehicles may deviate from the kinematic field defined path.

7

Fig. 4 details how the scaling of the kinematic field's rotational component affects the modified global field for one vehicle. This figure highlights that H equals 0 at the threshold distance ($\rho = 1000$ mm), chosen as turbulence dominates the vehicle's motion when the inter-vehicle distance is less than a metre. The rotational component of the kinematic field is shown in Fig. 4 to act in opposite directions either side of this threshold distance. Each vehicle only considering the closest vehicle ahead of it when modifying its kinematic field. In the case that one vehicle holds its position, all of the following drones will form a queue along the circular trajectory as seen in Fig. 5. Any final implementation of this system would require an override to allow the system to carry on if the leading vehicle remains stationary due to malfunction.

2.3. Altitude Control

A proportional controller is implemented to control the altitude of the vehicles, which operates in conjunction with the quadcopter's on-board, ultrasound dependant, altitude controller. The output of the proportional controller is converted from the global to the body reference frame as shown in Fig. 1. The requirements for altitude control were simple with the vehicle decreasing its height by a set interval after target coverage was complete at the current altitude, allowing a simple Z-axis waypoint to be used in conjunc-

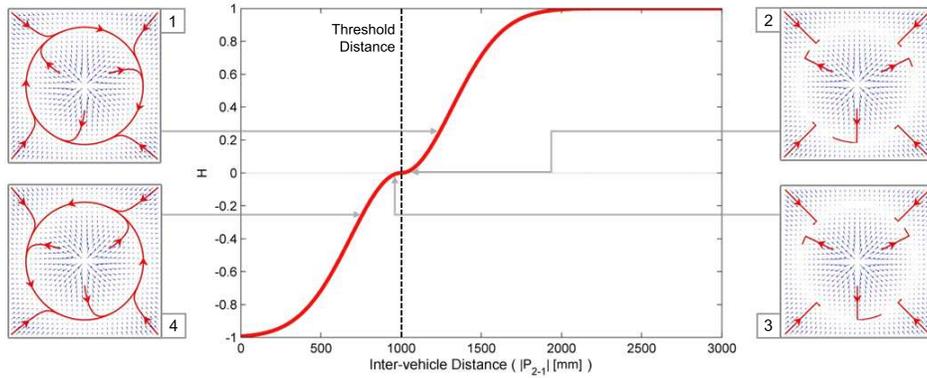Figure 4: Centre: $H(|P_{2-1}|)$ according to (10) where $\rho = 1000$ and $c_s = 15 \times 10^4$ m^2. Side: Instantaneous snapshots from above the modified kinematic field (centred on the target with streamlines displaying flow in the field) at different values of H; rotating clockwise (1), reduced rotational field strength as $H \to 0$ (2), anti-clockwise with reduced (3) and increased strength (4).

tion with this controller to transition to different altitudes. To determine if complete coverage has been achieved at any given altitude each vehicle has to have knowledge of the vehicle ahead of it. In particular its starting position at that coverage altitude, i.e. the point it has to reach before transitioning to another altitude. In this way the control scheme remains easily scalable with collision avoidance and coverage determination requiring each vehicle to only observe the one ahead of it.

2.4. Attitude Control

For a typical quadcopter, pitch and roll angles are coupled with the forward and lateral motion respectively. This design enables forward or side force components to be produced by tilting the vehicle. When no forward or side movements are commanded, the vehicle hovers and in this phase the attitude is controlled in closed loop by the on-board controller only. This is overridden by control commands when altering the yaw angle, which is controlled in the same closed loop manner as the altitude. For the inspection task, discussed in the following case study, the attitude controller keeps the vehicle's x-axis pointing in the direction of the target whilst the quadcopter manoeuvres around it. As a consequence, the desired azimuth changes with position. This is defined as

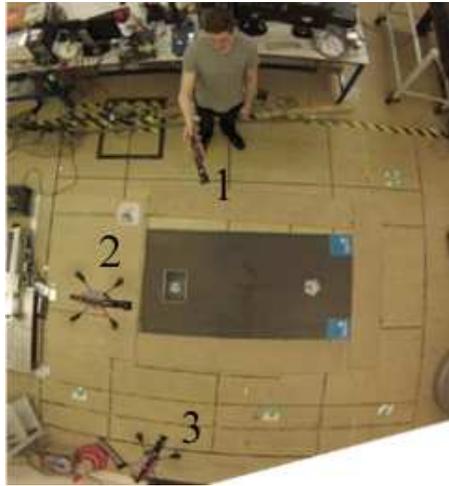$$\psi_{des(i)} = atan2(y_i, x_i) \pm \pi \tag{11}$$

Figure 5: Testing collision avoidance with vehicle 3 queuing behind 2 and 2 queuing behind the shell of vehicle 1, all of which are being tracked by the Vicon system.

where x_i and y_i are the coordinates of the vehicle in the global reference frame that is centred on the target and $atan2(y_i, x_i)$ is similar to calculating the arc tangent of y_i/x_i, except that the signs of both arguments are used to determine the quadrant of the result in the range $[-\pi, \pi]$. To set the desired vehicle angle towards the centre of the field $\pm\pi$ is applied to ensure the result remains within the range $[-\pi, \pi]$. The error in the actual angle is then mapped to an angular rate through a linear controller that selects the shortest rotation direction to reach the desired angle.

2.5. Linear Control

The linear controller maps the desired velocity of each vehicle to commanded pitch and roll angles. The desired velocity vector is decomposed along its forward and lateral components in the body reference frame and these are scaled by a proportional controller. The result is then filtered to output in the range $[-1, 1]$, required for the AR.Drone on-board software, by using the hyperbolic tangent function

$$\gamma^* = tanh(\gamma) \tag{12}$$

where γ is the vector of the controlled variables (including the roll angle φ, pitch angle ϑ, vertical velocity v_z and yaw rate ψ) and γ^* is the normalised output.

A proportional controller is used to map from desired forward and lateral velocities according to the kinematic field, vertical velocity and azimuth angle to commanded pitch and roll angles, vertical velocity and yaw rate. The controller is expressed by

$$\begin{pmatrix} \vartheta \\ \varphi \\ v_z \\ \psi \end{pmatrix} = \begin{pmatrix} k_\vartheta v_{des(x)} \\ k_\varphi v_{des(y)} \\ k_z(Z_{des} - z) \\ k_\psi(\psi_{des} - \psi) \end{pmatrix} \tag{13}$$

where, $v_{des(x)}$ and $v_{des(y)}$ are the forward and lateral velocities in the body reference frame produced by the kinematic field, ψ_{des} is the desired azimuth angle that varies with time, ψ is the actual one, and k_ϑ, k_φ, k_z, k_ψ are the gains of the proportional controller. The values of the gains were sought empirically and specifically to achieve smooth dynamics with the Parrot AR.Drones with $k_\vartheta = 0.7$ sm^{-1}, $k_\varphi = 0.7$ sm^{-1}, $k_z = 1$ s^{-1} and $k_\psi = 1.5$ used. The difference in the gains defined are a result of the quadcopter design with similar control movements resulting in the same gains for k_ϑ and k_φ.

3. Remote Inspection Case Study

A number of experiments were performed to validate the use of AKFs for 3D visual inspection using aerial vehicles:

1. Characterisation of the selected aerial vehicle's flight performance.
2. Two-vehicle inspection of ILW storage drum; characterising flight performance and capturing images for a 3D reconstruction of the drum.
3. A manual visual inspection of the drum, using the Parrot AR.Drone's camera, to ascertain the influence of flight motion on the 3D imaging performance.
4. A manual visual inspection of the drum, using a 24.2MP digital single-lens reflex (D-SLR) camera, to quantify the effect of the vehicle's camera quality on system performance.

3.1. Hardware

The Parrot AR.Drone 2.0 [34] was used in the experiment, a diagram of which is shown in Fig. 6. The AR.Drone 2.0 is a four rotor helicopter, or quadcopter, with a wingspan of 517 mm. It contains four brushless 28,500 RPM motors with long blades of radius 98.5 mm.

The AR.Drone 2.0 includes an embedded 720p, 30 fps, camera with a 92° wide angled lens, which produces JPEG compressed images. It uses a rolling shutter which will lead to some motion blur. A D-SLR camera, equipped with a 24.2MP CMOS sensor [35] and 18–55 mm f/3.5–5.6G VR [36] lens, was then used as the benchmark for experiment 4. It also employed a fine JPEG, compression ratio 4:1, and was used in conjunction with a tripod to obviate motion blur.

A six camera Vicon MX motion capture system [33] was used, as depicted in Fig. 7, for positional tracking, providing coverage for a volume of approximately 6×3 m. The Vicon system tracks 12 mm spherical reflectors, seven of which can be seen on the top of the quadcopter in Fig. 6. Patterns of 3 or more reflectors placed in a unique pattern for identification can be defined as objects and tracked in 6 degrees of freedom at 100 Hz with an estimated error of less than ±3 mm throughout the coverage volume [30].

The inspection sample was an intermediate level waste packaging and encapsulation plant liquor drum, from the Sellafield nuclear reprocessing site, that has a diameter of 800 mm and a height of 1200 mm. It was constructed from 316 Stainless Steel.

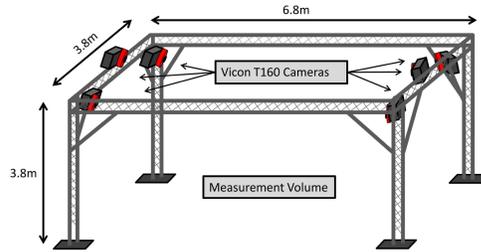Figure 6: Parrot AR.Drone 2.0 with 7 reflective markers.

Figure 7: Vicon tracking system environment.

The experiment was designed to be laboratory based and hence is of a smaller scale than the envisaged application. To inspect a large outdoor structure, such as a chimney, the Parrot AR.Drones could be replaced by commercial grade inspection drones such as the Falcon 8 from Ascending Technologies [37] and the Vicon MX positioning system could be replaced by GPS or for improved accuracy, differential GPS [38]. Pose uncertainty would increase when comparing a GPS tracking system with that of Vicon, but the control scheme can be adapted to compensate, with the collision avoidance threshold distance increased as well as increasing the safety margin on the set distance to target.

3.2. Inspection Setup

Given the size of the target and volume restrictions in the laboratory a two quadcopter system was chosen for performing the inspection. The Parrot AR.Drones are equipped with an ultrasound sensor for performing on-board height stabilisation control. Only two frequencies are available for the ultrasound, therefore a three-vehicle system in close proximity, although tested in the volume and capable of operating, is prone to disruptive ultrasonic sensor interference. Using a greater number of vehicles to inspect a larger target would significantly reduce inspection times. For this inspection, the two-vehicles required just over 100 s, which is a reduction of around a third in time taken compared with that required by a solitary vehicle. The inspection duration is also dependent on the initial position, before the vehicles are subject to autonomous control, and how often the intelligent collision avoidance is required. Even in a two drone scenario the variation in starting

conditions and the limited flight volume for the two vehicles often results in some, often minor, collision avoidance being required.

To complete the inspection, two vehicles takeoff and then enter autonomous flight, rising above the top of the drum to the first coverage band at a height of 1.5 m. As covered in Section 2.3, the CPP is achieved by using coverage bands, whereby each vehicle registers the position of the vehicle ahead and once all the vehicles have reached their registered position they transition to another band to repeat the process. For this case each band was separated by 30 cm with the vehicles landing after completing the final band at a height of 30 cm above the floor. Each band took 10-15 seconds to complete with five bands used for the drum inspection.

The visual inspection is carried out by the 720p HD camera, rigidly incorporated into the AR.Drone's main structure, with in-flight video recorded onto an on-board USB memory storage device. This device is then removed and the footage processed after the flight.

3.3. Photogrammetry Analysis

The recorded video footage is processed to enable the creation of a 3D model using photogrammetry analysis carried out by Autodesk's 123D Catch software [39]. First, frames were taken from the videos recorded by both of the quadcopters (70 frames were taken for this reconstruction). These frames were manually selected for their clarity at approximate, but consistent, time intervals to ensure the selection of around 35 frames per vehicle and complete drum coverage.

The distortion present in the selected frames, due to the lens on the Parrot AR.Drone 2.0, has to be corrected before proceeding. A calibration was carried out before the flight using the standard chessboard procedure [40], which estimates the distortion parameters from a number of images of a chessboard pattern taken from different angles at 1 m (nominal camera to target distance) from the vehicle camera [41]. This calibration enables objects that appear curved because of the lens distortion to be straightened to more accurately represent their real shape. Fig. 8 (a) shows a vehicle-recorded image and (b) is the post-processed output, where some of the image at the borders of (a) has been lost in the distortion correction process.

To successfully create an estimation of the drum's geometry using the Parrot AR.Drones, preparation of the drum and the surrounding area was required. This was necessary as the axisymmetric and monochromatic nature of the target provided very few features, required for image matching,
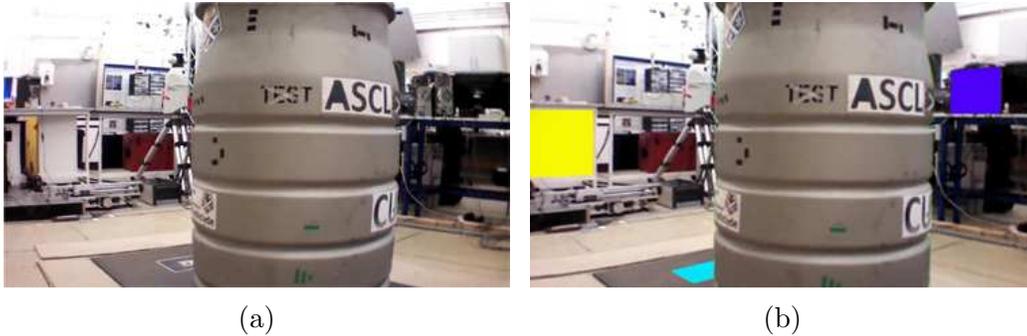
13

<center>(a)   (b)</center>

Figure 8: (a) Image from inspection footage and (b) distortion corrected image with artificial background markers.

when using AR.Drone's camera and resulted in image placement errors. The camera's image resolution (720p HD camera) and lack of isolation from the AR.Drone's motion, creating blur. Fig. 8 (b) also demonstrates that high contrast artificial markers, added in post-processing as in Fig. 8, were used successfully in conjunction with markers placed in the environment before the flight. However, attempts to place the drum in an entirely artificial environment, where the background is replaced with a single colour and high contrast markers, inserted in place of notable background objects, resulted in an unsuccessful model generation attempt.

A 3D surface-meshed model of the drum's lid and side is created as a product of the image matching and stitching with the recorded images mapped to the 3D mesh point cloud. The final textured CAD model, seen in Fig. 9, was achieved with manual input during the stitching process, where points could be selected that appeared in at least three images. The manual inputs and artificial markers had an influence on the final model by ensuring there were no gaps in the model construction and to prevent image placement errors that could produce significant but localised errors.

3.4. Model Error Analysis

The resultant mesh point cloud was exported to Geomagic Qualify [42] for metrological analysis. Due to the inherent lack of scale information available with single camera photogrammetry technology, it was concluded that a reliable scale approach and technique should be developed to ensure consistent and adequate model generation. The authors has access to the nominal CAD model, for the ILW storage drum, that allowed reference metrological

<center>14</center>

Figure 9: Generated 3D model produced from two-vehicle flight footage.

comparisons to be made at points and features of the drum. Therefore the point cloud model, generated from photogrammetry analysis and exported by Autodesk 123D Catch, was manually scaled, across a minimum of 10 points, according to the CAD model nominal data at the flange of the waste canister lid. This location was selected based on the density of coverage. Areas affected by poor density of coverage are the lower bottom section of the canister, due to the challenge of flying the vehicle at low altitudes close to the ground plane, and the topside of the lid due to the AR.Drone's camera angle being dependent on the pitch angle, which was nominally held parallel to the ground/horizontal plane.

The resultant scaled point cloud is then automatically aligned to the reference CAD model, using a least squares fitting approach. A comparison was then undertaken on every measured point within the cloud in reference to the nominal CAD model.

4. Results

4.1. 3D Model Errors

A baseline model was created with a 24.2-megapixel D-SLR camera [35] using similar shooting locations as achieved by the vehicles in flight. Additionally to quantify the effect of flight motion, a model was created by manually imaging the drum, from similar locations again, with an AR.Drone camera held by hand to keep it steady. Table 1, depicts the errors from each

15

of the three models where the Kolmogorov-Smirnov test determined that the data did not conform to standard normal distribution and, hence, is represented using interquartile range. The interquartile range defines where 50% of the samples reside with Q1 and Q3 marking the limits of this range. The baseline model, created with the D-SLR camera, provides an estimate for the errors resulting from the 3D model production software. This D-SLR produced model has the smallest error, and hence interquartile range, followed closely by the handheld AR.Drone camera. The interquartile range then increases to 28.74 mm for the model from in-flight footage. More details are given for this two-vehicle flight model, in Table 2, where the points beyond 60 mm deviation were discarded as outliers. In Fig. 10, a comparison is shown of the deviation distribution for the model produced by the two-vehicle flight and the handheld vehicle camera, which confirms the improved accuracy achieved with the steady camera.

Table 1: Comparison of Generated Model Point Clouds (Handheld Vehicle Camera and Two-Vehicle Flight

	Two-Vehicle Flight	Handheld	D-SLR Camera
Interquartile range [mm]	28.74	15.14	13.37
Q1 [mm]	-14.94	-7.36	-8.07
Q3 [mm]	13.80	7.78	5.30

Table 2: Analysis of Model Point Cloud Generated from Two-Vehicle Flight Footage

Maximum Upper Deviation [mm]	58.48	Maximum Lower Deviation [mm]	-59.95
No. of Data Points	5815	Outliers	75

4.2. Trajectory Errors

The focus of the AKF approach is to generate smooth trajectories automatically. This field then enables autonomous flight for multiple vehicles simultaneously. However, the trajectory errors are still of interest providing insight into the method's accuracy and the performance of the quadcopters in a restricted indoor environment.
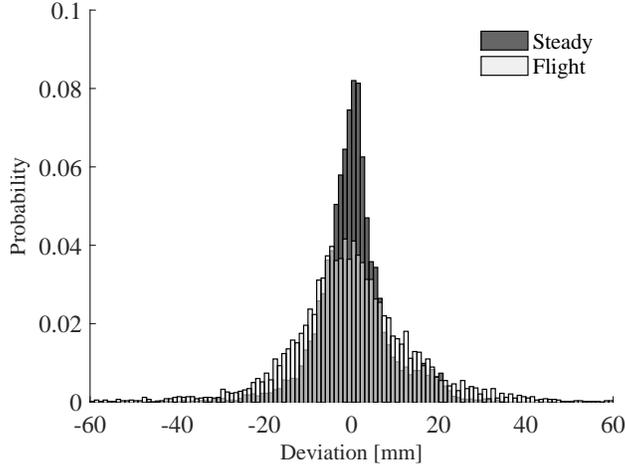
Figure 10: Deviation distribution for vehicle camera in-flight and held steady

For a drum inspection flight, where the vehicle is required to complete five coverage bands starting at 1.5 m and finishing at 0.3 m, the standard deviation of the error for each band is detailed in Table 3 where the errors in height position clearly increase with decreasing flight altitude. A similar trend can be seen in the radial error for one vehicle, but in the two-vehicle case this trend is not obvious due to complex aerodynamic interactions resulting in noisier data. The noise in the two-vehicle case is probably also responsible

Table 3: Standard Deviation of Positional Error from Multiple (1 & 2 Vehicle) Drum Inspection Flights

Nominal Height [mm]		1500	1200	900	600	300
1 Vehicle:	Mean of radial error Standard Deviation [mm]	134	168	202	219	197
1 Vehicle:	Mean of height error Standard Deviation [mm]	20	65	65	65	89
2 Vehicles:	Mean of radial error Standard Deviation [mm]	205	167	201	160	206
2 Vehicles:	Mean of height error Standard Deviation [mm]	38	67	75	85	104

for the reduced error in the 600 mm height case, when compared with one vehicle.

The starting position is not controlled, therefore the deviation in this may well have contributed to the large radial error in the two-vehicle case at height 1500 mm. The plot in Fig. 11 supports this claim with the initial positions seen to be offset from the radial path and the control method unable to prevent overshoot when attempting to maintain a radial distance. It also appears that the time to complete a coverage band decreases as the vehicles descend to lower altitudes with this fast coverage associated with a larger error elipsoid as depicted in Fig. 12.

4.3. Collision Avoidance

Due to the variable starting positions and turbulent flight environment, the collision avoidance mechanism is required even for a two-vehicle flight. In Fig. 13, the angular velocity of the trailing vehicle (determined from the angular position recorded by Vicon) is compared with the collision avoidance function H, see Fig. 4. It can be seen that there is a slight delay between command and execution but it is also clear that the rotational component of the vehicle's velocity decreases and begins rotating in the opposite direction in the places where the H falls below zero.
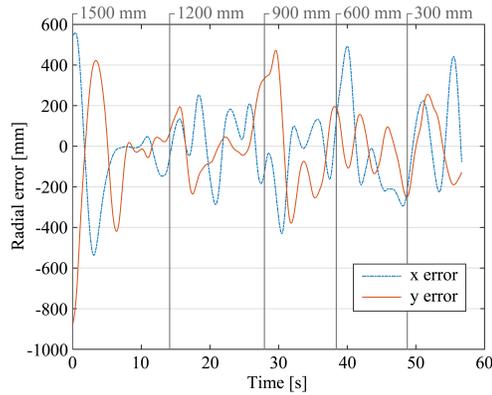

Figure 11: Representative flight data from a two-vehicle inspection flight with the altitude transitions marked above the plot.
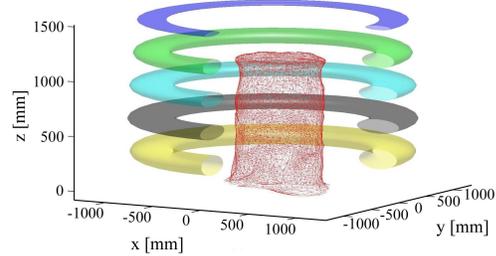

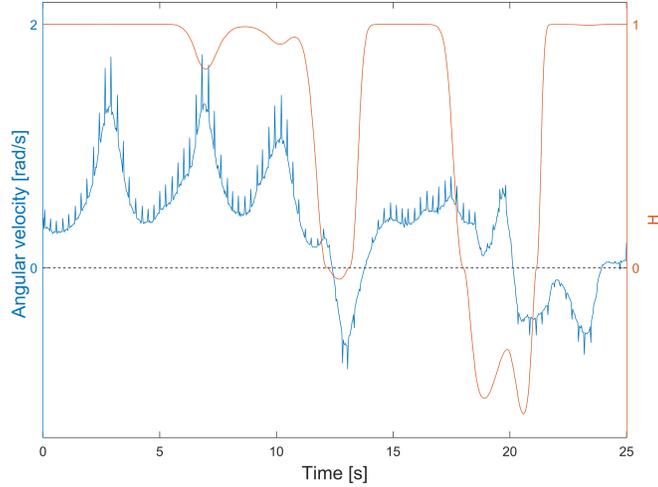Figure 12: Mean error flight elipsoids at the five coverage bands for one vehicle completing the drum inspection.

18

Figure 13: Comparison of flight data with collision avoidance function.

5. Discussion

5.1. Case Study

In the model, created from flight footage, the photogrammetry system performed worse in areas of sharp transition, at the lid edge and the four circumferential indentations. The difficulty in interpreting these ridges is, at least in part, due to the limited range of viewing angles achieved by the quadcopters that are flown roughly parallel to the floor. This restriction is a result of quadcopter dynamics with a manoeuvrable camera capable of achieving better coverage and, hence, a more accurate model.

Another source of error, also present in the handheld case, is the lens distortion, where errors will persist even after image correction. The most accurate corrections to the footage will be carried out on objects that are at the 1 m calibration distance, with slight distortions present elsewhere. However, a comparison between the handheld vehicle and D-SLR cameras highlights that the errors resulting from lens distortion and image resolution are far less significant than the errors resulting from the camera motion blur.

Table 1 shows that the D-SLR camera has a larger Q1 value, in terms of magnitude, but this is a result of the drum scaling process where the median value of the data set is slightly offset from zero. Therefore, the accuracy achieved with the D-SLR camera can be considered to be the upper limit of

19

what is achievable with the photogrammetry software, selected target and lighting conditions. The monochromatic and axisymmetric nature of the ILW storage drum make it a non-ideal target for photogrammetry based reconstruction. It is clear that with a more sophisticated vehicle the flight errors could have been reduced, the image quality improved and as a result the model errors reduced to be comparable with the D-SLR camera model errors.

5.2. Control

More precise positional control is possible with waypoint based trajectories, but the control precision is sacrificed to enable distributed and scalable control as discussed previously. That is not to say that the errors cannot be reduced as the P controller, described in Section 2, can be substituted with a PD to enhance the control performance, in particular the overshoot could be reduced. Currently the vehicles fly at a relatively low speed, which allows the desired pitch and roll to be commanded sufficiently on the basis of the AKF defined velocity only (see Eq. 13) rather than its difference with respect to the actual velocity. The introduction of a PD controller would instead consider the error in velocity in the horizontal plane and its derivative. This would introduce further refinements of the system that are beyond the scope of this paper. The error along the vertical axis and the yaw angle could also be controlled by a more straightforward PD controller.

5.3. Improvements & Automation

The authors manually added 2D-2D feature correspondence data to aid the reconstruction process used in 123D Catch. This step was necessary due to the low texture associated with the target which in turn was compounded by the compression used by the on-board camera. A fully automated system could be realised through modifications to the hardware and algorithms. Such as using a camera that produces uncompressed images, with potentially higher resolution, that would enable the feature extraction algorithms to better detect correspondences. For reference, the suitability of feature matching for plain surfaces found in stuctural inspection is considered in [30]. The availability of an on-board Inertial Measurement Unit gives rise to the potential for sensor fusion with the camera, as in [43], which would enhance the reconstruction process. In addition, recent state-of-the-art dense reconstruction methods, such as that reported in [44], make use of all pixels in the
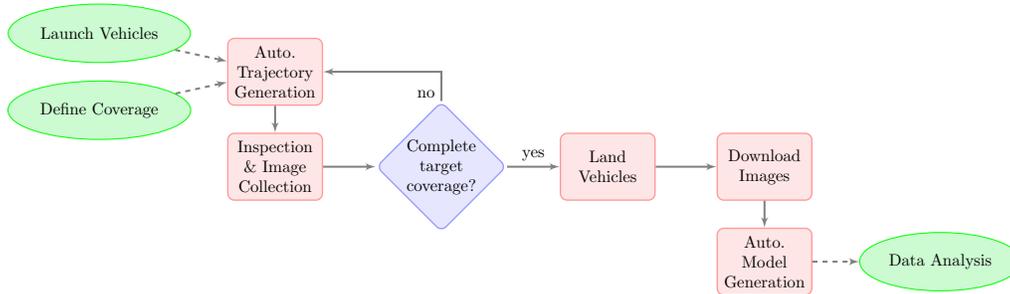
Figure 14: Flowchart of a feasible fully automated version of the system. Green requires user input, red is autonomously executed and blue is a decision process.

images rather than a sparse set of points. Such approaches are more suited for low texture environments.

Therefore, a fully automated system would be similar to that presented in Fig. 14. It is envisioned that this system would carry on generating trajectories until making the decision autonomously that coverage was complete, with target coverage requirements being manually inputted before the flight. The use of wireless data transfer and a higher resolution imager, with some form of image stabilisation, would enable full automation of the model generation process. Finally, a combination of automated and manual inspection of the final data would probably be necessary to complete the inspection.

6. Conclusion

A scalable and autonomous control scheme has been presented for CPP applied to aerial inspection of a central target. The strengths of the proposed technique are in its relatively simple control computation requirements, contributing to its distributable and scalable capabilities that include vehicle and target collision avoidance, whilst providing complete coverage with reduced inspection times. The scheme has been implemented for two vehicles, where the collision avoidance was tested with the vehicles forming a queue while maintaining their defined distance from the target. A two-vehicle system performed an inspection on a nuclear ILW storage drum. The inspection footage was processed, with the aid of high contrast artificial markers in the background and actual markers attached to the drum, to successfully generate a 3D surface-meshed model. The results achieved highlight that this autonomous trajectory generation approach is applicable for remote inspec-

tion with system improvements primarily dependant on hardware and not control limitations.

Acknowledgment

[1] F. Fraundorfer, Building and site reconstruction from small scale unmanned aerial vehicles (UAV's), in: Urban Remote Sensing Event (JURSE), 2015 Joint, IEEE, 1–4, 2015.

[2] S. Esposito, P. Fallavollita, W. Wahbeh, C. Nardinocchi, M. Balsi, Performance evaluation of UAV photogrammetric 3D reconstruction., in: IGARSS, 4788–4791, 2014.

[3] F. Bonnin-Pascual, E. Garcia-Fidalgo, A. Ortiz, Semi-autonomous visual inspection of vessels assisted by an unmanned micro aerial vehicle, in: Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on, IEEE, 3955–3961, 2012.

[4] J. Nikolic, M. Burri, J. Rehder, S. Leutenegger, C. Huerzeler, R. Siegwart, A UAV system for inspection of industrial facilities, in: Aerospace Conference, 2013 IEEE, IEEE, 1–8, 2013.

[5] M. Burri, J. Nikolic, C. Hurzeler, G. Caprari, R. Siegwart, Aerial service robots for visual inspection of thermal power plant boiler systems, in: Applied Robotics for the Power Industry (CARPI), 2012 2nd International Conference on, IEEE, 70–75, 2012.

[6] A. Ortiz, F. Bonnin-Pascual, E. Garcia-Fidalgo, Vessel Inspection: A Micro-Aerial Vehicle-based Approach, Journal of Intelligent & Robotic Systems 76 (1) (2014) 151–167.

[7] R. Khanna, M. Moller, J. Pfeifer, F. Liebisch, A. Walter, R. Siegwart, Beyond point clouds-3D mapping and field parameter measurements using UAVs, in: Emerging Technologies & Factory Automation (ETFA), 2015 IEEE 20th Conference on, IEEE, 1–4, 2015.

[8] F. Remondino, L. Barazzetti, F. Nex, M. Scaioni, D. Sarazzi, UAV photogrammetry for mapping and 3d modeling–current status and future perspectives, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences 38 (1) (2011) C22.

[9] J. Liu, H. Wang, X. Liu, F. Li, G. Sun, P. Song, An automated 3D reconstruction method of UAV images, in: Applied Optics and Photonics China (AOPC2015), International Society for Optics and Photonics, 96780S–96780S, 2015.

[10] J. S. Bellingham, M. Tillerson, M. Alighanbari, J. P. How, Cooperative path planning for multiple UAVs in dynamic and uncertain environments, in: Decision and Control, 2002, Proceedings of the 41st IEEE Conference on, vol. 3, IEEE, 2816–2822, 2002.

[11] M. Shanmugavel, A. Tsourdos, B. White, R. Żbikowski, Co-operative path planning of multiple UAVs using Dubins paths with clothoid arcs, Control Engineering Practice 18 (9) (2010) 1084–1092.

[12] E. Galceran, M. Carreras, A survey on coverage path planning for robotics, Robotics and Autonomous Systems 61 (12) (2013) 1258–1276.

[13] L. Babel, Flight path planning for unmanned aerial vehicles with landmark-based visual navigation, Robotics and Autonomous Systems 62 (2) (2014) 142–150.

[14] A. Barrientos, J. Colorado, J. d. Cerro, A. Martinez, C. Rossi, D. Sanz, J. Valente, Aerial remote sensing in agriculture: A practical approach to area coverage and path planning for fleets of mini aerial robots, Journal of Field Robotics 28 (5) (2011) 667–689.

[15] G. Murtaza, S. Kanhere, S. Jha, Priority-based coverage path planning for aerial wireless sensor networks, in: Intelligent Sensors, Sensor Networks and Information Processing, 2013 IEEE Eighth International Conference on, IEEE, 219–224, 2013.

[16] O. Khatib, Real-time obstacle avoidance for manipulators and mobile robots, The international journal of robotics research 5 (1) (1986) 90–98.

[17] D. F. Gordon, W. M. Spears, O. Sokolsky, I. Lee, Distributed spatial control, global monitoring and steering of mobile agents, in: Information

23

Intelligence and Systems, 1999. Proceedings. 1999 International Conference on, IEEE, 681–688, 1999.

[18] W. M. Spears, D. F. Gordon, Using artificial physics to control agents, in: Information Intelligence and Systems, 1999. Proceedings. 1999 International Conference on, IEEE, 281–288, 1999.

[19] N. E. Leonard, E. Fiorelli, Virtual leaders, artificial potentials and coordinated control of groups, in: Decision and Control, 2001. Proceedings of the 40th IEEE Conference on, vol. 3, IEEE, 2968–2973, 2001.

[20] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, R. E. Davis, Collective motion, sensor networks, and ocean sampling, Proceedings of the IEEE 95 (1) (2007) 48–74.

[21] D. A. Lawrence, E. W. Frew, W. J. Pisano, Lyapunov vector fields for autonomous unmanned aircraft flight control, Journal of Guidance, Control, and Dynamics 31 (5) (2008) 1220–1229.

[22] E. W. Frew, D. A. Lawrence, S. Morris, Coordinated standoff tracking of moving targets using Lyapunov guidance vector fields, Journal of guidance, control, and dynamics 31 (2) (2008) 290–306.

[23] D. J. Bennet, C. McInnes, Verifiable control of a swarm of unmanned aerial vehicles, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering 223 (7) (2009) 939–953.

[24] S. B. McCamish, M. Romano, X. Yun, Autonomous distributed control of simultaneous multiple spacecraft proximity maneuvers, Automation Science and Engineering, IEEE Transactions on 7 (3) (2010) 630–644.

[25] C. Vasile, A. Pavel, C. Buiu, Integrating human swarm interaction in a distributed robotic control system, in: Automation Science and Engineering (CASE), 2011 IEEE Conference on, IEEE, 743–748, 2011.

[26] C. Harris, M. Stephens, A combined corner and edge detector., in: Alvey vision conference, vol. 15, Citeseer, 50, 1988.

[27] P. Hansen, H. Alismail, P. Rander, B. Browning, Pipe mapping with monocular fisheye imagery, in: Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on, IEEE, 5180–5185, 2013.

[28] S. El Kahi, D. Asmar, A. Fakih, J. Nieto, E. Nebot, A vison-based system for mapping the inside of a pipe, in: Robotics and Biomimetics (ROBIO), 2011 IEEE International Conference on, IEEE, 2605–2611, 2011.

[29] J. Shi, C. Tomasi, Good features to track, in: Computer Vision and Pattern Recognition, 1994. Proceedings CVPR'94., 1994 IEEE Computer Society Conference on, IEEE, 593–600, 1994.

[30] G. Dobie, R. Summan, C. MacLeod, S. G. Pierce, Visual odometry and image mosaicing for NDE, NDT & E International 57 (2013) 17–25.

[31] M. L. Brutto, P. Meli, Computer vision tools for 3D modelling in archaeology, International Journal of Heritage in the Digital Era 1 (1 suppl) (2012) 1–6.

[32] T. P. Kersten, M. Lindstaedt, Image-based low-cost systems for automatic 3D recording and modelling of archaeological finds and objects, in: Progress in cultural heritage preservation, Springer, 1–10, 2012.

[33] Vicon, Vicon motion tracking system, `http://www.vicon.com/Software/Tracker`, [Online; accessed January-2015], 2015.

[34] Parrot, Parrot AR.Drone 2.0 Technical Specifications, `http://ardrone2.parrot.com/ardrone-2/specifications//`, [Online; accessed January-2015], 2015.

[35] Nikon, Nikon D3200, `http://www.europe-nikon.com/en\_GB/product/digital-cameras/slr/consumer/d3200/`, [Online; accessed January-2015], 2015.

[36] AF-S DX NIKKOR 18-55mm f/3.5-5.6G VR, `http://www.europe-nikon.com/en\_GB/product/nikkor-lenses/auto-focus-lenses/dx/zoom/af-s-dx-nikkor-18-55mm-f-3-5-5-6g-vr`, [Online; accessed January-2015], 2015.

[37] Ascending Technologies, `http://www.asctec.de/`, [Online; accessed February-2015], 2015.

[38] P. Misra, P. Enge, Global Positioning System: Signals, Measurements and Performance Second Edition, Lincoln, MA: Ganga-Jamuna Press, 2006.

[39] Autodesk, Autodesk 123D Catch, `http://www.123dapp.com/catch`, [Online; accessed February-2015], 2015.

[40] P. Azad, T. Gockel, R. Dillmann, Computer Vision: principles and practice .

[41] Autodesk, Calibration toolbox, `http://www.123dapp.com/catch`, [Online; accessed February-2015], 2015.

[42] Geomagic Control, `http://www.geomagic.com/en/products/qualify/overview`, [Online; accessed January-2015], 2015.

[43] J. Kelly, G. S. Sukhatme, Visual-inertial sensor fusion: Localization, mapping and sensor-to-sensor self-calibration, The International Journal of Robotics Research 30 (1) (2011) 56–79.

[44] J. Engel, T. Schöps, D. Cremers, LSD-SLAM: Large-scale direct monocular SLAM, in: Computer Vision–ECCV 2014, Springer, 834–849, 2014.