J Intell Robot Syst (2011) 61:251–265
DOI 10.1007/s10846-010-9496-6

Vision-Based Landing of Light Weight Unmanned Helicopters on a Smart Landing Platform

Mohammad H. Mahoor · R. Godzdanker · K. Dalamagkidis · K. P. Valavanis

Received: 1 February 2010 / Accepted: 1 September 2010 / Published online: 6 November 2010

Abstract This paper presents a simple and efficient solution to vision guided autonomous landing of a light-weight (<150 Kg) unmanned helicopter on a smart landing platform, called *ISLANDS*—*I*ntelligent *S*elf-*L*eveling *a*nd *N*odal *D*ocking *S*ystem. The advantage of ISLANDS is that it may allow the helicopter upon docking to recharge its batteries or refuel, thus, indirectly increasing endurance and flight range. In order for the helicopter to dock with ISLANDS, an on-board 'vision module' coupled with the helicopter attitude controller is developed. This 'vision module' detects the location and orientation of ISLANDS and feeds back information to the helicopter attitude controller, which commands the helicopter to descent onto the landing platform at a desired orientation and speed. The Scale Invariant Feature Transform (SIFT) is used for automatic detection of the landing platform based on images captured by a single camera mounted on the helicopter. The detected SIFT features are used to estimate the 3-D orientation of the platform relative to the helicopter using Homography and RANSAC techniques. The focus of this paper is on the vision-guided landing technique in a predefined orientation and not on controller details, which may be found in Shim et al. (1998).

Keywords Unmanned helicopter · Landing platform · Vision-guided landing

1 Introduction

Most existing experimental light weight unmanned helicopters have a common limitation related to endurance and mission range. To overcome such limitations, the idea of docking an unmanned helicopter on a landing platform, although very

M. H. Mahoor (✉) · R. Godzdanker · K. Dalamagkidis · K. P. Valavanis
Unmanned Systems Laboratory Electrical and Computer Engineering,
University of Denver, Denver, CO 80208, USA
e-mail: mmahoor@du.edu

challenging and difficult, offers a viable solution to the problem. One such system was proposed by SPAWAR system where the landing and docking system were designed to work with only one type of ducted fan vertical take-off and landing (VTOL) vehicle, the iSTAR MAV [2], thus, making it VTOL-specific and not versatile.

The proposed *ISLANDS* (*I*ntelligent *S*elf-*L*eveling *a*nd *N*odal *D*ocking *S*ystem) is the outgrowth of the work initially proposed by Dalamagkidis et al. [3, 4], which is a more versatile platform capable of accommodating different families of light weight helicopters such as the *RAPTOR* series, the *Maxi Joker*, the *BERGEN Observer/ Twin*, and the *Yamaha R-50*, to name just a few types (http://www.rotomotion.com/). All of these helicopters are standard rotary craft helicopters composed of a main and tail rotor. ISLANDS can be easily modified to be a standalone node or mounted on top of a UGV (Unmanned Ground Vehicle) or any other ground vehicles. Additionally ISLANDS has the ability to 'level with the horizon' regardless of the ground vehicle orientation (on or off the road) as shown in Fig. 1.

In order to benefit from ISLANDS or any other docking system the autonomous helicopter must land in the correct orientation; this is achieved with the aid of a comprehensive on-board vision system that includes a vision module coupled with the helicopter attitude controller. In essence, the on-board vision system is capable of functioning in either of two modes: (1) Independently of the helicopter navigation controller; (2) In a coupled mode with the navigation controller, assisting in vision-guided waypoint navigation, collision avoidance, etc. However, in this paper, the objective is centered on vision-based autonomous landing of a light weight helicopter for the purpose of improving (indirectly) endurance and extending the mission range of autonomous helicopters. The extended mission range is achieved because the helicopter can refuel mid mission, hence the autonomous helicopter no longer has to conserve fuel for the return trip. Long-term, it is envisioned that a set of such landing platforms (a set of ISLANDS) will function as a 'sensor network' as shown in Fig. 2. Since in the future ISLANDS themselves will be equipped with long range communication capabilities.

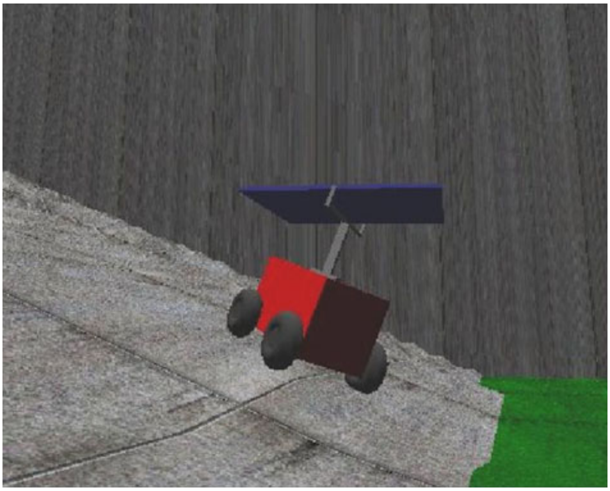

Fig. 1 A 3-D representation of ISLANDS mounted on top of an ATRV-Jr leveling with the horizon regardless of the orientation of the UGV

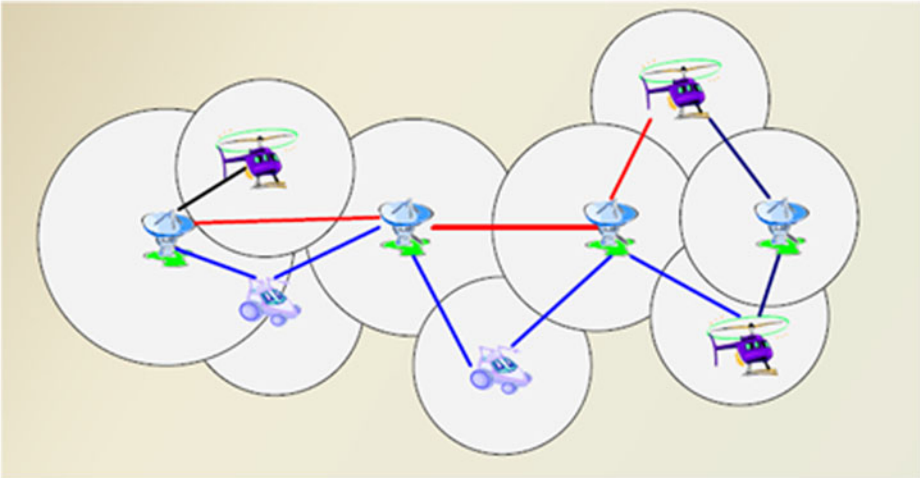

Fig. 2 A network of unmanned systems with landing platforms. It may function as a swarm and/or as a wireless sensor network

Considerable research has been conducted in the area of vision assisted landing. In the early 1990s, Dickmanns and Schell presented a vision system for landing an airplane [5]. Garcia-Pardo et al. proposed a vision-based system for finding safe landing sites on the ground. Reported research in [6] refers to vision-based techniques for landing on an artificial helipad of known shape with slow movement. Caballero et al. [7] proposed a visual odometry and SLAM for medium and high altitude UAVs based on monocular imagery. They used Homographic models and Homography decomposition to extract the real camera motion and a range sensor to obtain the scale factor of the motion. Researchers at Berkeley [8, 9] have used computer vision techniques to estimate the pose of a UAV relative to a planar landing target attached on a moving deck. Multiple view geometry techniques were used to compute the real motion of the UAV with respect to the planar landing target based on rank deficiency of the so called multiple view matrix. For feature matching, they used a landing target with a known rectangular geometry such that the corners of the rectangles can easily be extracted. Simple image processing techniques such as thresholding grayscale images to binary images, image segmentation, and corner detection using simple connected components labeling algorithms were used for feature point extraction and matching.

Compared to the related research, the presented approach in this paper uses single view geometry to estimate the 3-D orientation of the helicopter relative to a planar landing target. Scale Invariant Feature Transform (SIFT) is used for feature extraction and matching between the captured images of the landing platform on the ground and a reference image of the landing platform with a know orientation. This implies that the geometric pattern on the landing platform is not of importance and that this method may be applied to a planar landing platform with any shape or geometry. The navigation controller is not considered explicitly in this paper; however, it is assumed that any developed nominal controller such as [1] may be used as the baseline controller to fly the unmanned helicopter.

The remainder of this paper is organized as follows. Section 2 presents an overview of the proposed vision and platform integration. The vision-guided approach for detecting the landing platform and its orientation is described in Section 3. Section 4 discusses the ISLANDS. Results are shown in Section 5, followed by conclusions and future research in Section 6.

2 System Overview

The goal of the developed vision module is to detect the orientation of a planar target on the ground using visual data (i.e., regular optical images) relative to the helicopter. Then feed this information to the helicopter attitude controller which controls the decent onto the ground aligning the helicopter with the planer target. In the work presented in this paper the planer target is located on the landing surface of ISLANDS leading to an aligned landing of the autonomous helicopter. Once landed in the correct orientation ISLANDS can then initiate a centering routine of the helicopter which accommodates for the errors associated with vision system and attitude controllers. Once oriented and centered ISLANDS docks with the helicopter and re-fuels/re-chargers or performs data exchange as necessary. It is envisioned that the unmanned helicopter will fly using GPS via point to the vicinity of an ISLANDS landing platform at which point the vision system will take over and allow for a safe landing/docking. Figure 3 depicts a block diagram overview of the developed vision system and Fig. 4 shows the first rough prototype of the landing platform surface with orientation markers.

As seen in Fig. 3 the vision system is capable of outputting the correction needed in all Euler angles; roll pitch and yaw. During landing only yaw is necessary as changes in roll and pitch will cause the helicopter to fly either sideways or forward and backwards depending on the sign on the correction of roll and pitch. The values of roll and pitch during testing were observed to be always near zero since the position of the helicopter was placed directly above the landing platform with only a yaw offset. For this reason our experiments only deal with yaw perturbations and noise is only injected into the yaw sensor (we refer the reader to Sections 4 and 5 for more detail on this).

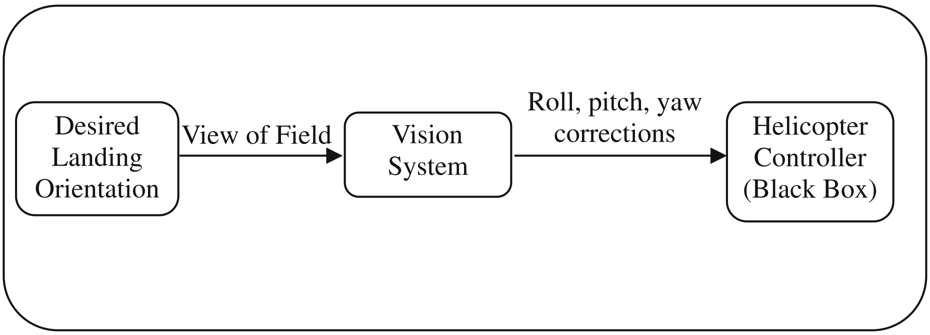

Fig. 3 A general block diagram of the autonomous vision-based landing system

Fig. 4 The landing platform with a T label used in this paper

3 Platform Detection Using SIFT

In order to estimate the orientation of the target landing platform with respect to a reference frame using Homography, we need to find a suite of corresponding points between a reference image and the platform images captured by a camera mounted on the bottom of the helicopter. In this paper we employ SIFT [10] for feature extraction and image matching. SIFT has been successfully used in various computer vision applications such as object recognition, robotic mapping and navigation, image stitching, 3-D modeling, gesture recognition, and video tracking [12–14]. The SIFT features are local and based on the appearance of the object at particular interest points. In addition, they are invariant to rigid transformation (translation, scale, and rotation) and illumination variations. They are well localized in both spatial and frequency domains, thus reducing the chance of disruption by occlusion, noise, and clutter. Based on this approach, a set of key points, called SIFT keys are selected from an image by multi-resolution analysis. Afterwards, at each key point, 2-D gradient magnitude, $m(x, y)$ and orientation $\theta(x, y)$ for each image resolution are calculated [11]:

$$m(x, y) = \sqrt{L_x^2 + L_y^2}, \theta(x, y) = \tan^{-1}\left(\frac{L_y}{L_x}\right) \tag{1}$$

where L_x and L_y are the approximate differences computed as $L(x+1, y) - L(x-1, y)$ and $L(x, y+1) - L(x, y-1)$, respectively.

At each keypoint, the magnitude and orientation of the image gradients are used to create a descriptor -feature vector- for matching the keypoints between two images. To create the keypoint descriptor, first a Gaussian weighting function with equal to one half the width of the descriptor window is used to assign a weight to the magnitude of each sample point. Second, the weighted magnitudes are used to create orientation histograms over 4×4 sample regions. Then, the descriptor is formed from a vector containing the values of all the orientation histogram entries, corresponding to the lengths of the gradient arrow.

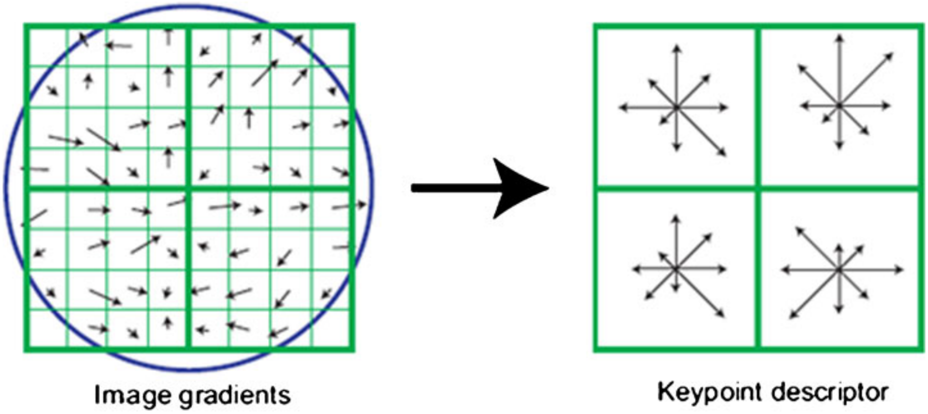

Fig. 5 The gradient magnitude and orientation at each image sample point in a region around the keypoint location is shown on the left figure. These are weighted by a Gaussian window, indicated by the overlaid circle. These samples are then accumulated into orientation histograms summarizing the contents over 4 × 4 sub-regions, as shown on the right, with the length of each arrow corresponding to the sum of the gradient magnitudes near that direction within the region. The experiments in this paper use 4 × 4 descriptors computed from a 16 × 16 sample array [11]

Figure 5 shows a 2 × 2 array of orientation histograms, with 8 orientation bins in each. In our experiments, a 4 × 4 × 8 = 128 element feature vector for each keypoint has been used. The best candidate match for each keypoint between two images (reference image, I_r, and target image I_t), is found by identifying its nearest neighbor of the detected keypoints from frame I_t. We refer our reader to [11] for further detail on descriptor matching.

The coordinates of the descriptor and the gradient orientations are rotated relative to the keypoint orientation to achieve orientation invariance. The SIFT descriptors are also semi-invariant under the affine transformation and slightly affected by changes in the camera angle. Invariant features are very important since the UAS flies in a low altitude above the ground and the 2-D captured images are distorted. In other words, the images captured by the UAS camera have different views at different roll/pitch/raw orientation of the helicopter. Thus, the SIFT features are robust in matching images captured in a range of camera angles.

Figure 6 illustrates two images of the target landing platform used in this paper. The green points represents the position of the SIFT key points extracted from each image and the blue lines connect every pair of the corresponding matches established using SIFT.

We find a set of corresponding points (matched keypoints) between every video frame, I_t, and a reference planar frame, I_r, and use them for robust estimation of the rotation and translation parameters (extrinsic parameters):

$$P_c = \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} = RP_w + T; \quad p = \begin{bmatrix} x_c \\ y_c \end{bmatrix} = \frac{f}{Z_c} \begin{bmatrix} X_c \\ Y_c \end{bmatrix} \tag{2}$$

where matrix R is a 3×3 rotation matrix, and T is a 3×1 translation vector; P_c is the 3D point coordinates on the planar platform in camera coordinate system; p is

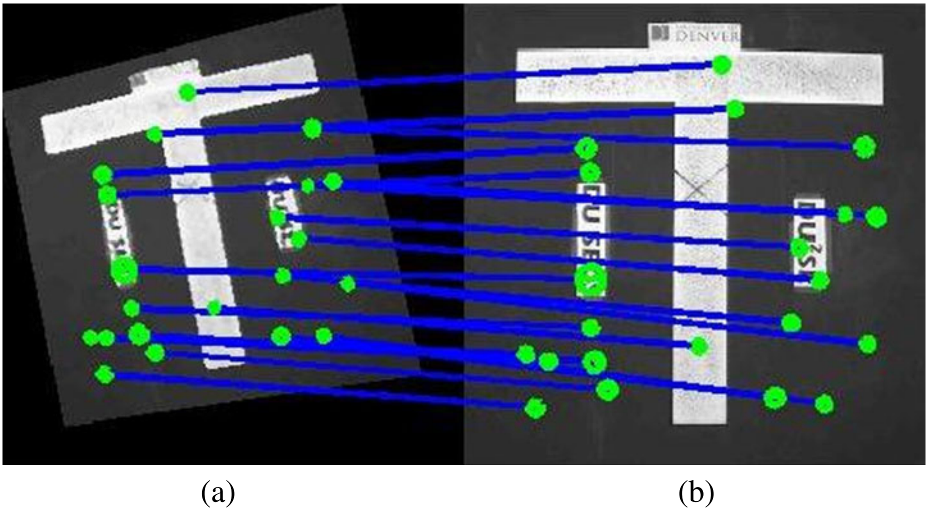

Fig. 6 SIFT correspondences between a reference image of the platform (roll = 0, pitch = 0, yaw = 0) (**b**) and a rotated version (roll = 0, pitch = −.1, yaw = + .2). The green points are the SIFT key points and the blue lines represent the matches between the points calculated using SIFT

the image point coordinates in image I_t, $P_w = [X_w, Y_w, Z_w]'$ is the 3D point coordinates on the planar platform I_r in world coordinate system; f is the camera focal length. Since the platform is planar, therefore, Z_w can be constant (e.g., zero).

In this paper, we used Caltech camera calibration toolbox (http://www.vision.caltech.edu/bouguetj/calib_doc/) to estimate the extrinsic parameters, R and T, using the SIFT corresponding points (at least four corresponding points are needed). However, there is a high possibility that the SIFT matches are not true matches (i.e., outliers). Thus, we use Random Sample Consensus (RANSAC) algorithm [15] to eliminate the outliers and find a suite of true corresponding key points (i.e., the inliers) that best estimates the transformation matrices. RANSAC technique is a well known iterative algorithm technique in computer vision to estimate parameters of a mathematical model (i.e., transformation) from a set of observed data/correspondences which contains outliers. Once the matrix R is estimated using the SIFT matches and the RANSAC algorithm, the roll, yaw and pitch angles can readily be extracted from the rotation matrix R.

4 ISLANDS

The idea behind ISLANDS is to provide a self leveling mobile platform that can either stand alone in the field or be mounted to a surface/air vehicle. A level platform is required for a helicopter due to rotor dynamics effects of unlevel surfaces and the simple case that the helicopter could slip off the platform if not level. For a platform to be capable of self leveling with the horizon, it will require a minimum of two degrees of freedom; one degree of freedom to align the platform with the incline, which will be called the azimuth correction axis, and one degree of freedom to correct the incline or elevation.

Table 1 Existing autonomous helicopters with their respective rotor diameters and power source

	Main rotor diameter (inches)	Power source
Yamaha R-50	125	Gas
Bergman observer	80	Gas
Maxi joker 2	71	Electric
Rototmotion platforms		
SR200	118	Gas
SR100	79	Gas
SR30	78	Gas
SR20	75	Electric

As stated previously a level landing platform is important for any helicopter so even upward thrust is kicked back up to the rotors; this is referring to the ground effect phenomenon associated with any VTOL vehicle. In addition to being level the platform must be large enough such that a vertical projection of the rotor area is completely on the platform area. In making ISLANDS as versatile as possible handling autonomous helicopters of up to 150 kg or 333 lbs, a survey of existing autonomous helicopters platforms was performed examining in particular their main rotor diameter. Table 1 presents the rotor diameter of existing autonomous helicopters, hence representing the minimum dimension for the top deck ISLANDS.

As one of the design objectives of ISLANDS to be able to recharge electric helicopters a way of generating electricity and storing that energy on board ISLANDS is required. By designing ISLANDS to accommodate up to a Yamaha Rmax sized helicopter leads to landing deck with a surface area of 10 m^2 or 100 ft^2, as the main rotor diameter of the Yamaha Rmax is 3.2 m. Having such a large surface area led us to the design decision of implementing photovoltaic cells on the surface which could conservatively produce 1200 W. During times when ISLANDS is not docking with a helicopter on board batteries can be charged allowing for longer deployment periods of ISLANDS and resources of re-charging helicopters.

Figure 7 depicts the initial design proposed for ISLANDS by Dalamagkidis [3, 4] and is a two DOF mechanism using two actuators one on the base to align the platform with the azimuth and one at the center of the platform to level the platform with the elevation. This design is very similar to a pan and tilt unit used for cameras.

This initial design will work for light loads but when this design scaled up to the 150 kg future load requirement is not the most efficient. The part of the design that is inefficient is the location of the actuator for leveling with the gradient. As can be seen in Fig. 7 the leveling to the gradient actuator is located at the pivot point of the landing deck. Hence if the deck becomes large and the helicopter was to land further anywhere other than the center of the deck the moment created will become

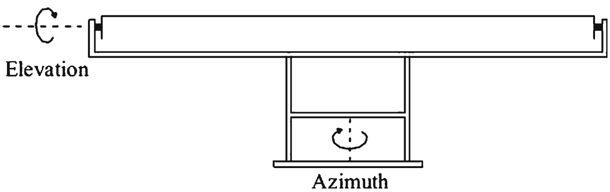

Fig. 7 Cross sectional view of original gimble design of ISLANDS proposed by Dalamagkidis

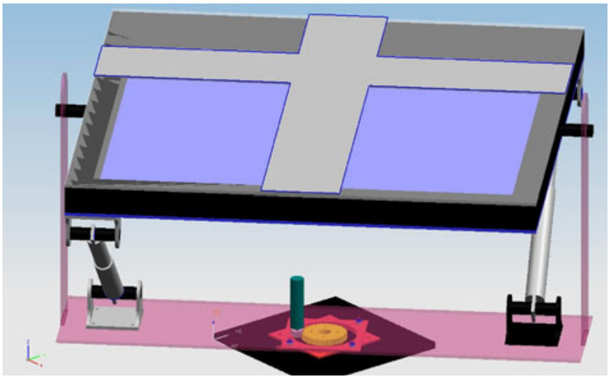

Fig. 8 CAD design of ISLANDS prototype capable automatically leveling to horizon with landing marker

large hence requiring a very powerful actuator. Additionally even small vibrations at the end of the landing decks will translate to large torques that will be need to be overcome by the motor. For this reason an alternative design of ISLANDS is proposed and built according to the cad drawing shown in Fig. 8.

The design depicted in Fig. 8 still uses two DOF but the single actuator that was previously placed on the center of the landing surface has been replaced by two linear actuators that are connected to the base of the platform. This important modification provides a more stable landing surface since the forces are now distributed along the vertical supports of the landing platform, the actuator itself and back to the ground. This design option additionally provides redundancy incase one of the actuators fails. This design still maintains the pivot of the elevation degree of freedom at the center of the platform and the azimuth correction is done by using a turn table with a gear system.

The platform originally proposed in [3, 4] is designed to interact with the ATRV-Jr receiving IMU (Inertial Measurement Unit, which provides attitude data) data to level out the platform. This design requires that any system that ISLANDS is attached must have some tilt sensors on board. Otherwise if ISLANDS is to be self contained and a bolt on node to any fixed or moving system it must have an IMU or other means of determining orientation on board.

5 Experimental Setup and Results

Matlab Simulink was used to test whether the rotation matrix produced by the SIFT pattern matching algorithm can be successfully integrated with an autonomous helicopter controller for the purpose of alignment with a desired surface target. Simulink was used to integrate the SIFT based vision system with a model of an autonomous helicopter attitude controller shown in Fig. 9. The VLFeat open source vision library was utilized for SIFT feature extraction and matching (http://www.vlfeat.org/). For the detail of the attitude controller, we refer the readers to [1].

The physical system implementation of this vision system integrated with an autonomous helicopter would have continuous video input to the ‘vision module’ captured by an on-board camera. However, this initial work is being performed

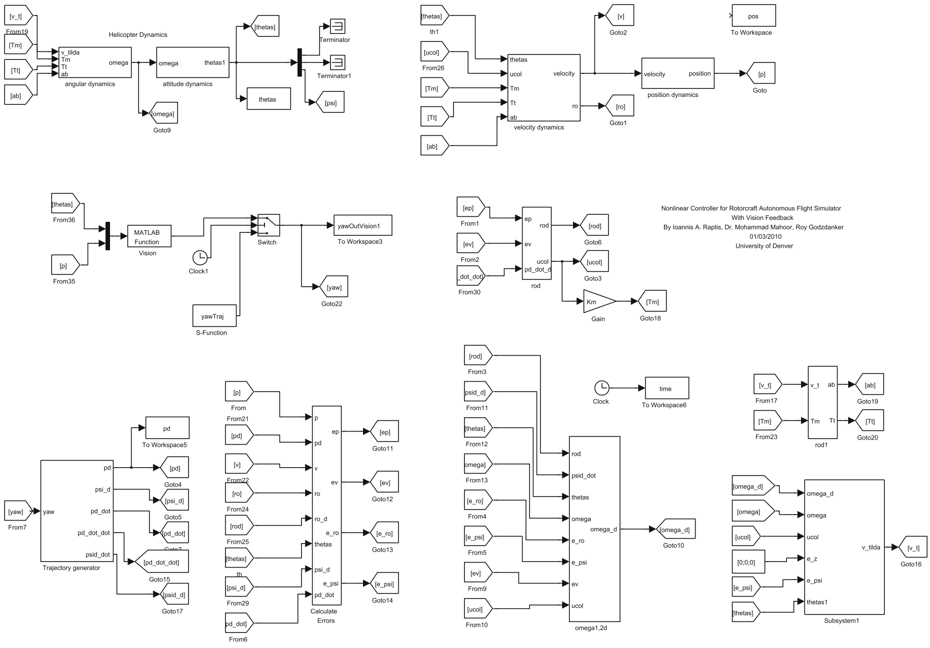

Fig. 9 Simulink model

in Matlab to model continuous updated images captured by a virtual camera. The virtual camera projects an original image—image in Fig. 4—based on the X, Y, Z, and the roll, pitch, yaw angles specified by the current helicopter position. These attitude parameters where extracted from the helicopter model. The elements blocked off in red in the Simulink model of Fig. 9 are the ones associated with the vision module. On the left side inputs of X, Y, Z position and roll, pith, yaw body orientation of the simulated helicopter are fed to the vision module. The output of the vision module is the rotation matrix required to align the camera and therefore the helicopter with the platform. The helicopter controller maintains a zero degree pitch and roll representing level flight during hovering or elevation change. This means that these angles are essentially zero (less than .01° in radian) as shown in Fig. 10. However, the vision system is capable of estimating the pitch and roll angles with a high accuracy based on SIFT features.

The roll and pitch angles are maintained by the helicopter controller at zero therefore only the corrected yaw angles from the 'vision module' are required as input to the helicopter attitude controller, as seen by the elements blocked off in blue.

The experiment was conducted as follows. The model is initially set to navigate to a predetermined altitude and yaw angle that is not aligned with the landing platform. Once this yaw angle and altitude is achieved a switch in the model is triggered to start using the vision system to drive the helicopter into an orientation that is aligned with the landing platform (the vision module is functioning all the time during the simulation). The landing platform is assumed to be at a yaw angle of zero, for this reason the next yaw inputted to the controller from the vision system is the negative

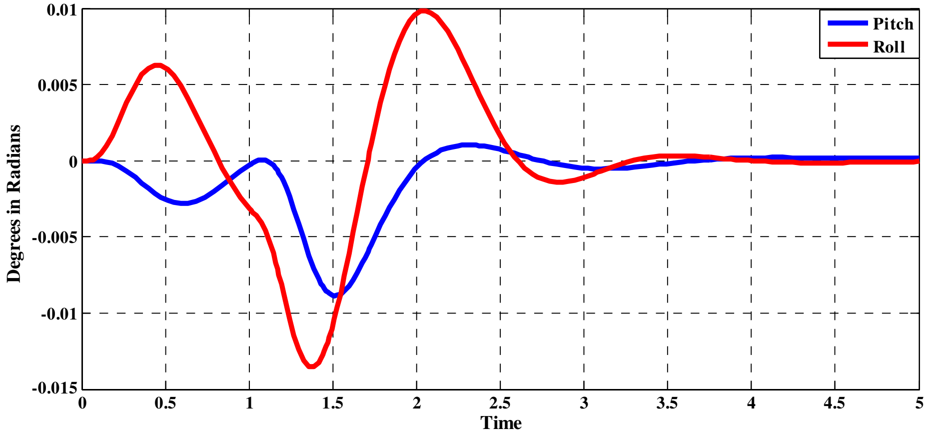

Fig. 10 Roll pitch angle maintained at zero by the helicopter controller

of the estimated yaw angle by the vision module. This then drives the helicopter to a yaw angle of zero aligning the helicopter with the landing platform.

Figure 11 shows the result of one of these experiments where the system is initially driven to a yaw angle of 45° and then the vision system is turned on at t = 1 s. The blue line represent the current yaw angle of the helicopter while the red line represent the input yaw angle being generated by the 'vision module'.

As can be seen in Fig. 11 the vision system was initiated at t = 1 s which is exactly when the helicopter model reached the desired position of 45°. Figure 12 also illustrates a 45° initial yaw angle but waits for t = 2 s before initiating the vision system hence the plateau seen in the figure. The spikes seen in the vision yaw response curve are due to the outliers in the SIFT matching step which causes the estimated yaw angle to be erroneous. However, these small errors can be eliminated

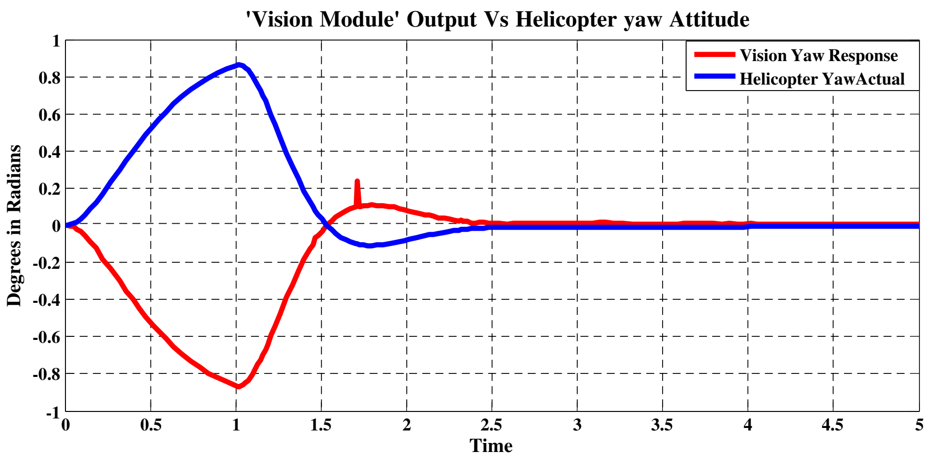

Fig. 11 Roll pitch angle maintained at zero by helicopter controller

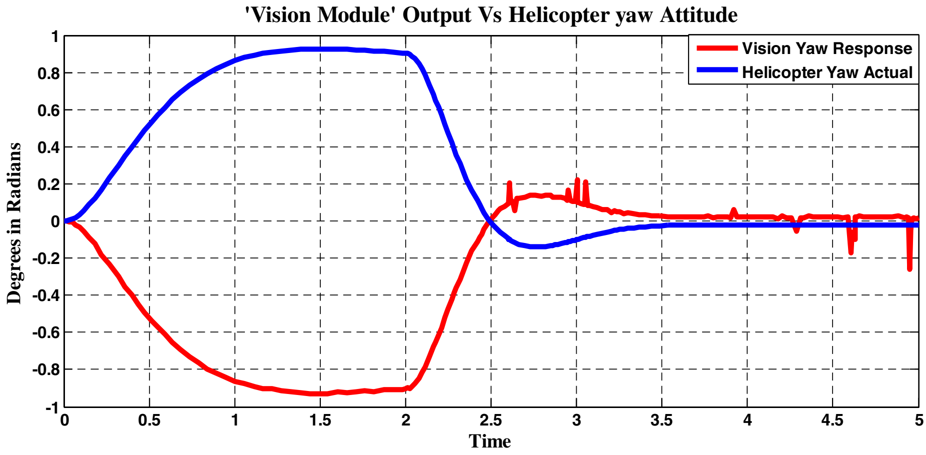

Fig. 12 Longer wait till initiations and initial yaw drive to −30°

either by increasing the number of SIFT features or readily by low pass filtering the estimated yaw angle.

Additional experiments were conducted by adding noise to the position and angle measurement inputs to both the vision and the helicopter controllers. Figure 13 shows the results of a simulation where only the position data was noisified. A Normal Gaussian noise with zero mean and .1 variance was added to the position sensor simulating position error caused by wind gust pushing the helicopter around.

As can be seen position data error does not significantly affect the robustness of the proposed system and the helicopter yaw angle approaches zero. In another experiment the yaw angle data coming into both the helicopter controller and vision

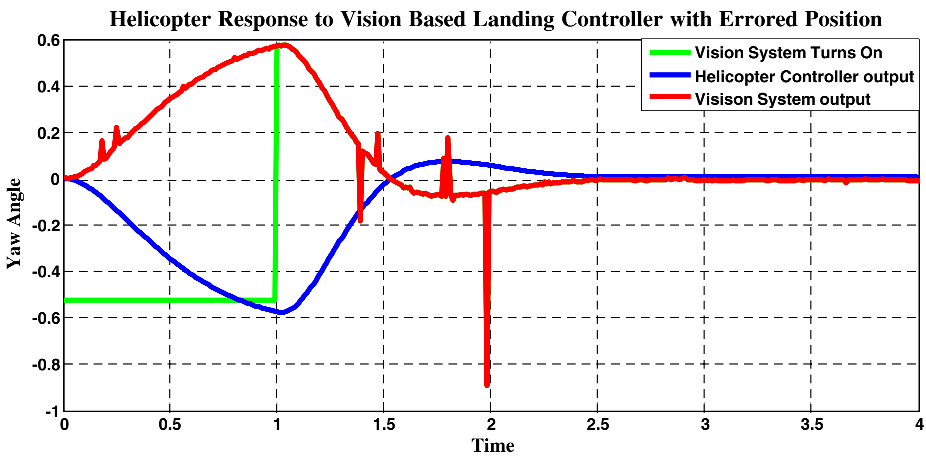

Fig. 13 Error introduced to the position sensor data

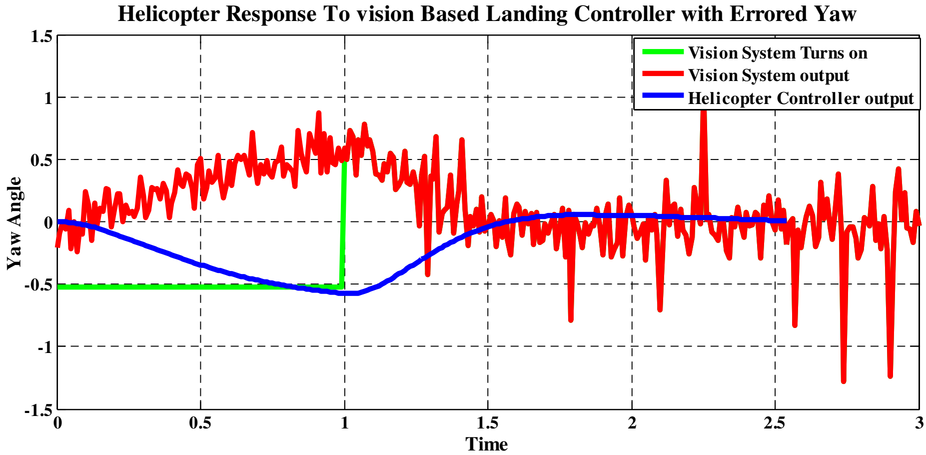

Fig. 14 Error introduced to yaw sensor data

controller was noisified (Normal Gaussian noise with mean zero and variance .01). The results of this experiment are shown in Fig. 14.

Finally the last experiment was a combination of the previous two experiments (adding noise to both the yaw angle and position) with results shown in Fig. 15.

As can be seen the system is more sensitive to angle variations as they have greater effect on the yaw angle estimated by the vision module. Yet the system is still robust enough to drive the helicopter back towards the required yaw measurement of zero degree of the desired reference orientation.

The computational complexity of SIFT is as the computational complexity of the other features [10]. SIFT matching approach can be efficiently implemented by using the gradients and orientations for each level of the pyramid that are used for

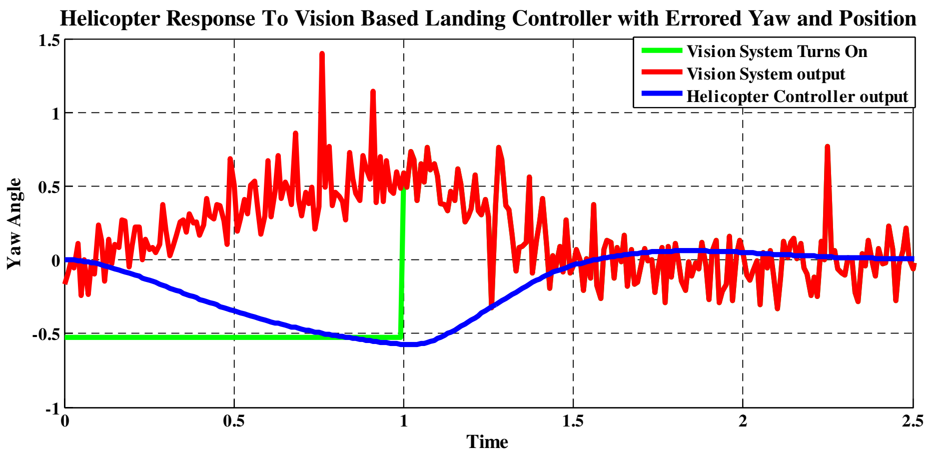

Fig. 15 Error introduced to both yaw and position sensor data

orientation selection. For each keypoint, we use the pixel sampling from the pyramid level at which the key was detected. The SIFT features for the reference image I_r are extracted only once and SIFT features for the I_t image are extracted using VLFeat toolbox (C++ implementation). The run time for finding four corresponding points (out of 80 detected SIFT features) using the RANSAC technique is about 1 second on a regular computer with a 2.5 GHZ Pentium processor (the SIFT matching is carried out using C++ and the RANSAC is implemented using Matlab).

6 Conclusions and Future Work

This paper showed the successful integration of a SIFT based 'vision module' to with any arbitrary autonomous helicopter controller for the purpose of landing in a desired orientation. As this work is in its initial stage, the experiments were performed in Matlab to show the feasibility of the approach. This paper additionally presented a prototype platform being developed for the purpose of docking with autonomous helicopters independent of form factors. In parallel with the landing system, work is being done on developing another prototype of ISLANDS based on the CAD drawing presented in this paper. In the future, this developed 'vision module' will be implemented on a Maxi Joker electric class helicopter to demonstrate a real system integration.

References

1. Shim, H., Koo, T.J., Hoffmann, F., Sastry, S.: A comprehensive study of control design for an autonomous helicopter. In: Proceedings of the 37th IEEE Conference on Decision and Control, Tampa, FL, pp. 3653–3658 (1998)
2. Mullens, K.D., Pacis, E.B., Stancliff, S.B., Burmeister, A.B., Denewiler, T.A. (SAIC), Bruch, M.H., Everett, H.R.: An automated UAV mission system. In: AUVSI Unmanned Systems in International Security 2003 (USIS 03), London, England (2003)
3. Dalamagkidis, K., Ioannou, S., Valavanis, K.P., Stefanakos, E.: A mobile landing platform for miniature vertical take-off and landing vehicles. In: Proceedings of 14th Mediterranean Conference on Control and Automation, Ancona, Italy (2006)
4. Ioannou, S., Dalamagkidis, K., Valavanis, K.P., Stefanakos, E.K., Wiley, P.H.: Improving endurance and range of a UGV with a gimbaled landing platform for small unmanned VTOL vehicles. In: Revised, Submitted to the Journal of Intelligent and Robotic Systems
5. Dickmanns, E.D., Schell, F.R.: Autonomous landing of airplanes using dynamic machine vision. In: Proc. of the IEEE Workshop Applications of Computer Vision, pp. 172–179 (1992)
6. Garcia-Pardo, P.J., Sukhatme, G.S., Montgomery, J.F.: Towards vision-based safe landing for an autonomous helicopter. Robot. Auton. Syst. **38**(1), 19–29 (2001)
7. Caballero, F., Merino, L., Ferruz, J., Ollero, A.: Vision-based odometry and SLAM for medium and high altitude flying UAVs. J. Intell. Robot. Syst. **54**(1–3), 137–161 (2009)
8. Shakernia, O., Vidal, R., Sharp, C., Ma, Y., Sastry, S.: Multiple view motion estimation and control for landing an aerial vehicle. In: Proceedings of the International Conference on Robotics and Automation, ICRA, vol. 3, pp. 2793–2798, IEEE (2002)
9. Vidal, R., Sastry, S., Kim, J., Shakernia, O., Shim, D.: The Berkeley aerial robot project (BEAR). In: Proceeding of the International Conference on Intelligent Robots and Systems, IROS, pp. 1–10. IEEE/RSJ (2002)
10. Lowe, D.G.: object recognition from local scale-invariant features. In: International Conference on Computer Vision, Corfu, Greece, pp. 1150–1157 (1999)
11. Lowe, D.: Distinctive image features from scale-invariant key points. Int. J. Comput. Vis. **60**(2), 91–110 (2004)

12. Brown, M.L., Lowe, D.: Recognizing panoramas. In: Proceedings of the 9th IEEE International Conference on Computer Vision (2003)
13. Niebles, J.C., Wang, H., Li, F.-F.: Unsupervised learning of human action categories using spatial-temporal words. In: Proceedings of the British Machine Vision Conference (BMVC) (2006)
14. Se, S.L., David, G., Little, J.: Vision-based mobile robot localization and mapping using scale-invariant features. In: Proceedings of the IEEE International Conference on Robotics and Automation (2001)
15. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM **24**, 381–395 (1981)