

First Results in Autonomous Landing and Obstacle Avoidance by a Full-Scale Helicopter

Sebastian Scherer, Lyle Chamberlain, and Sanjiv Singh

Abstract— Currently deployed unmanned rotorcraft rely on carefully preplanned missions and operate from prepared sites and thus avoid the need to perceive and react to the environment. Here we consider the problems of finding suitable but previously unmapped landing sites given general coordinates of the goal and planning collision free trajectories in real time to land at the “optimal” site. This requires accurate mapping, fast landing zone evaluation algorithms, and motion planning. We report here on the sensing, perception and motion planning integrated onto a full-scale helicopter that flies completely autonomously. We show results from 8 experiments for landing site selection and 5 runs at obstacles. These experiments have demonstrated the first autonomous full-scale helicopter that successfully selects its own landing sites and avoids obstacles.

I. INTRODUCTION

Unmanned aerial systems (UAS) are now common for high altitude surveillance, weapons designation, or communication relay missions. In these missions fixed-wing aircraft are the most relevant technology since they have the largest endurance. In such missions aircraft takeoff and land at prepared runways which are set up to ensure safe flight up to elevations at which they can fly safely without need for perceiving the environments. More recently, rotorcraft UAS are being suggested as viable solutions for missions such as casualty extraction, and cargo supply.

Before low-altitude missions can be feasible, there must be a reliable way of operating near incompletely mapped terrain in proximity to obstacles such as power lines, buildings, or trees. Landing in rugged unprepared terrain can be catastrophic if the ground crew incorrectly identifies a landing zone as safe, or if GPS errors cause the vehicle to land in a rough or sloped area. Hence, a rotorcraft UAS should have the “common sense” to be safe even if directed along an approach that would nominally cause a collision or if directed to land on rough or sloped terrain.

A motivating scenario that we used is one of autonomous casualty evacuation. The rotorcraft flies to the reported position of the casualty and searches for a safe landing zone in the vicinity. It then routes itself to land at the chosen site, ensuring that there are no collisions with near terrain objects as it descends.

Such a scenario requires two key capabilities: landing zone selection, and obstacle avoidance. Here we report on recent developments that address these needs. Specifically, we report on an incremental model-based landing zone evaluation algorithm, and a method of obstacle avoidance.



Fig. 1: Perception system mounted on the nose of the Boeing Unmanned Little Bird. Our system is capable of detecting obstacles as well as locating safe landing zones during flight.

We report in detail on flight tests that have accomplished the first autonomous landing in previously unmapped terrain by a full-scale helicopter.

II. RELATED WORK

A recent survey article by Kendoul [1] describes significant advances for intelligent rotorcraft. In particular automated landing has been studied for small UAS platforms. Machine vision systems have been popular on small radio-controlled helicopters. Bosch et al. [2] used monocular images from a camera to identify planar areas for landing. Johnson et al. [3] used structure-from-motion to build 3D models of the terrain for guidance and landing. Templeton et al. [4] also demonstrated a landing system based on a similar structure-from-motion concept.

While passive vision does work for many smaller aircraft that can hover very close to terrain, larger full-scale rotorcraft must operate outside the range in which passive stereovision is useful.

LADAR systems have been used as high-resolution solutions to the autonomous landing problem. Johnson et al. [5] did early analysis and simulation that suggested that LADAR would be well suited for the similar problem of landing on Mars. Whalley et al. [6] implemented an autonomous landing system on a 4-meter RMax unmanned helicopter platform. The results of this work demonstrated the promise of compact 3D LADAR as a navigation aid for a UAS. In a similar vein, we have developed a landing analysis system that operated on a full-scale EC-135 helicopter, but only in open-loop manned runs [7]. This paper describes our efforts to implement and test a 3D LADAR-based perception and planning system that accommodates the high speeds, shallow glide slopes on approach, and limited-hover requirements of a large cargo-carrying rotorcraft.

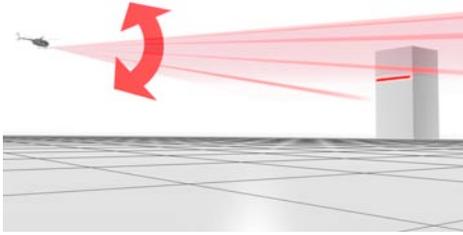


Fig. 2: Obstacle-Detection Scan Configuration. Scanner faces forward and nods up and down to detect obstacles as the aircraft descends to lower altitudes.

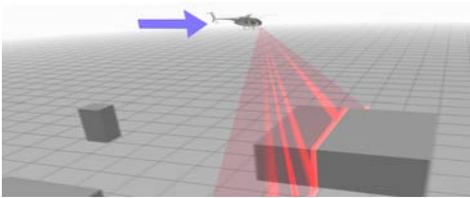


Fig. 3: Landing Zone Survey Scan Configuration. Planar scanning sweeps terrain as the aircraft overflies potential landing zones from a safe altitude.

We also set out to extend our earlier work in obstacle avoidance, previously demonstrated with 3D LADAR on an RMax helicopter [8], to a full-scale platform.

Below we outline our perception and system design, show our approach to obstacle avoidance, and present the results for multiple experiments on the unmanned helicopter.

III. PERCEPTION AND SYSTEM DESIGN

The sensing system for an unmanned platform operating in an unmapped environment must provide two types of information: terrain information for analysis of potential landing zones, and proximity information about obstacles it may encounter en-route. Several sensing modalities are possible. However, the driving requirement in our case was high-resolution measurement of the terrain while flying at a reasonable speed and altitude for a full-scale helicopter. This terrain measurement requirement rules out RADAR and machine vision. RADAR with sufficient angular resolution would require a large antenna, and the depth and the reliability of vision systems is inadequate for this application. Recent advances in scanning LADAR technology have improved the range and bandwidth significantly, opening this modality as a prime candidate for a viable perception system on a full-scale helicopter.

We have designed and built a 3-D scanning LADAR that operates in two modes: Forward-scanning for obstacle detection during low-altitude flight (Fig. 2), and downward scanning for terrain mapping and landing zone search from a higher altitude (Fig. 3). The scanner operates in two axes: a fast 100 Hz, 100-degree transverse scan and a slower “nodding” scan that cycles up and down. The LADAR (Class 1) pulses at 83kHz, providing sufficient point density for a high-resolution terrain map. The sensor has a range of 150 meters with centimeter accuracy, and can process multiple returns for obscurant and vegetation penetration (See Fig. 4).

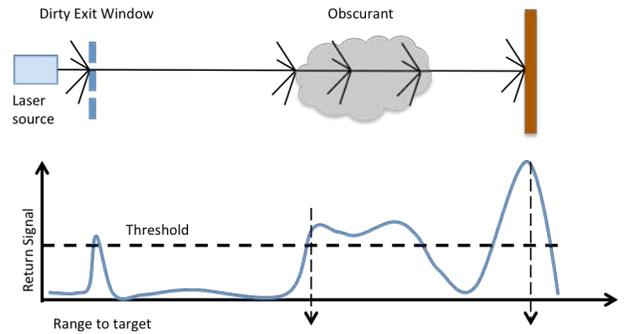


Fig. 4: Example of multi-return processing. A multi-return lidar can process multiple returns to penetrate reflections from obscurants.

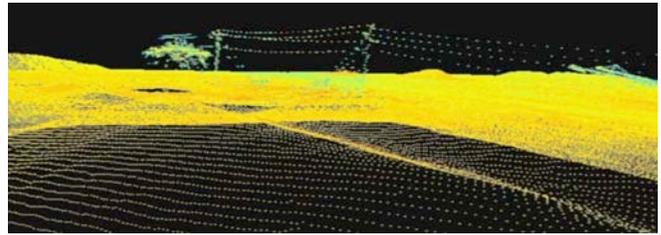


Fig. 5: Example point cloud showing terrain and power lines.

Motion control for actuation of both axes is precisely synchronized to a global timeframe so that for each measurement, the time is known and can be fused with INS data for point registration

The LADAR is rigidly coupled to a RTK GPS/INS system so that the attitude of the sensor head is known at all times despite vibration. Otherwise, the map will become “blurred” due to errors in the measurement of where the instrument was pointed when each point was measured. Each measured point is registered into a global reference frame. The INS uses a ring-laser gyro with <10-deg/hour drift, differential GPS with 1cm accuracy, and an Extended Kalman filter that timestamps the measurements into a timeframe synchronized

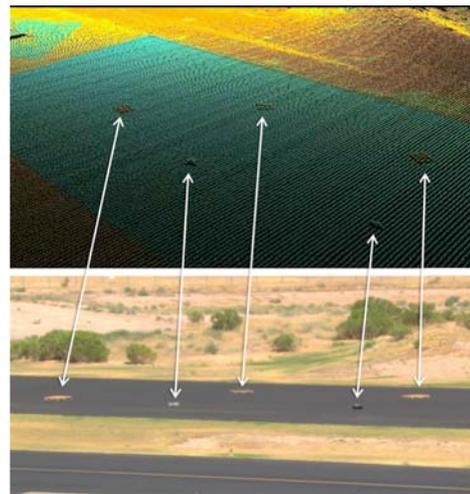


Fig. 6: Obstacle registration near landing area. Detection of small objects (here, forklift pallets and plastic storage boxes) requires high resolution and precise coupling with the inertial measurement system in order to resolve small features from long distances. This map was made at 40 knots from 150 feet AGL.

with the rest of the perception and computing system.

The system is capable of detecting chain link fences, wires (Fig. 5), and 4-inch flat pallets from up to 100 meters away (Fig. 6) during flight on a helicopter. could we say more about the sensor?

Figure 7 shows the overall system design. The LADAR/INS feeds into the perception module of the software, which classifies the lidar point cloud into landing sites. Perception data feeds into the planning software. The planning system then plans a feasible approach path while keeping the aircraft safe. The result is a trajectory that is sent to the autopilot as often as necessary to respond to new information arriving from the perception system. The software is running in real-time on two Intel Core 2 Quad PCs running at 2.2 GHz.

The Boeing Unmanned Little Bird (ULB) aircraft has the capability of waypoint-mode autonomous flight and landing. Our system provides the aircraft with obstacle-free trajectories to follow, and with verified landing zones.

The best landing site is initially unknown. At the beginning of a mission, the aircraft is only given a GPS coordinate that marks the location of a casualty that needs to be picked up. No other information is given a-priori.

Figure 8 shows the steps the landing zone evaluation algorithm follows to find a landing site. The approach is an extension of the algorithm we presented in [7]. First, the aircraft overflies the known position of the casualty in downward-scan survey configuration (see Figure 3) and builds a map of the surrounding terrain in the global coordinate frame. This map is shown in artificial color with reflectance values in Figure 8a. Note that in this example there are several obstacles on what would normally be an ideal flat and level place to land on the tarmac (See Figure 5 for a close up of the obstacles in a similar scan).

The first data-processing step, a “coarse” evaluation,

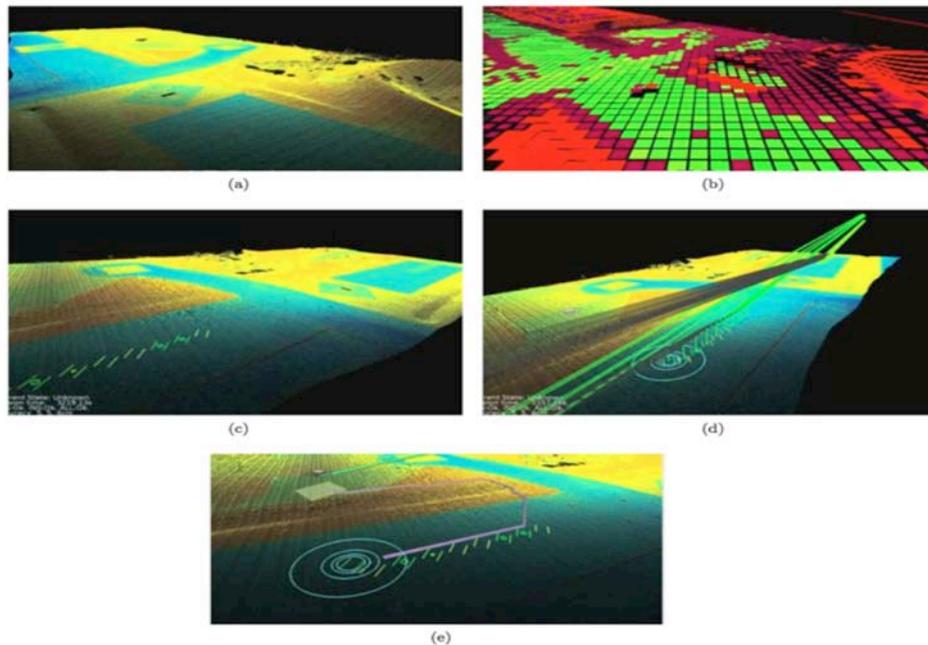


Fig. 8: Landing Zone Evaluation is performed in 5 steps: (a) Overflight survey point cloud, (b) rough terrain evaluation, (c) fine terrain evaluation, (d) approach/abort path validation, and (e) ground path for traversal to the casualty

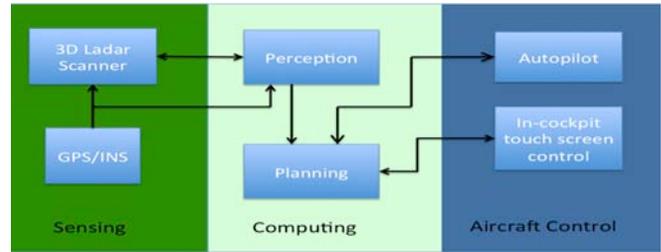


Fig. 7: Hardware Functional Diagram. The 3D scanning lidar registers its measurements to a global reference frame using GPS/INS. The computers both evaluate terrain and obstacles, and also plan safe trajectories to complete the mission. These trajectories are sent to the flight control system autopilot on the rotorcraft.

shown in Figure 8b, immediately removes rough or sloped areas from consideration. The green grid squares (or light, for those viewing in greyscale) are considered smooth enough for a closer look. Most of the terrain not on the tarmac is immediately discounted (red, or dark) from further consideration. Note that the clutter on the runway is also discounted in this step.

With the list of potential landing sites narrowed down to a few options, the system can perform a fine evaluation of each to decide which is best. The fine evaluation places a 3-dimensional virtual model of the helicopter on each grid cell and evaluates skid/ground contact, position of center-of-gravity within the support points (static tipover stability), wind direction, and clearance with adjacent terrain both beneath and adjacent to the rotor and tail. The software evaluates multiple headings at each potential landing spot. In this example, the fine evaluation step found good landing spots shown as parallel skid positions in Figure 8c. The relative suitability of a given position is indicated by the size of the circle rendered between the landing skid markers. Note that these sites are located in the center of the tarmac.

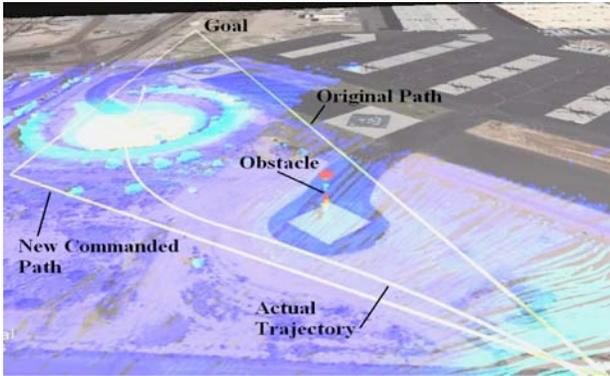


Fig. 9: Data render from an autonomous obstacle avoidance maneuver at 21 knots ground speed. Points measured by the LADAR in-flight are overlaid against a satellite image for this visualization.

The fine evaluation step discarded the sites near the edge of the flight line or near the clutter obstacles because there was less clearance for the rotor/tail compared with the center locations.

With a list of acceptable touch-down points, the system evaluates the glide slope to each one from multiple directions, as well as the corresponding abort trajectory (Fig. 8d). These approach/abort paths are compared against the surveyed point cloud model and evaluated for proximity to obstacles such as buildings or power lines. Also, glide slopes with a headwind are preferred, crosswind is less acceptable, and downwind is rejected.

Simultaneous to the approach/abort step, the system also evaluates the path that a ground medical crew would have to follow to get to the landing zone from the reported casualty position. Figure 8e shows a path from a selected site (marked by concentric rings) to the casualty (diamond). This ensures that the helicopter won't land on the wrong side of a ravine or fence.

The actual landing site is chosen by a multiple-objective optimization function that evaluates the results from the approach/abort and ground path metrics and comes up with a compromise between the two decision criteria. Once the final site is chosen, a path planner computes a landing pattern trajectory that will start its decent at the beginning of the approach path corresponding to the winning landing zone.

When the vehicle starts its decent at the beginning of the final glide to the landing zone, it switches the perception mode to forward-facing for obstacle detection (Figure 2). This helps ensure that new or previously occluded objects in



Fig. 10: The system autonomously guides the Unmanned Little Bird around a previously unmapped obstacle.

the glide path are detected and avoided by aborting the landing.

IV. OBSTACLE AVOIDANCE

We extended our earlier work on a Yamaha RMax helicopter for obstacle avoidance in unmapped environments [8]. Obstacles are registered separately from the landing map in an evidence-grid representation of the world [9]. As the helicopter flies, its planned trajectory is compared to a C-Space expansion of the evidence grid [10]. If the trajectory intersects an obstacle, then a new path is planned. In simulations, the path planner uses a hybrid trajectory optimization/path-planning algorithm utilizing a forward model of the helicopter dynamics to create a viable avoidance path. In actual flight tests, the limited waypoint-style interface to the flight controller necessitated a different planning method. We used the LIDT and D* lite [10] algorithm to create a path around obstacles, and then forward-simulated the resulting waypoint paths to ensure that they did not intersect the obstacle.

Figure 9 and 10 show an example obstacle avoidance test and illustrates the resolution of commands we were able to send the Unmanned Little Bird controller. The original path was defined using a single waypoint, and the projected “original path” is the forward-simulated prediction of where the vehicle would fly. The new obstacle avoidance path is given using just three intermediate waypoints generated and verified using the D* lite and forward simulation approach outlined above. Even though the vehicle does not closely follow the connecting lines, the behavior is similar to what the path planner predicts using its forward simulation. It can be seen that operation in more cluttered airspace with slalom-like maneuvers will require a tighter integration between the flight controller and the path planner (as simulated in ideal system design earlier in the project).

V. EXPERIMENTAL RESULTS

Tests were performed on the flight line of the Boeing test facility in Mesa Arizona.

A. Landing

After system integration and testing, a total of 8 landing missions were demonstrated with varying ground clutter and approach obstructions. Due to the heavy air traffic and flyover restrictions in the area, we were constrained a survey flyover at 150 feet AGL and 40 knots.

Figure 11 shows a typical autonomous landing mission. After autonomous takeoff and climb out, the aircraft approaches the flight line from the southwest (see Figure 3). After deciding on a landing zone, the system sets up an approach from the north and lands near the casualty. Note, that the computation time for finding the landing zones is low enough that a landing zone is found almost immediately after finishing the survey flyover.

A comparison of the landing approaches in Figure 11 and 12 shows the effect that changing the priorities of the optimization process has on the decisions the system makes. In Figure 11, the optimization process was not heavily

penalized for approaches that descend through unmapped terrain. This is sufficient, since the system switches to forward-looking scanning (Figure 2) when flying at lower altitudes and could respond to unexpected obstacles. In this experiment, the system avoided the mining equipment in the quarry (upper right of scan) and a manlift placed close to the approach; consequently, it chose to land from the North.



Fig. 11: A typical mission including survey overhead, descent, and touch-down. Here, no penalty is levied for flying the final approach through unsurveyed terrain, and a descent from the north is chosen. The forward-facing scanner is responsible for detecting any obstacles during the final approach in such cases.

In contrast, Fig. 12 shows a mission where the weights for the landing zone evaluation were set to penalize unknown areas. In this case the helicopter flew a buttonhook turn to land at the selected site. Here the rotorcraft is forced to get closer to obstacles but it only flew through areas that it already had mapped.



Fig. 12: A mission where the system is penalized for descending through unsurveyed terrain. It chooses to perform a "button-hook" turn to final and descend through the area just mapped.

Figure 13 shows a close-up of two missions flown with no penalty for descending outside the surveyed terrain. In Figure 13a, the system chooses to land in a large clear area on the flight line. It is a bit distant from the casualty (the red "+"), but the multi-objective optimization chooses to land there because it is a better landing site overall. Next the best landing areas were cluttered with low-profile obstacles as shown in Figure 13b. (The obstacle locations are annotated with the red X's because they are difficult to see in a small image. See Figure 6 for a close-up view of the type of clutter used and the resulting point cloud). In this case, the system decides to land just north of the casualty at a constrained landing site. The red circle around the casualty marker "+" denotes the safety radius around the casualty position. The

vehicle has to use a north approach to land there due to the no-fly restrictions and the manlift obstacle to the northeast (see Figure 11).

Figure 14 shows a pair of tests where the optimization was penalized for any flight through unmapped areas. Like the previous example, the system decides to land in the large open area shown in Fig. 14a when there is no clutter on

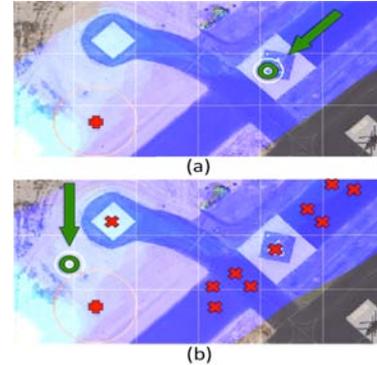


Fig. 13: Example of chosen landing sites during a test series. (a) shows the landing decision and approach direction (green arrow and circle) for a casualty (red cross) located on sloped terrain but near large flat areas. (b) landing site chosen when these large areas are covered with debris (denoted with red x's. See Figure 5 for close-up). The system lands from the north to avoid a structure that would occlude the glide slope from the North East.

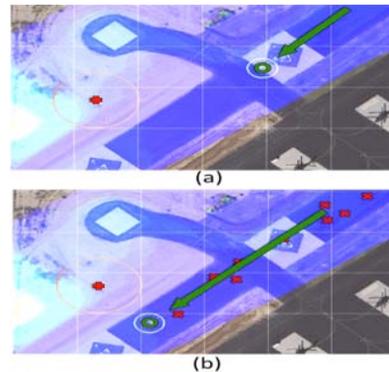


Fig. 14: Example of chosen landing sites when the system is constrained to approach only through measured terrain. Similar to the tests in Figure 12, the system initially chooses the large open area near the helipad. However, when faced with clutter, it lands to the south of the casualty because this is the only clear area that has an obstacle-free glide slope available.

the flight line. However, with clutter placed similar to the previous example, the system decides to land further down the flight line using the approach from the northeast (see Fig. 14b). In this case, the landing zone to the north of the casualty was discounted because the manlift obstacle would block the northeast approach. For a video visit of this landing see: <http://youtu.be/BJ3RhXjucsE>

B. Obstacle Avoidance

The obstacle avoidance tests demonstrated the ability of the system to detect and react to obstacles in real-time while flying at low altitude with the sensor in forward-scan mode. We conducted a total of five obstacle-avoidance runs against a manlift obstacle. In these tests, the aircraft flew towards a GPS waypoint goal at approximately 20-meters altitude

above ground (AGL). Unbeknownst to the system, we extended a 20-meter manlift crane directly into the flight path needed to achieve this waypoint. The system needs to detect the obstacle, recognize the threat, and change course to avoid a collision and still arrive at the goal point.

Figure 9 shows a snapshot of the data during one of these runs. Here, the manlift obstacle has been detected, a new

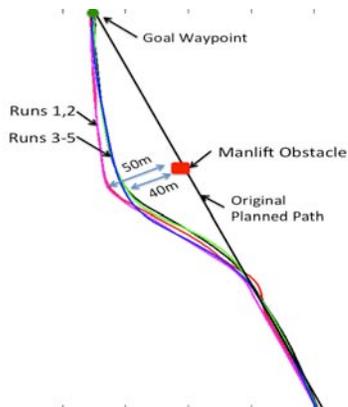


Fig. 15: Five obstacle-avoidance runs against an unmapped obstacle. Runs 1 and 2 were below 10 knots groundspeed, and runs 3-5 were at 21 knots. (Top-down view)

route has been planned, and the helicopter has successfully avoided the obstacle and is arriving at its waypoint destination.

The first two runs at 25 meters AGL were flown below 10 knots ground speed to verify successful avoidance. Figure 14 shows the flight traces from these runs. At low speed, the vehicle avoids the obstacle by about 50 meters (45 meters from edge of rotors). The third run was done at 21 knots ground speed at 25 meters AGL. The system responded in time and avoided the manlift by 40 meters (35 meters from edge of rotors). Since the margin was sufficient we lowered the altitude to just below the manlift height, approximately 18 meters AGL. Test runs 4 and 5 were run at 21 knots ground speed. Again the system responded appropriately and avoided with 35 meters of clearance to the rotor blades (Fig. 14)..

VI. CONCLUSIONS

We have shown a perception and planning system capable of guiding a full-scale helicopter in previously unmapped environments. The system successfully completed 8 landing missions with varied obstacle configurations. We also demonstrated 5 obstacle avoidance runs at speeds up to 21 knots against a tower-style obstacle.

Our results show that high-accuracy inertial measurement and precise pointing registration allow a LADAR system to create maps at the higher speeds and longer ranges needed by full-scale aircraft. While a 40 knot speed can be supported by the current sensor range, longer-range sensing will be needed for faster flight during take-off and landing.

While the full system was demonstrated successfully, the airspace constraints of the test environment precluded thorough methodical testing of all of the capabilities of the system. For example, aborted landings due to changing

obstacle configurations were not tested. The effects of other terrain and casualty locations should also be investigated.

Obstacle avoidance was shown to work well, but the limitations of the interface to the autopilot on the aircraft limited the types of avoidance maneuvers to simple swerving-style trajectories. A UAS design that will incorporate a perception system should have a flight guidance system designed to accept a continuously updated trajectory from the planning/perception unit.

ACKNOWLEDGMENTS

We would like to thank TATRC (Telemedicine and Advance Research Center) from the US Army Research and Materiel Command and Piasecki Aircraft Corp. The Boeing Company provided the Unmanned Little Bird and the test facilities in Mesa, AZ.

References

- [1] Farid Kendoul, "A Survey of Advances in Guidance, Navigation and Control of Unmanned Rotorcraft," *Journal of Field Robotics*, 2012, DOI: 10.1002/rob.20414
- [2] S Bosch, S Lacroix, and F Caballero, "Autonomous Detection of Safe Landing Areas for an UAV from Monocular Images," IROS 2006
- [3] A Johnson, J Montgomery, and L Matthies, "Vision Guided Landing of an Autonomous Helicopter in Hazardous Terrain," ICRA 2005
- [4] T Templeton, D Shim, C Geyer, and S Sastry, "Autonomous Vision-based Landing and Terrain Mapping Using an MPC-controlled Unmanned Rotorcraft", ICRA 2007
- [5] AE Johnson, AR Klumpp, JB Collier, and AA Wolf, "Lidar-based hazard avoidance for safe landing on Mars," *AIAA Journal of Guidance Control and Dynamics*, 25(6):1091-1099, 1/2002.
- [6] M Whalley, G Schulein, and C Theodore, "Design and Flight Test Results for a Hemispherical LADAR Developed to Support Unmanned Rotorcraft Urban Operations Research," AHS Forum, 2008.
- [7] Sebastian Scherer, Lyle Chamberlain, and Sanjiv Singh, "Online Assessment of Landing Sites," AIAA Infotech, 2010.
- [8] Sebastian Scherer, Sanjiv Singh, L Chamberlain, and M Elgersma, "Flying Fast and Low Among Obstacles: Methodology and Experiments," *The International Journal of Robotics Research*, 27(5):549-574, May 2008.
- [9] Martin C Martin and Hans Moravec. (1996, Mar) Robot Evidence Grids, Tech Report The Robotics Institute, 1996
- [10] S Scherer, D Ferguson, and S Singh, "Efficient C-Space and Cost Function Updates in 3D for Unmanned Aerial Vehicles," ICRA , 2009.
- [11] Peter Tsenkov et al., "A System for 3D Autonomous Rotorcraft Navigation in Urban Environments," *AIAA 2008*, p. 23, Sep 2008.
- [12] Shane Tierney and Jack W Langelaan, "Autorotation Path Planning Using Backwards Reachable Set and Optimal Control," *AHS Forum*, Apr 2010.
- [13] J Sprinkle, J Eklund, and S Sastry, "Deciding to land a UAV safely in real time," ACC 2005
- [14] D Barber, S Griffiths, T McLain, and Randal W Beard, "Autonomous Landing of Miniature Aerial Vehicles," AIAA Infotech, 2005.
- [15] M Whalley, M Freed, R Harris, and M Takahashi, "Design, Integration, and Flight Test Results for an Autonomous Surveillance Helicopter," Proc. AHS Int. Spec. Meet. on Unmanned Rotorcraft 2005.