Terrain Modeling and Following Using a Compliant Manipulator for Humanitarian Demining Applications

Marc Freese[1], Surya P. N. Singh[2], William Singhose[3],
Edwardo F. Fukushima[1], Shigeo Hirose[1]

Abstract Operations with flexible, compliant manipulators over large workspaces relative to the manipulator are complicated by sensor noise, vibration, and measurement bias. These difficulties are compounded by unstructured environments, such as those encountered in humanitarian demining. By taking advantage of the static structure of the terrain and the manipulator's fundamental mechanical characteristics, a series of adaptive corrections and filters refine noisy topographical measurements. These filters along with a shaped actuation scheme can generate smooth and well-controlled trajectories. Experimental testing was performed on a field robot with a compliant, 3 *m* long-reach hybrid manipulator and a stereo vision system for terrain sensing. The proposed method provides a vertical tracking precision of ±5 *mm* on a variety of ground clearings, with scanning speeds of up to 0.5 *m/s*. As such, it can agilely move the attached sensor(s) through precise scanning trajectories that are very close to the ground. This method improves overall detection and generation of precise maps of suspected mine locations**.**

1. Introduction

While robust manipulation in difficult field conditions is still in its infancy, environmental modeling using computer vision has progressed with several applications, including autonomous Martian mapping [10]. With regards to manipulation, variation and noise are routinely minimized by stiffening the structure or constraining the operation [1, 2]. In humanitarian demining, the increased sensitivity needed by the metal detector precludes the addition of proximal metal content and requires minimal mean and variance in the air gap to the ground. In demining, a myriad of sensing technologies have been proposed [7]. However, relatively little attention has been directed towards the field manipulation requirements for automating the dangerous and tedious ground scanning task.

[1] Department of Mechanical and Aerospace Eng., Tokyo Institute of Technology, e-mail:{marc@sms, fukusima@mes, hirose@mes}.titech.ac.jp

[2] Australian Centre for Field Robotics, University of Sydney, e-mail: spns@usyd.edu.au

[3] School of Mechanical Eng., Georgia Institute of Technology, e-mail: singhose@gatech.edu

Fig. 1 Photograph of Gryphon in a test humanitarian demining field.

An autonomous mobile robot named Gryphon [6] has been developed to assist the mine detection process. As shown in Figure 1, the robot is based on an all terrain vehicle (ATV) [4] to which a custom hybrid robotic manipulator [5] is added. This lightweight and counter-balanced 3-DOF arm is made from glass-fiber reinforced plastic (GFRP) which is compatible with sensitive metal detection (MD) operation that requires minimal metal near the sensor [3]. As shown in Figure 2, this design has a tip flexure of up to 5 *cm* (for a 2 *m* link) due to gravitational forces only. While the structure could have been stiffened, the compliance also provides some safety in the event of a collision. Secondary motion from the ATV suspension is reduced, but not eliminated, through counterbalancing the manipulator. Hence, there is uncertainty in sensor location with respect to the ground.

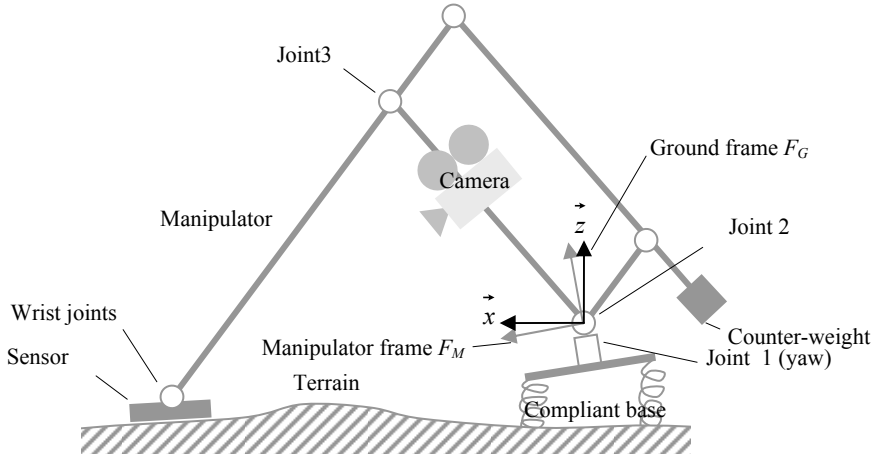

Fig. 2 Schematic of Gryphon and its principal frames of reference.

As illustrated by the control architecture in Figure 3, Gryphon operates by driving close to a region of interest, then while the ATV is stationary, generates a **stereo map**. As detailed in Section 2, these are iteratively refined to construct a geo-

metric **terrain model**. By iteratively operating using a local model, absolute rectification is not required because later processing stages account for aberrations through command shaping [6, 11]. This approach adds robustness without the need to identify the origin of imprecision. However, its use of linear models means its highest accuracy is near regions used to perform the system-level calibration where errors are small, hence defining the *zone of effectiveness*. From this terrain model a desired **path** for the manipulator endpoint is generated and corrected for the height errors and the travel of a detector body (as opposed to end-effector frame), as detailed in Section 3. The final path is close to the ground while maintaining the best possible orientations for the detector. Experimental tests of the manipulator and the control architecture are presented in Section 4.

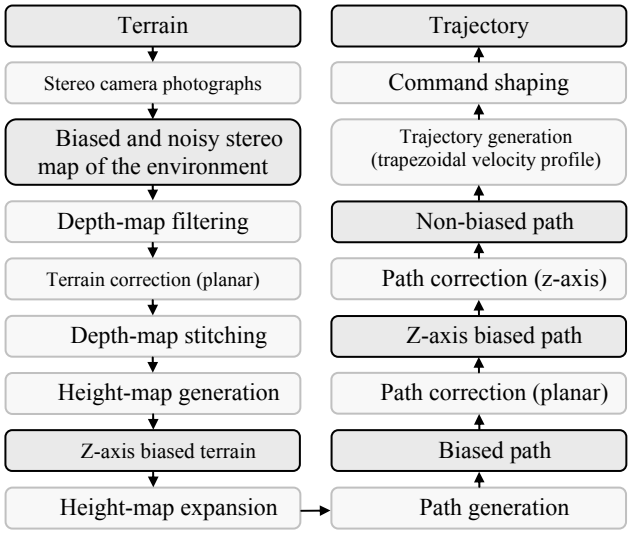

Fig. 3 Terrain scanning motion generation architecture. (Darker boxes indicate process outputs).

2. Terrain Modeling

In order to plan operations that map the location of suspected mines found by the sensor, it is necessary to form a geometric model of the terrain.

The terrain modeling begins by sensing the environment using stereo vision. A *BumbleBee* (model BB-HICOL-60, Point Grey Research) stereo vision camera (location shown in Figure 2), acquires several depth maps from different manipulator configurations to cover the region of interest. Compared to laser mapping, this generates maps quickly, but more noise, especially in regions lacking texture. Map accuracy depends on object range and camera aberrations. Acquired raw depth maps (or disparity maps) are checked for consistency by subdividing them into *coherent* patches. Patches that fail to comply with certain criteria (e.g. do not represent possible terrain locations) are then discarded.

2.1 Conditional Planar Filtering

Due to limitations on camera resolution and calibration, a raw model generated from the camera data is significantly degraded by noise. Median filtering techniques are primarily effective against the shot noise, but do not remove high-frequency components. Simply smoothing (i.e., spatial low-pass filtering) is insufficient as this results in a degradation of features, especially at obstacles boundaries, which could lead to a collision between the end-effector and the terrain.

Thus, an adaptive filter based on the average region planarity is used to adjust filter kernel sizes for both Gaussian smoothing and median filters (i.e., the *conditional planar filter [CPF]*). The planarity of the selected region is determined by calculating the mean deviation between each point in the region and the region's corresponding (least-squares) best-fitting plane. If the deviation is small, the region is assumed to be planar, hence giving a conservative, but rapidly computed terrain classification. Based on this, the strength of the Gaussian and median kernels are varied depending on the deviation magnitude. Applying this to the perceived data yields a less noisy, but still potentially offset, or biased, map.

2.2 Height Map Generation

To simplify and speed-up terrain data processing, the depth maps are transformed to a common height map function. This map represents the terrain as a series of height (or z-coordinate) values at point locations specified by a uniform mesh in the ground plane of the ground frame. This is done, for instance, via Delauney triangulation methods with increased precision obtained through spatial weighted averages of the sensed data. The obtained height map offers the advantage of two dimensionally indexed queries on the terrain model and facilitates, by linear interpolation, a mechanism to fill holes and patches that may have arisen due to occlusions or lack of texture.

2.3 Height Map Expansion

At this stage of the process, the height map is a good approximation of the underlying topographical information; however, it is often desired to perform the scanning at a constant distance from the ground (i.e., a scanning gap). For that purpose we expand the height map via an envelope expansion. That is, we solve for a new surface, f_{env}, whose distance from f_{terr} is given by the scanning gap. This is illustrated in Figure 4. Note that the simple approach of shifting the height map vertically would be a bad approximation for highly curved or inclined terrain.

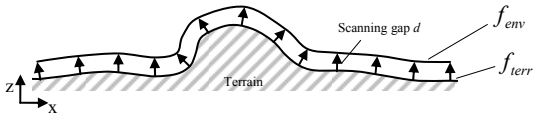

Fig. 4 Terrain surface expansion.

Considering the continuous case, the desired surface is easily obtained with the following parametric equations:

$$x_{env} = x_{terr} - \lambda \cdot \partial f_{terr}(x_{terr}, y_{terr}) / \partial x \tag{1}$$

$$y_{env} = y_{terr} - \lambda \cdot \partial f_{terr}(x_{terr}, y_{terr}) / \partial y \tag{2}$$

$$z_{env} = f_{terr}(x_{terr}, y_{terr}) + \lambda = f_{env}(x_{env}, y_{env}) \tag{3}$$

$$\lambda = d \cdot \left(\left(\partial f_{terr}(x_{terr}, y_{terr}) / \partial x \right)^2 + \left(\partial f_{terr}(x_{terr}, y_{terr}) / \partial y \right)^2 + 1 \right)^{\frac{-1}{2}} \tag{4}$$

3. Path Generation

The path generation for the manipulator takes as inputs the expanded terrain model from the previous section and the manipulator configuration, taking into account joint limitations of the wrist mechanism.

3.1 Scanning Scheme

Two different linear incrementing scanning schemes are available: in joint space (giving a circular profile since the base joint is revolute) or in workspace (giving a rectangular profile). The trajectories are smoothly combined at their extremities to reduce unnecessary slowdowns during direction changes. Performing this in the joint-space of the robot simplifies joint coordination by reducing simultaneous motions and velocity variation, and thus is dynamically more stable. This reduces power consumption and arm vibration and improves detection performance by reducing sensor location uncertainty.

3.2 Terrain Sampling

Either of the above described scanning schemes produce a planar path, where only the *x*-*y* coordinate components are defined. The *z* coordinate is obtained by sampling the corresponding position on the expanded terrain. Orientation of the detector at this position is calculated using the normal vector to the expanded terrain using the following relationships:

$$z_{path} = f_{env}(x_{path}, y_{path}) \tag{5}$$

$$\vec{n}_{path} = \begin{bmatrix} -\gamma \cdot \partial f_{env}(x_{path}, y_{path})/\partial x & -\gamma \cdot \partial f_{env}(x_{path}, y_{path})/\partial y & \gamma \end{bmatrix}^T \tag{6}$$

$$\gamma = \left(\left(\partial f_{env}(x_{path}, y_{path})/\partial x \right)^2 + \left(\partial f_{env}(x_{path}, y_{path})/\partial y \right)^2 + 1 \right)^{\frac{-1}{2}} \tag{7}$$

3.3 Advanced Terrain Following (ATF)

The planned path has focused on guarantying that the center of the sensor follows the expanded terrain. With the consequence that the rest of the sensor body is still free to enter the scanning gap or even to collide with the terrain at positions where the curvature is concave, thus, it is important to consider the sensor as an extended body attached to the manipulator end point. This is performed by modeling the sensor via a series of control points along its lower surface; where, the number and position of points has been chosen to best approximate the contour and surface.

For each sensor position on the trajectory, all control points are tested for possible collisions with the expanded terrain. For each point under the grid, a small correction rotation (e.g., ½°) of the sensor is performed in the order of vertical deviation as illustrated in Figure 5a. This is iteratively repeated until an equilibrium orientation is reached (cf. Figure 5d). The sensor orientation is then corrected to respect configuration-space constraints by limiting wrist joint values to within allowable ranges and the sensor's *z*-coordinate is modified so as to have no control point in the scanning gap (cf. Figure 5e). Such an approach can lead to large gaps between the detector center and the terrain at concave positions, but it guarantees a constant minimum distance between the sensor as a whole and the terrain.

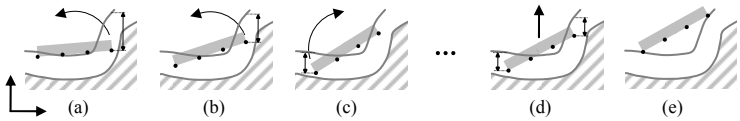

Fig. 5. Path correction scheme with the dots indicating control points.

A reactionary approach, by comparison, would simply lift the sensor without changing its orientation and thus is not sufficient as it would lead to non-ideal configurations with large and changing air gaps and potentially excessive motion.

3.4 Partial Path Correction

While perception inaccuracies were corrected for the horizontal plane in the *Terrain Modeling* Section, the model only partially accounted for mechanical inaccuracies. The type of mechanical inaccuracies corrected for are typically those arising from mechanical compliances (in the links and base) and calibration errors.

Within the limitations of the *zone of effectiveness*, the path's x and y components are modified appropriately. This second part of the system-level calibration consists of applying a radial offset and scaling:

$$\rho' = \rho \cdot scaling_{\rho} + offset_{\rho} \tag{8}$$

and an angular offset:

$$\varphi' = \varphi + offset_{\varphi} \tag{9}$$

The above correction parameters are obtained by measuring the discrepancy between the real and computer-model manipulator tip position. Orientation of the detector is kept unchanged. The path becomes:

$$\begin{bmatrix} x_{path} & y_{path} & z_{path} \end{bmatrix}^T = \begin{bmatrix} \rho' \cdot \cos\varphi' & \rho' \cdot \sin\varphi' & f_{em}(\rho \cdot \cos\varphi, \rho \cdot \sin\varphi) \end{bmatrix}^T \tag{10}$$

3.5 Final Path Correction

The last modification of the path, and the final step of the system-level calibration, is the correction along the z-axis. Each x-y position is assigned an individual vertical correction factor that is obtained by linear interpolation between values of an Overall Calibration Matrix (OCM):

$$z'_{path} = z_{path} + f_{ZCorrection}(x_{path}, y_{path}) \tag{11}$$

The OCM is obtained by mapping a relatively flat terrain, generating a non-expanded height map, and then manually driving the sensor to touch the terrain at various spots within the workspace. At each spot, the necessary correction factor is given by the deviation of the computer model of the sensor from the terrain contact. Spots are then used to generate a Delaunay triangulated surface whose height at a given position gives the amount of correction needed. The OCM is directly extracted from that surface at regular grid intervals along the x-y plane in the manipulator frame.

4. Experiments and Results

Several experiments were conducted with Gryphon to assess performance of the terrain following method described in the previous sections. Precise quantification was performed by replacing the detector payload (i.e., metal detectors and/or ground penetrating radar) with a laser rangefinder (SICK DME 2000, SICK) to track the scanning gap at the region of maximum detector sensitivity. The effects of command shaping, the implementation of which on Gryphon is detailed in [6], were analyzed by attaching a 3-axis accelerometer to the tip of the manipulator.

4.1 Filter Performance

Other operating conditions, where possible, were adjusted to match those of a typical payload (e.g., weight was added to the laser rangefinder). A sandy terrain under natural light conditions (i.e., a fair weather day) was used with a circular scanning scheme with a scanning rate of 100 *mm/s* to evaluate the filter performance, and the effects of the advanced terrain following.

This set of tests used a 10 cm height map expansion and no ATF. The map filter was varied between three algorithms: (a) *minimal filtering* (a control case based on a raw depth map with median filtering), (b) *Gaussian Smoothing,* and (c) *Conditional Planar Filter (CPF)*.

Filter performance was tested on two different terrain profiles: (i) relatively flat terrain and (ii) the same terrain with a rather challenging hump (or obstacle). In both cases, the scanning procedure comprised several passes.

The obstacle, shown in Figure 6, represents extreme slopes and contours for expected demining conditions. The pass selected for comparison was the most challenging and includes a 70° curvature.

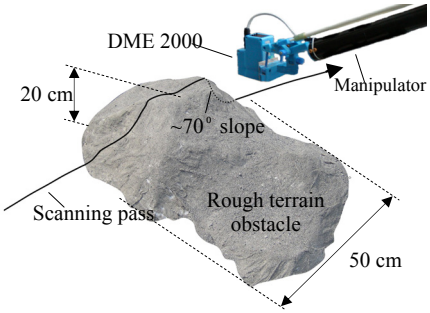

Fig. 6 Scanning pass used during experiments.

The advantage of the CPF over Gaussian smoothing is visible in Figure 7. The Gaussian degrades terrain features, which results in a loss/gain of obstacle height.

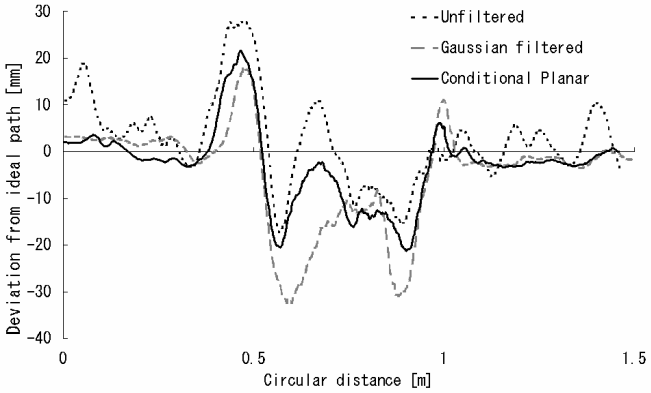

Fig. 7 Scanning pass over obstacle, filter comparison.

4.2 Effects of Command Shaping

The experiments also validate the use of robust control to address sensing and planning aberrations. In the case of Gryphon, an unshaped trajectory is first calculated by generating a trapezoidal velocity profile in joint space. The trajectory is then modified by a Zero Vibration and Derivative (ZVD) shaper (cf. ref. [11]). The sensor loading is large enough to that arm compliance is visible. In particular, measured residual vibration at the tip of the arm is approximately ±20 *mm*. This makes it difficult for a feedback control system relying on the joint encoders to adequately control the endpoint vibration. With command shaping, the residual vibration is approximately ±2 *mm*.

Smooth trajectories are of importance for best metal detector sensitivity, especially for straight-line rectangular motion as this requires coordinated motion of the arm joints. As shown in Figure 8, command shaping is significantly smoother.

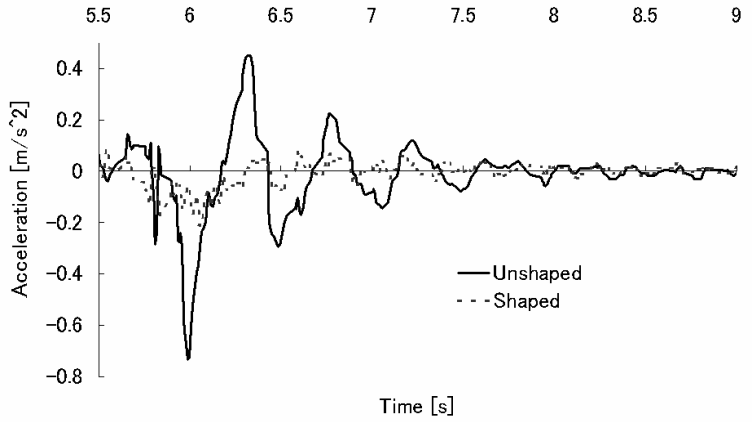

Fig. 8 Endpoint acceleration at the end-effector for a typical scanning pass for a straight line scanning pass of 2.4 *m* with a maximum speed and acceleration of 0.8 *m/s* and 2.66 m/s^2.

4.3 In-Field Testing

Gryphon is an integrated, weather-proof system built to be robust against dust, humidity and rain, and resistant to extended temperature ranges. It has been field tested for 95 days on flat ground, bumpy terrain, and slopes. This includes operations on various test mine fields, including tests conducted under the supervision of the Japan Science and Technology Agency (JST) in Japan in early 2005 [8] and early 2006 in Croatia [9]. Two Gryphon robots (each with a different detector configuration) also took part in extensive trials in Cambodia (in cooperation with by the Cambodian Mine Action Center). Operated by Cambodian deminers, the Gryphon machines performed tests for more than 150 hours of semi-autonomous operation.

5. Conclusion

Field operations of long-reach manipulators are complicated by noise, measurement bias and vibration that lead to end-effector positional uncertainty. This drastically reduces the number of field applications that can be performed accurately and safely. The architecture described in this paper combines a novel filtering method, a decoupled system-level calibration procedure and a vibration reduction technique to yield an effective framework for obstacle identification, trajectory planning and generation. Experimental testing on Gryphon shows considerable deviation reduction when applying the framework. Combined with an advanced terrain following technique, it effectively avoids collision with the terrain for a successful scanning operation. The methods demonstrated improved automated mine detection performance and tracking on the Gryphon robot.

Acknowledgments We warmly appreciate the cooperation with Cambodian deminers and Cambodian Mine Action Center staff. This work is supported by the Japan Science and Technology Agency and the Grant-in Aid for the 21st Century COE Program by the Ministry of Education, Culture, Sports, Science and Technology. The second author also acknowledges the support of the Rio Tinto Centre for Mine Automation and the ARC Centres of Excellence Programme funded by the Australia Research Council and New South Wales State Government.

References

1. W. J. Book, "Modeling, Design, and Control of Flexible Manipulator Arms: A Tutorial Review", *Proc. of the 29th Conf. on Decision and Control*, Honolulu, December 1990.
2. R. R. Craig and M. Bampton, "Coupling of Substructures for Dynamic Analysis", *Journal of AIAA*, vol. 6, No. 7, pp. 1313-1319, 1968.
3. CEIA Corporation. CEIA Metal Detector MIL-D1, Operator Manual.
4. P. Debenest, E. F. Fukushima, Y. Tojo, and S. Hirose, "A New Approach to Humanitarian Demining, Part 1: Mobile Platform for Operation on Unstructured Terrain", *Journal of Autonomous Robots*, vol. 18, pp. 303-321, 2005.
5. P. Debenest, E. F. Fukushima, Y. Tojo, and S. Hirose, "A New Approach to Humanitarian Demining, Part 2: Development and Analysis of Pantographic Manipulator", *Journal of Autonomous Robots*, vol. 18, pp. 323-336, 2005.
6. M. Freese *et al.*, "Robotics Assisted Demining with Gryphon", *Advanced Robotics*, vol. 21, no. 15, pp. 1763-1786, 2007.
7. M. Ghaffari, D. Manthena, A. Ghaffari, and E. L. Hall, "Mines and human casualties, a robotics approach toward mine clearing," in *SPIE Intelligent Robots and Computer Vision XXI: Algorithms, Techniques, and Active Vision*, vol. 5608, 2004.
8. J. Ishikawa, M. Kiyota, and K. Furuta, "Evaluation of Test Results of GPR-based Anti-personnel Landmine Detection Systems Mounted on Robotic Vehicles", *Proc. of the IARP Int. Workshop on Robotics and Mechanical Assistance in Humanitarian Demining,* Tokyo, June 2005.
9. J. Ishikawa, M. Kiyota, N. Pavkovic, and K. Furuta, "Test and Evaluation of Japanese GPR-EMI Dual Sensor Systems at Benkovac Test Site in Croatia", *technical report JST-TECH-MINE06-002*, Japan Science and Technology Agency, 2006.
10. C. F. Olson, L. H. Matthies, J. R. Wright, R. Li, and K. Di, "Visual Terrain Mapping for Mars Exploration", *Proc. IEEE Aerospace Conf.*, vol. 2, pp. 762-771, March 2004.
11. N. C. Singer and W. P. Seering, "Preshaping command inputs to reduce system vibration," *J. of Dyn. Syst., Meas. Contr.*, vol. 112, pp. 76-82, 1990.