

Outdoor Downward-facing Optical Flow Odometry with Commodity Sensors

Michael Dille, Ben Grocholsky, and Sanjiv Singh

Abstract Positioning is a key task in most field robotics applications but can be very challenging in GPS-denied or high-slip environments. A common tactic in such cases is to position visually, and we present a visual odometry implementation with the unusual reliance on optical mouse sensors to report vehicle velocity. Using multiple kilometers of data from a lunar rover prototype, we demonstrate that, in conjunction with a moderate-grade inertial measurement unit, such a sensor can provide an integrated pose stream that is at times more accurate than that achievable by wheel odometry and visibly more desirable for perception purposes than that provided by a high-end GPS-INS system. A discussion of the sensor's limitations and several drift mitigating strategies attempted are presented.

1 Introduction

Accurate knowledge of position is critical to successful completion of field robotics tasks. In known or highly structured environments, localization relative to a pre-determined or progressively-refined map is typically performed using sensors appropriate for registering map features to observations. In general outdoor scenarios, absolute positioning using Global Positioning and Inertial Navigation systems (GPS-INS) is frequently performed, often in conjunction with input from odometry integration, and may augmented further by continuous registration of local terrain or obstacle maps. Indeed, recent GPS-INS devices advertise errors as low as a few centimeters in position and hundredths of a degree in attitude after alignment [13].

Many applications present significant challenges for these positioning strategies. GPS may be unavailable or ineffective if too few satellites are visible during urban or subterranean operations. Wheel-based odometry depends on an accurate kinematic model and can degrade greatly in the presence of wheel slip typical of low-friction

The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA, e-mail: {mdille3, grouch, ssingh}@andrew.cmu.edu

surfaces and skid-steered vehicles. In poorly lit or very low-texture environments, systems based on vision will see reduced performance.

Planetary rover missions pose particularly challenging cases. GPS is unavailable and alternatives such as star tracking do not offer comparable accuracy. Loose surface dust easily impedes wheel odometry. Further, visual odometry and terrain mapping methods are frustrated by poor lighting, few visual surface features, and tight computational constraints limiting their implementation complexity.

The specific need we seek to fulfill is that of a pose estimation system for Scarab, the lunar rover prototype developed at Carnegie Mellon University pictured in Figure 1 [4] [20]. This 280kg skid-steered vehicle is designed to explore permanently shadowed lunar polar craters, which combines the greatest weaknesses of most positioning methods: the unavailability of GPS, constant near-total darkness, occlusion of most stars by the high crater walls, soft lunar regolith, and strictly limited computing and power facilities. Exclusive reliance on wheel odometry was first planned but then abandoned in recognition of the errors that would be induced by surface slip, the eventual implementation of softer wheels that would thereby be of varying radius, and the passive differenced rocker suspension, which while excellent for maintaining stability on bumpy terrain, makes estimating heading changes from odometry inadequate. Certainly strategies can be applied for detecting slip some proportion of the time—and [16] provides several methods for doing just this on planetary rovers—but the design of Scarab’s suspension frustrates this approach and arguably more complicated physical modeling might best be replaced by an alternative sensing modality.

Conventional forward-facing visual odometry as used on the Mars Exploration Rovers (MERS) [11] was considered but deemed impractical because lighting the surrounding area would require more power than the rover can provide and induce complex shadows. Instead, we decided upon using downward-facing visual odometry. This provides tractable lighting requirements, however these forward-facing techniques cannot be directly applied because the situation is ill-posed to compute the 3-D incremental pose differences and the terrain beneath the rover is likely to be too homogeneous for point-based feature tracking to work effectively. Rather, we chose to rely on optical flow to provide an estimate of vehicle speed, which could then be integrated to estimate incremental distance-traveled as part of a broader odometry framework. Various methods, such as the Horn & Schunck algorithm [7], have long existed for computing so-called dense optical flow over regions between images. Such algorithms have been used, for instance, for autonomous heading control for obstacle avoidance for fixed-wing aircraft [21], estimating distance to the ground and canyon walls from unmanned aerial vehicles [6], and autonomous landing for helicopters [17]. In the odometry realm, downward-facing cameras have been successfully used for positioning in pre-explored environments by correlating the visible area against an existing database [8], however such methods are inapplicable to planetary rovers observing most patches for the first time.

Although an optical flow implementation using a typical camera would have been straightforward, an intriguing alternative presented itself in the form of a custom-built optical flow sensor from AirRobot GmbH [1] used for stabilization on its quad-

rotor helicopters. This device, shown in Figure 2, contains four commodity optical mouse sensors each attached to a lens of a different focal length and reports 2-D velocity (in m/s) along the planar surface at which it is pointed and an estimate of distance to this plane. The precise details of its operations are proprietary to the manufacturer, but given that a scalar “tracking quality” value is known to be returned by mouse sensors, we conjecture that height is derived by interpolating across the quality values returned by each of the four sensors, with speed similarly interpolated across the four focal-distance-compensated speed values. The sensor nominally reports velocity with a resolution of 0.3mm/s within a range of ± 10 m/s and height up to 2.5m with a resolution of 1cm.



Fig. 1 The Scarab lunar rover. The glow beneath **Fig. 2** The optical sensor used on Scarab, in a custom ruggedized enclosure.

The use of optical mouse sensors for measuring ground velocity has a number of benefits leveraging over a decade of commercial refinement. Key among these are remarkably robust operation on a wide variety of surfaces, significant lighting insensitivity, and extremely low cost (several US dollars) due to the volume in which they are produced. These typically contain a 15 to 100 pixels-square camera sensor and are believed to implement a fast hardware version of Horn & Schunck [14] reporting sub-pixel flow rates at on the order of 1kHz. Designs based on laser interferometry are becoming prevalent, however they are not as adaptable as they do not use simple lenses. Due to their low cost and simplicity, there exists wide interest in the hobbyist robotics community in these sensors. Several examples of their use in limited indoor scenarios exist in the robotics literature (eg [15], [18], and [19]), but they have yet to see use in a field robot application. A method similar to that described here using a webcam and image correlation matching was presented in [12].

In this paper, we describe our odometry method using this sensor, present results from extensive field testing, and draw conclusions on the effectiveness of commodity optical flow sensors for pose estimation.

2 Odometry Method

We first assume that the vehicle can only instantaneously move along its current heading vector, or equivalently assume the non-holonomic constraint of no wheel side-slip. We believe this to be valid because except for rare cases such as sliding laterally down a slope, wheel slip should occur only when skidding in place during turns or failing to grip the ground during forward travel. Position is integrated based

on the current linear speed and heading vector at all times. Linear speed is estimated from optical sensor readings, and the heading vector comes from the attitude determined by integrating IMU angular rate values. Uncertainty in the integrated pose is propagated by modeling uncertainty for the optical sensor and the IMU. The next section formalizes this procedure followed by methods for improved accuracy.

2.1 Basic Odometry Model

We begin by defining the vehicle position \mathbf{x} and orientation θ in a chosen world frame:

$$\mathbf{x}_t = [x, y, z]^T, \quad \theta_t = [\theta_x, \theta_y, \theta_z]^T \equiv \mathbf{q}_t = [q_s, q_x, q_y, q_z]^T \quad (1)$$

where the latter equivalence denotes that we may refer to the orientation as the Euler angles θ_t or a unit quaternion \mathbf{q}_t .

At each time-step, the IMU provides angular rates and linear accelerations in the body frame:

$$\boldsymbol{\omega}_t = [\omega_x, \omega_y, \omega_z], \quad \mathbf{a}_t = [a_x, a_y, a_z]. \quad (2)$$

For a given state \mathbf{x}_t we may then write the vehicle unit heading vector $\hat{\mathbf{u}}_t$ in the world frame. The non-holonomic velocity constraint can then be embedded in the definition of the vehicle's linear velocity as

$$\mathbf{v}_t = v_{body} \hat{\mathbf{u}}_t \quad (3)$$

where v_{body} is the scalar linear velocity, which is assumed to be directly provided by the optical sensor. In actuality, it reports velocity in two dimensions, but to avoid having to very precisely calibrate for any rotational offset, we simply take the norm of the velocity vector it reports to be the speed and determine the sign from that of the velocity reported along the axis most nearly aligned with the body.

At time $t + 1$, given \mathbf{v}_{t+1} and the previous values \mathbf{x}_t and \mathbf{v}_t , the new position may be estimated as

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \frac{1}{2} \Delta t (\mathbf{v}_{t+1} + \mathbf{v}_t), \quad (4)$$

where Δt is the time elapsed between t and $t + 1$, computed via trapezoidal integration of the velocity vector. A similar procedure for orientation is used. Namely,

$$\begin{aligned} \dot{\mathbf{q}}_t &= \frac{1}{2} \mathbf{q}_t * [0, \frac{1}{2} (\boldsymbol{\omega}_t + \boldsymbol{\omega}_{t+1})]^T \\ \mathbf{q}_{t+1} &= \mathbf{q}_t + \Delta t \dot{\mathbf{q}}_t \end{aligned} \quad (5)$$

estimates the new orientation, where $*$ denotes quaternion multiplication [10].

Dead-reckoning error uncertainty is propagated by the linearized covariance prediction equation ¹

$$\mathbf{P}_{k+1} = \mathbf{F}_k \mathbf{P}_k \mathbf{F}_k^T + \mathbf{G}_k \mathbf{Q}_k \mathbf{G}_k^T, \quad (6)$$

where

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & -V \sin \psi \Delta t \\ 0 & 1 & V \cos \psi \Delta t \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{G} \mathbf{G}^T = \begin{bmatrix} \cos^2 \psi \sigma_V^2 \Delta t + \frac{1}{3} V^2 \sin^2 \psi \sigma_\omega \Delta t^3 & \cos \psi \sin \psi (\sigma_V^2 \Delta t - \frac{1}{3} V^2 \sigma_\omega^2 \Delta t^3) & -\frac{1}{2} V \sin \psi \sigma_\omega^2 \Delta t^2 \\ \cos \psi \sin \psi (\sigma_V^2 \Delta t - \frac{1}{3} V^2 \sigma_\omega^2 \Delta t^3) & \sin^2 \psi \sigma_V^2 \Delta t + V^2 \cos^2 \psi \sigma_\omega^2 \Delta t^3 & \frac{1}{2} V \cos \psi \sigma_\omega^2 \Delta t^2 \\ -\frac{1}{2} V \sin \psi \sigma_\omega^2 \Delta t^2 & \frac{1}{2} V \cos \psi \sigma_\omega^2 \Delta t^2 & \sigma_\omega^2 \Delta t \end{bmatrix}.$$

Approximate trends in state uncertainty growth over time can be observed by reducing Equation 6 to straight line constant speed motion with initial covariance $\mathbf{P}_k = \text{diag}(\mathbf{P}_x \ \mathbf{P}_y \ \mathbf{P}_\psi)$. For x-axis-aligned motion, Equation 7 indicates along-track

¹ For brevity, we here present uncertainty propagation for planar location and heading. The equations for the full 6-D model as we have implemented are a straightforward extension.

and heading error increases proportional to the square-root of elapsed time. Cross-track error exhibits faster error growth, linear in time due to initial heading uncertainty plus growth at three-halves power of time due to heading rate noise. This highlights the importance of low heading-rate uncertainty in achieving accurate dead-reckoning.

$$\mathbf{P}_{k+\Delta t} = \begin{bmatrix} \mathbf{P}_{x,k} & 0 & 0 \\ 0 & \mathbf{P}_{y,k} + V^2 \Delta t^2 \mathbf{P}_{\psi,k} & V \Delta t \mathbf{P}_{\psi,k} \\ 0 & V \Delta t \mathbf{P}_{\psi,k} & \mathbf{P}_{\psi,k} \end{bmatrix} + \begin{bmatrix} \Delta t \sigma_V^2 & 0 & 0 \\ 0 & \frac{1}{3} V^2 \Delta t^3 \sigma_\omega^2 & \frac{1}{2} V \Delta t \sigma_\omega^2 \\ 0 & \frac{1}{2} V \Delta t \sigma_\omega^2 & \Delta t \sigma_\omega^2 \end{bmatrix} \quad (7)$$

Distance and heading error drift rates for the system have been determined through ground-truthed experimental trials. Figure 3 shows the difference in changes observed between the dead-reckoning solution and ground-truth for varying length-time ensembles. Twenty thousand samples were compared over a two hour period at thirty second ensemble increments. This trial indicates drift rates of $1.2m/\sqrt{hour}$ in distance traveled and $2.5^\circ/\sqrt{hour}$ in heading.




Fig. 3 Estimated drift rate for distance traveled (*left*) and heading (*right*) determined by comparing changes over increasing time ensembles between the dead-reckoning solution and ground-truth.

2.2 Improved Attitude Estimation

While the error in the pose estimate provided by the above integration procedure will necessarily grow without bound, several simple measures can be taken to greatly reduce attitude error. Many of these are intuitive and even commonplace in commercial navigation systems, however they receive rare mention in the robotics literature and are too often ignored in implementation.

A first strategy critical for long-term operation is to remove the approximately 15 degrees per hour on Earth (or 0.56 on the moon [2]) angular velocity the gyroscope will inherently pick up due to the rotation of the earth, which is done by rotating the known angular velocity of the planet into the local coordinate system and subtracting. When stationary, this measured angular velocity may even be used to perform gyrocompassing. The gyroscope bias may be computed by averaging the measured angular velocity vector while stationary (after subtracting the planetary rotation rate). With some IMUs this may be needed as often as every ten minutes for

even short-term performance [5]. This step was not required in our case since this value typically averaged to be negligible.

Another common technique is to use the gravity vector (weaker but still useful at about 0.16g on the moon [3]) as measured by the IMU’s accelerometers when stationary to estimate roll and pitch by taking arctangents of the accelerations along the axes. We made use of this to reset drift in roll and pitch during stationary periods. Given now two complementary sources of roll and pitch—this direct computation most accurate when the attitude is slowly changing and integration of angular rates just the opposite—we implemented a matched complementary Butterworth filter pair to continuously merge these two streams, with a cutoff frequency of 0.05Hz found to be most appropriate given the very slow motion of the rover. This provides excellent results with errors in roll and pitch of typically much less than a degree at all times.

Finally, anticipating missions consisting of long stationary periods (primarily intended for battery recharging) followed by short periods of motion, we additionally clamped the measured angular velocity vector to zero when the vehicle is known to be stationary to avoid blatantly unnecessary noise integration.

2.3 Velocity Scale Calibration

Ideally, the optical flow sensor returns a correctly scaled lateral speed regardless of the distance between the sensor and the ground. We model the actual corrupted scaled velocity returned as

$$v_{measured} = s(h) \cdot v_{actual} + b + v, \quad (8)$$

where $s(h)$ is a scaling factor that varies with the height h , and v is noise. At most a trivial bias term b was observed, and this was effectively dealt with by clamping the measured velocity to zero when the vehicle is known to be stationary.

The height of the sensor is not fixed as Scarab is equipped with an active suspension used to lower the body so that a core-drilling apparatus may operate. Calibration trials at varying body heights were performed. Figure 9 shows a number of such runs over distances varying between 10 and 30 meters. As it indicates, rather different scaling values were found for daytime and nighttime operation, but within each class, a simple quadratic fit provides reasonable compensation.

The sensor also exhibited a scale dependence on the surface type with values tightly clustered for a given surface. An online auto-calibration scheme proved quite effective by relying on the observation that during consistent straight-line motion, wheel odometry can be very accurate. Every three seconds, a procedure runs that tests a battery of heuristics to decide whether the wheel odometry from that period is likely to be trustworthy. This includes for instance verifying that each wheel’s velocity agrees closely with the average, all agree on direction, the reported velocity is above a noise floor, and that the IMU is reporting minimal yaw rates. Primarily, this eliminates periods including heavy slip or turning (implying likely wheel odometry error). Each period is added as a learned data point over which a variation of a recursive locally-weighted linear least squares algorithm is run to compute the scaling factor as a function of body height. Each time a new velocity value from the optical sensor is received, the best estimate of the scaling factor for that height is

used. If insufficient calibration data points are available for the region surrounding that height the scaling factor determined by the original manual calibration trials is used as a fall-back.

An example run using the online scale calibration procedure is shown in Figures 10 and 11, which indicates that this process performs well, though unsurprisingly not as well as a post-hoc batch method computing the scaling function from all the data points at once (shown in Figure 12). However, it performed well enough and eliminated most difficulties associated changing surfaces. All data presented uses it unless otherwise indicated.

3 Experimental Results

During the course of system development and verification, we collected tens of hours of data from field testing over several kilometers of traversal, including extended simulated lunar terrain at Moses Lake, Washington and Mauna Kea, Hawaii. Field tests were conducted both during daytime and in total darkness, for which high-intensity LED bars were mounted to the vehicle's underbelly to illuminate the area seen by the optical sensor.

Anecdotally, the optical sensor worked remarkably well across a wide variety of surfaces, including dirt, mud, grass, asphalt, and sand. Tracking was poor on poured concrete surfaces and nearly useless on painted concrete. Most of the extended field testing took place over sand to best emulate lunar terrain, and it was noted that typical auto-calibrated scaling factors were somewhat higher than for other surfaces, likely explainable by somewhat poorer tracking on this relatively featureless terrain.

As the Honeywell HG1700 IMU used for odometric integration is part of a NovaAtel SPAN GPS-INS system, a convenient source of approximate ground truth was also available. At the slow speeds the rover moves, INS position errors on the order of tens of centimeters are large relative to short-term distances traveled, however over long distances, these relatively static uncertainties are small compared to accumulated odometry error.

In the plots below, full odometry results are provided as well as the result of using INS-provided attitude (the output of the INS Kalman filter that uses GPS data) in place of integrated attitude. The purpose of this is to independently demonstrate the effectiveness of the optical sensor (accumulating just distance-traveled error in the odometry) and because further methods of improving heading accuracy would be a priority in any future implementation.

An early observation was that even when the odometry-derived pose drifted significantly, as shown in Figure 8 having atypically high gyro drift, the smoothness of the integrated solution was much more useful to the on-board perception system than that from the INS. Laser scanners are used to build local terrain and obstacle maps. Pose jumps present false obstacles without resorting to complicated registration algorithms precluded by limited on-board computing. While a map built from the integrated solution will not be globally correct, the resulting higher-fidelity short-term maps are more valuable for local motion planning.

As odometry error is inherently path dependent, errors accumulated over different trials varied, however, the results shown in Figure 6 are representative. Data from this run, a nearly 700m nighttime traverse at Moses Lake, is also used as the examples for the described attitude improvement and scale calibration methods. Attitude comparisons between the INS-reported values and integrated results are given in Figure 5. Traces of easting, northing, and heading uncertainty, propagated using the model described and the estimated covariances on angular rates and optical flow velocity, are provided in Figures 13-15. Though this is just one example, the errors lie within the uncertainties, lending credence to our propagation model and process covariances.

4 Conclusions and Future Work

In this paper, we presented an implementation of vehicle odometry for a lunar rover prototype using an optical mouse sensor to provide vehicle velocity. Results show that with a moderate-grade IMU, errors can be small over long distances. Clearly the weakest point in such a design is the accumulation of heading error, the reduction of which would be a key focus of future implementation. The lunar application may present some such reduction opportunities, perhaps via the stop-start nature of motion during missions or observation from an orbiter. Other avenues of exploration include using multiple sensors to track heading or relax the kinematic assumptions [9] and designing a new sensor with focal lengths tuned for the mounting height.





Fig. 4 Comparison of INS-reported and integrated yaw angle after application of described attitude integration improvements. **Fig. 5** Error in integrated yaw angle relative to integrated yaw angle after application of described attitude integration improvements.





Fig. 6 Overhead comparison of odometry results against INS-reported position using both integrated angular rates and INS-reported attitude. **Fig. 7** Altitude comparison of odometry results against INS-reported position using both integrated angular rates and INS-reported attitude.





Fig. 8 Example run in which even the significantly drifting integrated estimate is preferable to over a number of short straight-line runs. **Fig. 9** Optical velocity scale calibration data from a very jumpy GPS signal for perception purposes.





Fig. 10 Overhead comparison of optical velocity scaling methods with wheel odometry and INS methods with wheel odometry and INS position. **Fig. 11** Comparison of optical velocity scaling methods with wheel odometry and INS methods with wheel odometry and INS position, zoomed in to the end of the run.





Fig. 12 Points used for batch velocity scale calibration. A linear fit—mostly just capturing a scale bias in this operation region—is overlaid. **Fig. 13** Easting position error and prediction.

Acknowledgements

This research was supported in part by NASA under grants, NNX07AE30G, John Caruso, Project Manager, and NNX08AJ99G, Robert Ambrose, Program Scientist. The authors wish to thank the Scarab project team for expressing interest in this idea and tolerating related experimentation.





Fig. 14 Northing position error and prediction. Fig. 15 Heading error and prediction.

References

1. AirRobot GmbH: AirRobot website [online] (2009). URL <http://www.airrobot.de>
2. Allen, C.W., Cox, A.N.: Allen's Astrophysical Quantities. Springer (2000)
3. Baldwin, R.: A Fundamental Survey of the Moon. McGraw-Hill, New York (1965)
4. Bartlett, P., W., D.W., Whittaker: Design of the scarab rover for mobility and drilling in lunar cold traps. In: International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS). Los Angeles, CA (2008)
5. Chung, H., Ojeda, L., Borenstein, J.: Accurate mobile robot dead-reckoning with a precision-calibrated fiber optic gyroscope. *Trans. on Robotics and Automation* **17**(1), 80–84 (2001)
6. Griffiths, S., et al.: Maximizing miniature aerial vehicles. *Robotics & Automation Mag.* (2006)
7. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial Intelligence* **17**, 185–203 (1982)
8. Kelly, A.: Mobile robot localization from large-scale appearance mosaics. *International Journal of Robotics Research* **19**(11), 1104–1125 (2000)
9. Kim, S., Lee, S.: Robust mobile robot velocity estimation using redundant number of optical mice. pp. 107–112 (2008). DOI 10.1109/ICINFA.2008.4607977
10. Kuipers, J.B.: Quaternions and Rotations Sequences. Princeton University Press (1999)
11. Maimone, M., Cheng, Y., Matthies, L.: Two years of visual odometry on the Mars Exploration Rovers: Field reports. *J. Field Robot.* **24**(3), 169–186 (2007)
12. Nourani-Vatani, N., Roberts, J., Srinivasan, M.V.: IMU aided 3D visual odometry for car-like vehicles. Canberra, Australia (2008)
13. NovAtel Inc.: HG 1700 SPAN-58 Specifications (2008). URL http://www.novatel.com/Documents/Papers/HG1700_SPAN58.pdf
14. Owens, R.L.: Optical mouse technology [online] (2006). URL <http://web.archive.org/web/20080112045313/http://www.mstarmetro.net/~rlovens/OpticalMouse/>
15. Palacin, J., Valgaon, I., Pernia, R.: The optical mouse for indoor mobile robot odometry measurement. *Sensors and Actuators A: Physical* **126**(1), 141–147 (2006)
16. Reina, G., Ojeda, L., Milella, A., Borenstein, J.: Wheel slippage and sinkage detection for planetary rovers. *Mechatronics, IEEE/ASME Transactions on* **11**(2), 185–195 (2006)
17. Saripalli, S., Sukhatme, G.S.: Landing on a moving target using an autonomous helicopter. In: *Field and Service Robotics*, pp. 277–286 (2003)
18. Sekimori, D., Miyazaki, F.: Precise dead-reckoning for mobile robots using multiple optical mouse sensors. *Informatics in Control, Automation and Robotics II* pp. 145–151 (2007)
19. Sorensen, D.: Online optical flow feedback for mobile robot localization/navigation. Master's thesis, Texas A&M University (2003)
20. Wettergreen, D., et al.: Design and experimentation of a rover concept for lunar crater resource survey. In: *AIAA Aerospace Sciences*. Orlando, FL (2009)
21. Zufferey, J.C., Floreano, D.: Fly-inspired Visual Steering of an Ultralight Indoor Aircraft. *IEEE Transactions on Robotics* **22**(1), 137–146 (2006). DOI 10.1109/TRO.2005.858857