Optical Kinematic State Estimation of Planetary Rovers using Downward-Facing Monocular Fisheye Camera

Sudharshan Suresh
National Institute of Technology
Tiruchirappalli, India
Email: sudharshan.nitt@gmail.com

Eugene Fang
Robotics Institute
Carnegie Mellon University
Pittsburgh, USA
Email: fangeugene@cmu.edu

William L. "Red" Whittaker
Robotics Institute
Carnegie Mellon University
Pittsburgh, USA
Email: red@cmu.edu

***Abstract*—Knowledge of the kinematic state of rovers is critical to navigation, path reconstruction and exploration, especially on rugged terrain like planetary surfaces. Existing methods employ many encoders, potentiometers and hall sensors. These add components and wiring to moving parts. The components are susceptible to mechanical and electronic failures, add mass, and some require thermal regulation. In addition, the sensor wires are susceptible to bending, flexing and wear. Where miniaturization counts, the limitations on mass, size and power encourage elimination of sensors wherever possible. This paper presents a method to estimate the kinematic state of rovers using only a downward-facing fisheye camera. This novel approach implements a vision algorithm to obtain kinematic state information in planetary rovers. The two additional benefits of the technique are (1) redundancy to proprioceptive measurements, (2) means for perceptive visual odometry. The method uses a single camera to estimate 10 degrees-of-freedom (associated with steering, driving and suspension) on the *AutoKrawler*, a rover test platform for planetary exploration. Motions are estimated by self-perception - combining fiducial marker tracking, optical flow techniques and the kinematic constraints of rover mechanics. Experimental results, obtained from the rover operating in an environment analogous to the lunar surface, are presented. The results obtained are compared with ground truth data to validate the approach.**

I. Introduction

Planetary rovers have locomotive capabilities designed for traversing highly uneven terrain. A rover's axles or suspension act as passive hinges, allowing the wheels to make contact with uneven surfaces. Steering is critical to negotiating terrain. Understanding the real-time kinematic state of a rover provides timely information about abnormal motion and supports operational decisions. Such knowledge proves useful for rover guidance, navigation and safeguard. This information is also fundamental to modeling rover odometry. Proprioceptive sensing has for long been the modus operandi to obtain the aforementioned data, providing accurate estimates of position, orientation and motion of links and joints.

Space technology typically employs a large number of sensors to obtain kinematic state information. For example, the states of the Mars Rover prototypes are determined by six wheel encoders, three accelerometers, joint potentiometers for bogey angles and a sun sensor [1]. The problems with such a system are manifold, and of significant interest in the domain of space robotics. Optical encoders have rotating mechanical disks that reduce system reliability. Atmospheric dust, similar to that of Mars, may hinder optical wheel encoders [2]. A pressing concern with encoders is the numerous electrical conductors required. The Sojourner rover comprised of ten motors, each with an optical encoder. The motors required two conductors each to drive, but the encoders required six wires each to moving parts- translating to 80 conductors fed from the body of the rover [3]. In an articulated rover, special care must be taken in routing cables to prevent wear of wire harnesses. In addition to flexing and bending with suspension, when exposed to the frigid environment, the wires can stiffen and break [3]. Sensing elements with electronic components may require thermal protection and this additional insulation can compromise total system mass. 50% of a typical rovers mass distribution is generally dedicated to its subsystems [4], and the myriad sensors reduce the viable scientific payload limit.

Figure 1: The AutoKrawler traversing uneven terrain in a lunar analogue site.

In the past, vision algorithms were used sparingly on planetary missions, primarily due to the on-board computational constraints. A major leap in this direction was the NASA/JPL Mars Exploration Rover (MER) mission, where algorithms performed tasks such as visual odometry, stereo vision and feature tracking [5]. All processing was performed off-board, either on Earth or the lander. Subsequent FPGA implementation of the same algorithms have shown performance improvements of three orders of magnitude and are now utilized on-board large rovers [6]. These encouraging results point to full-fledged vision systems in future missions with minimum latency, to perform rover odometry and safeguarding, even on small, simple rovers.

This research conceives and demonstrates a novel method of optical kinematic state estimation of planetary rovers using a downward-facing monocular fisheye camera. In this paper, 10 degrees-of-freedom of a planetary rover are estimated. Fiducial marker tracking is used to obtain pixels of interest from the camera image. Using a spherical camera model, 2-D image coordinates are projected to 3-D points on a unit sphere around the camera's focal point. These are subsequently mapped to real-world coordinates using planes defined by the rover's kinematic constraints. 8 defined positions are tracked and a geometric approach is developed to determine the axle roll and steering angles. Odometry is performed for all four wheels by using an optical flow algorithm directly on the treads, as opposed to traditional egomotion estimation. The methodology is exhibited and evaluated by application to the AutoKrawler (see Figure 1), a highly-versatile, double-Ackermann rover test platform for planetary exploration.

The aforementioned system was evaluated on datasets generated from field experiments conducted in a lunar analogue site. The rover was tele-operated in conditions that mimic risk-prone traversals, where knowledge of kinematic state can aid control decisions. Results show close agreement between data from our technique and ground truth data from proprioceptive sensors onboard.

II. Related Work

The kinematic study of complex rocker-bogey mobility systems have been a prominent research problem. [7] and [8] describe a body of work surrounding the Rocky-7 Mars Rover that estimates position, velocities, contact angles and orientation of the robot. Systems inputs are generally accelerometer, gyroscope and encoder data, without involving visual observation. Lamon and Seigwart demonstrate the benefit of kinematic knowledge by performing 3-D odometry, showing a significant improvement in results on rough terrain. A controller that minimizes slip based on state information is also developed [9]. Visual self-perception (or the strategy of observing oneself), has been used to generate kinematic models of robotic manipulators [10]. Its joint configuration however, is determined solely by its own actuation and not by the environment.

Optical flow methods using pyramidal Lucas-Kanade [11] have been previously explored using a downward-facing fisheye by Seegmiller [12] to perform robust visual odometry. The algorithms employ RANSAC outlier rejection, which has been incorporated in this body of work. Fang's work also highlights the importance of a minimally obstructing, robust camera support structure for the selected configuration [13]. Fiducial markers have been used, albeit sparingly, in planetary robotics for accurate instrument positioning and as calibration targets [14]. Scaramuzza obtained a generalized omnidirectional camera model for fisheye and catadioptric systems, that has been used here [15].

III. Kinematic Model

A. Rover Overview

The AutoKrawler is a four-wheel, double-Ackermann steered rover specialized to traverse adverse terrain. A passive body-axle suspension system allows the rover to maintain compliance with uneven, rugged terrain. For an illustration of the Cartesian coordinate conventions and degrees-of-freedom, refer to Figure 2. Each articulated axle of the dual-axle configuration can perform roll about the y-axis and translation about the z-axis. The steering swivel joints can rotate in the range of [-30°, 30°] about the axle frame's z-axis.

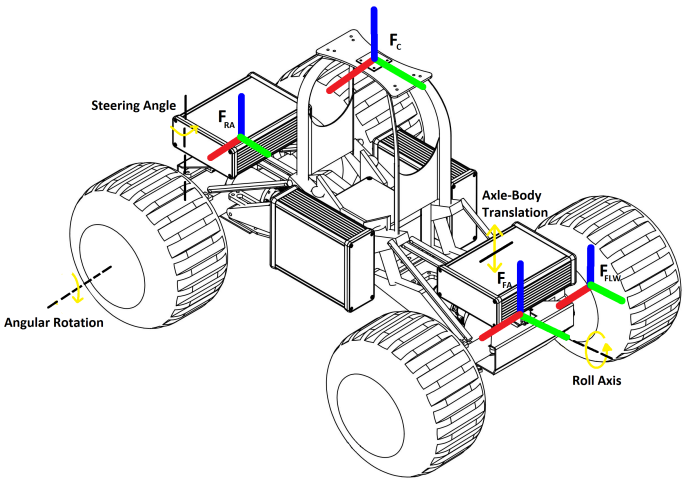

Figure 2: Depiction of (i) The coordinate frames of the camera, axle and wheels, illustrated with XYZ (**RGB**) axes. (ii) The permitted rotational and translational motions of the robot about their respective axes.

In conceptualizing the state estimation, several reasonable assumptions on rover kinematics are made:

- Angular movement of wheels on a common axle are assumed to be identical - giving each axle a single steering angle.
- The mechanical connection between the axle and body of the AutoKrawler comprises of a plurality of links and joints. However, for all practical purposes, the axles are

assumed to only exhibit one rotational degree-of-freedom (about y-axis).

- It is assumed that there is no relative translation in the x-direction between the axle and the body frame. This allows us to approximate the motion of each axle to be on a vertical plane perpendicular to the y-axis.
- The wheels of the rover are modelled as cylinders (of radius $r = 11.5cm$) that rotate with the wheel frame, centered at its origin.

B. Coordinate Frames and Variables

The translational and rotational displacements of rover joints and linkages are expressed via homogeneous transformation matrices. Examples of coordinate frame locations are provided in Figure 2. With transformation matrices, a kinematic chain of coordinate frames can be represented. Below, an equation is used to transform the center of the wheel (P) from the swivel joint frame to the camera frame. To do so, the 4×1 vector $[x \quad y \quad z \quad 1]^\intercal$ is successively pre-multiplied with ${}^A_S T$ (transform from swivel joint to axle frame) and ${}^C_A T$ (transform from axle to camera frame):

$$ {}^C P = {}^C_A T \, {}^A_S T P \tag{1}$$

In (1) ${}^C P$ represents the real-world coordinates of a wheel center, with the reference frame being the camera focal point. The degrees-of-freedom estimated are 10 in total:

(i) Two rotational (Axle roll angles) - ψ_F, ψ_R
(ii) Two translational (Axle-Body vertical distances) - d_F, d_R
(iii) Two steering - λ_F, λ_R
(iv) Four rotational - $\theta_{FL}, \theta_{FR}, \theta_{RL}, \theta_{RR}$

IV. System Description

A. Camera Model

The research problem necessitates that all motions (axle roll, steering and wheel rotation) of the rover be detected with a single camera. The best way to achieve this is with a downward-facing omnidirectional camera. The choice of a fisheye lens over a catadioptric lens is inconsequential, as a unified camera model [15] is applicable to both. The body frame is fit with a mount to house the downward-facing fisheye camera. The mount minimally occludes field-of-view, with four orthogonal plates along planes that intersect the camera center (as seen in Figure 3). Color images are acquired with a resolution of 640×480 at 30 frames per second. The chosen resolution provides adequate pixels to work with, while at the same time remaining computationally feasible.

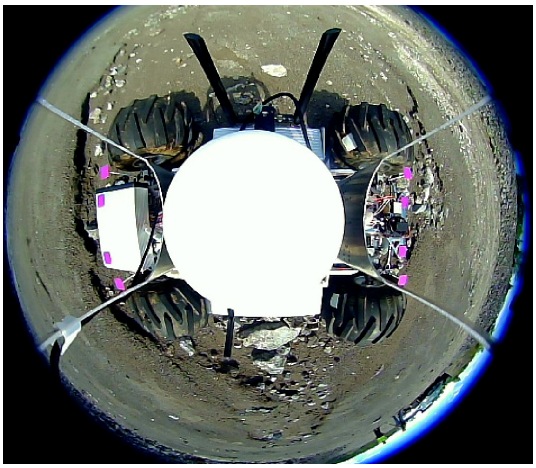

Figure 3: Fisheye lens's field-of-view with minimally occluding camera mount. The four thin plates in the image constitute the mount structure.

The MATLAB toolbox developed by Scaramuzza et al. [16] is used to calibrate the camera with a planar checkerboard pattern. The model uses affine transforms to handle misalignment between the camera's optic center and the focus point of the lens. A 4^{th} order polynomial is utilized to account for the camera's radial distortion. Any real-world point can be accurately projected to a point on a unit sphere, with the camera's focal point as its center.

B. Fiducial Markers

Passive, mono-colored square markers are used, and the paper does not attempt to validate tracking robustness. In our application, magenta is chosen due to its hue value being distinct to that of the rover's surrounding. Eight markers are placed on the rover - two collinear pairs on the front and rear axle plates, and two pairs on the Ackermann tie rods.

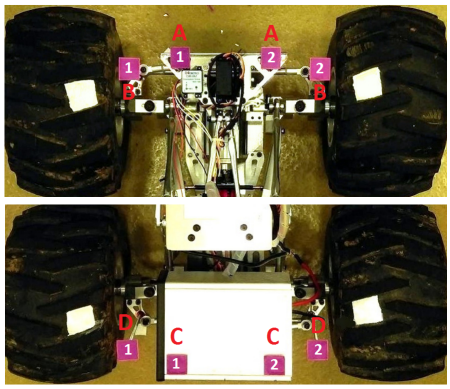

Figure 4: Mono-color markers visible on the front and rear axles. Pairs A and C track front and rear axle roll respectively, while pairs B and D track front and rear steering angles. Also seen are the white correction markers present on the wheels.

The process pipeline first performs color-based image segmentation in the HSV (Hue, Saturation, Value) color space. After noise-removal by erosion and dilation, pixel positions of the contour centers are obtained. These pixel positions are transformed to world coordinates based on the geometric constraints, as described in subsections V-A and V-B. Each wheel tread consists of a white marker, referred to henceforth as a correction marker. It is tracked via thresholding to correct wheel angular rotation, as described in subsection V-C.

C. System Design

A compositional view of the required variables is taken and the resulting system is simple and modular (see Figure 5). Axle roll, axle-body vertical distance and steering angles (ψ_F, ψ_R, $d_F, d_R, \lambda_F, \lambda_R$) are computed at once, requiring fisheye image data. Wheel angular rotations ($\theta_{FL}, \theta_{FR}, \theta_{RL}, \theta_{RR}$) require all the aforementioned variables *and* fisheye image data. This is because optical flow is performed only on dynamic observation windows, as described in V-C. If these variables are not present, due to momentary loss of marker tracking, flow operations are performed using the prior state of the rover. In the absence of computational overhead, the system operates at the frame rate of the camera. All vision tasks are optimized to defined regions-of-interest based on kinematic constraints, instead of being performed on the entire image.

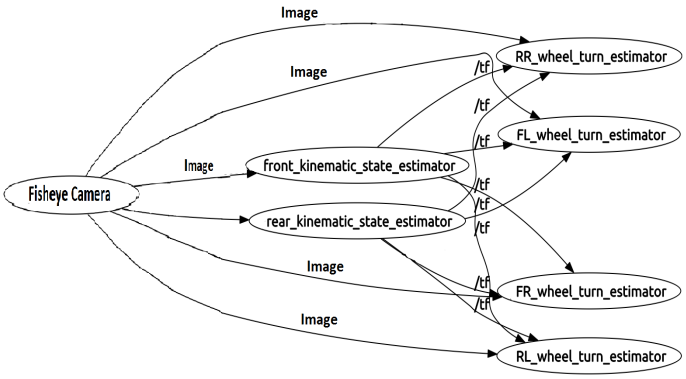

Figure 5: ROS Nodes and Topics and their relationship in the system. *tf* indicates transform matrices of the rover.

V. Estimation Methods

The following subsections describe how the degrees-of-freedom are obtained. All the estimation methods use **world coordinates** as inputs, and not pixel positions.

A. Estimation of Axle Configuration

The kinematic state estimation begins by calculating axle roll (ψ_F, ψ_R). By making the assumption that the axle has only one rotational and one translational degree-of-freedom (refer to Section III-A), markers are constrained to a plane. This is defined by:

$$y = \begin{cases} +c, & \text{front axle} \\ -c, & \text{rear axle} \end{cases} \tag{2}$$

where $c \in \mathbb{R}_{>0}$.

As explained in section IV-A, the omnidirectional camera projects a pixel to a point on a unit sphere - (a, b, c). The world coordinates of the marker centers are computed by finding the intersection point between the ray through the points $\{(0,0,0), (a,b,c)\}$ with the plane defined in equation 2. The two marker coordinates are subtracted to obtain a vector that represents axle orientation, $\vec{v} = x\hat{i} + y\hat{j} + z\hat{k}$. Finally, ψ is computed as the directed angle between $\vec{v}$ and vector $\vec{u} = \hat{i}$ (unit vector on x-axis). The midpoint of the two marker coordinates is the rotational center of the axle, and its vertical distance from the camera frame is the axle-body distance (d).

B. Estimation of Steering Angle

In steering angle estimation, a plane that rotates about the y-axis (similar to axle rotation) is considered. In the absence of axle excursion, the normal to the plane is $[0 \quad 0 \quad 1]^\intercal$. However, the axle roll must be accounted for by computing a new normal:

$$\vec{n'} = R_y(\theta) \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} \sin\theta \\ 0 \\ \cos\theta \end{bmatrix} \tag{3}$$

where θ is the angle of roll and $R_y(\theta)$ is an elemental rotation about the y-axis. The position vector of the axle's rotational center is represented by $\vec{r_o}$.

Thus, the plane considered is defined as:

$$\vec{n'} \cdot (\vec{r} - \vec{r_o}) = 0 \tag{4}$$

Similar to section V-A, the ray-plane intersection is computed to get centers of the steering markers in world coordinates. Euclidean distance between marker pairs is calculated and the given ratio is defined :

$$x = \begin{cases} \frac{d(\mathbf{B[1]},\mathbf{A[1]})}{d(\mathbf{B[2]},\mathbf{A[2]})}, & \text{front axle} \\ \\ \frac{d(\mathbf{D[1]},\mathbf{C[1]})}{d(\mathbf{D[2]},\mathbf{C[2]})}, & \text{rear axle} \end{cases} \tag{5}$$

where marker labels $(A, B, C, D; 1, 2)$ are assigned as per Figure 4. The geometric arrangement of the tie rod ensures that the ratio is ≈ 1 when the wheels aren't steered. The ratio is *greater* than 1 if steered in one direction, and *less* than 1 is steered in the other. To obtain a relation between the ratio x and wheel alignment, manual calibration is performed. x vs. steering angles is tabulated, where steering angle is set to values in the range of [-30°, 30°]. The calibration equation is a second-order polynomial $f(x)$ fit to the data, such that $\lambda = f(x)$ (see Figure 6).

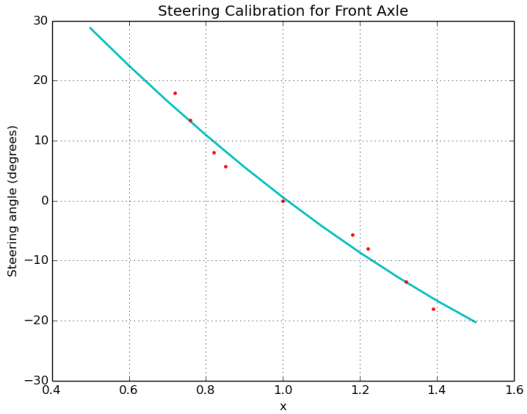

Figure 6: 2nd degree polynomial regression for steering angle vs. distance ratio. The red points indicate tabulated readings and the cyan curve is the polynomial fit.

C. Estimation of Wheel Angular Rotation

Optical flow provides vital knowledge of the arrangement of features in an image, and change in this arrangement over a sequence of frames. Egomotion estimates can be obtained from analysis of image sequences, as performed in [2], [12] and [13]. It has computational advantages over structure from motion algorithms, making it ideal to run online.

Lucas-Kanade optical flow is chosen, an established algorithm for correlated feature tracking [11]. Sparse feature sets are identified by the Shi-Tomasi corner detection algorithm [17]. Pyramidal implementation of Lucas-Kanade [18] relaxes the small displacement constraints of the algorithm and works with coarse variations. In our image sequences, flow is computed in the forward and backward direction with respect to time. All feature vectors that are inconsistent between frames are discarded.

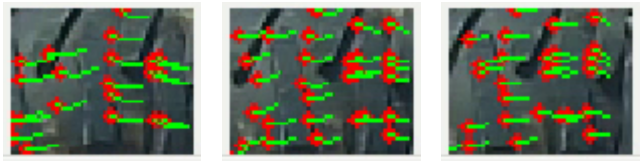

Figure 7: Lucas-Kanade feature tracking over three consecutive frames. The red dots indicate feature positions and the green trails show the direction of motion. There is a shift in features towards the left, and features may disappear and appear between frames.

The technique is applied on an unrectified fisheye image sequence to give 2-D displacement vectors. As highlighted in III-A, the rover's wheels are modelled as finite cylinders. It is defined by:

(i) A fixed radius of $r = 11.5cm$
(ii) Cylinder axis vector $\vec{v} = \vec{a} - \vec{b}$, where $\vec{a}$ and $\vec{b}$ are the base center position vectors in world coordinates. These points are obtained with knowledge of ψ and λ, along with the position of the swivel joints in the axle frame and the wheel centers in the swivel joint frame.

Similar to ray-plane intersection computed in V-A, the ray-cylinder intersection point is computed. The intersection point represents the world coordinate position of a point on the wheel's surface.

Feature search is performed in a dynamic observation window of fixed area, so as to only focus on the wheel tread (as seen in Figure 7). The position $w_{60} = [0 \quad \pm r\cos 60 \quad r\sin 60]$ with respect to the wheel coordinate frame is mapped back to a 2-D pixel coordinate and treated as the window center. This ensures that the flow is invariant to axle roll/translation and steering angle. Between frames, the arc-length travelled in world measurements is computed. To get angular displacement $\alpha(i)$ of a feature i, given its position on the circle's edge in consecutive frames (p_1 and p_2):

$$\alpha(i) = \cos^{-1}\left(\frac{2r^2 - |p_1 - p_2|^2}{2r^2}\right) \tag{6}$$

α is a $1 \times n$ size matrix, where n is the number of features tracked. Rather than computing the mean angular displacement directly, Random sample consensus (RANSAC) is used to reject outliers. These outliers may appear as a result of self-shadowing or occlusion. The RANSAC algorithm begins by selecting random angular displacements from the matrix α. The randomly selected value is compared with rest of the set, and all those that are approximately equal (within a tolerance range *t*) are treated as inliers for the selected model. After a defined number of iterations, the model with the maximum number of inliers is considered and its outliers are discarded. The mean of this set is taken to be the final angular displacement between the frames. The angle of rotation is incremented every frame, and the derivative of angular displacement with respect to time gives wheel angular velocity.

The demerit of a sparse feature set is the loss of tracking due to a lack of *interesting* features, or if there is image blur. When incremental angular rotation is observed, these inaccuracies accumulate. This can lead to θ drifting far from its ground truth over long traverses. Correction markers on the wheels (visible in Figure 4) clearly indicate each complete rotation. They are used to prevent this drift in angular rotation. The marker is segmented out by thresholding, and once every rotation, $theta$ is corrected:

$$\begin{aligned} m &= (\theta - c_o) \ div \ 2\pi \\ \theta_{corrected} &= 2m\pi + c_o \end{aligned} \tag{7}$$

where $m \in \mathbb{Z}$ and c_o is the value of θ when the first correction marker was encountered.

VI. Results

The performance of the method was evaluated on datasets acquired in an outdoor field experiment (see Figure 8). The rover was tele-operated at a location where the terrain characteristics and features served as an excellent analogue to lunar conditions.

Figure 8: Optical kinematic state estimation field test.

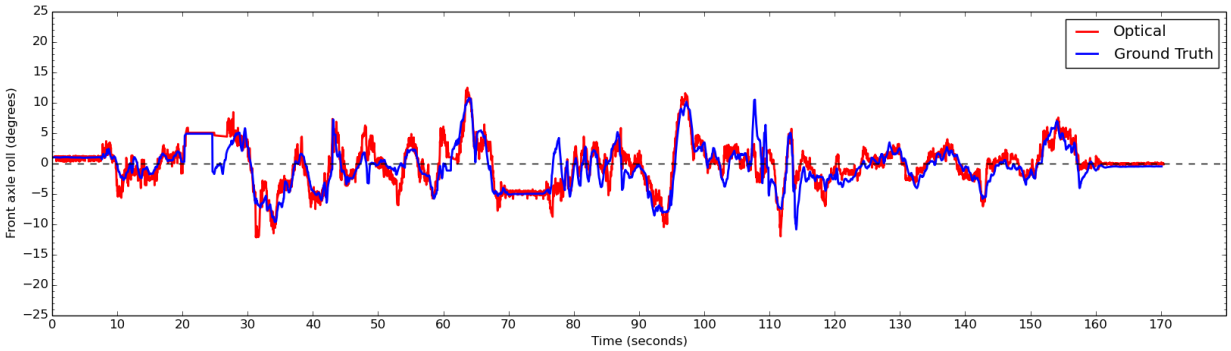

(a) Front Axle

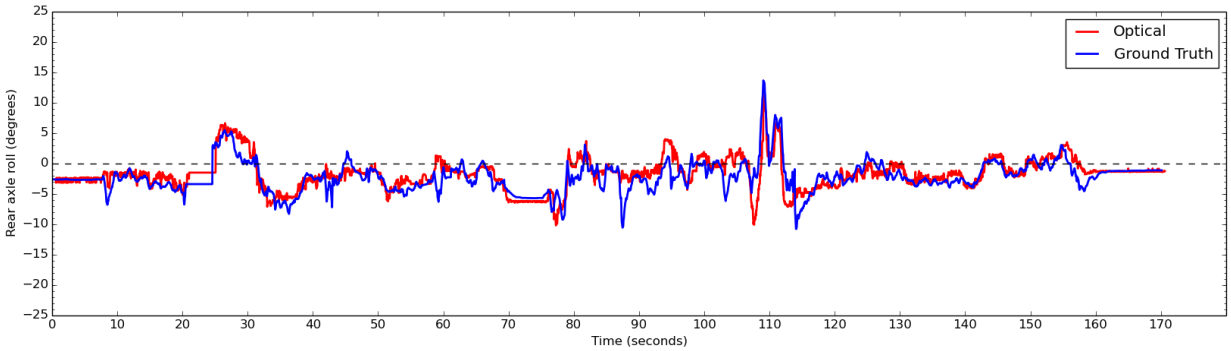

(b) Rear Axle

Figure 9: Estimated axle roll vs. Ground truth IMU data

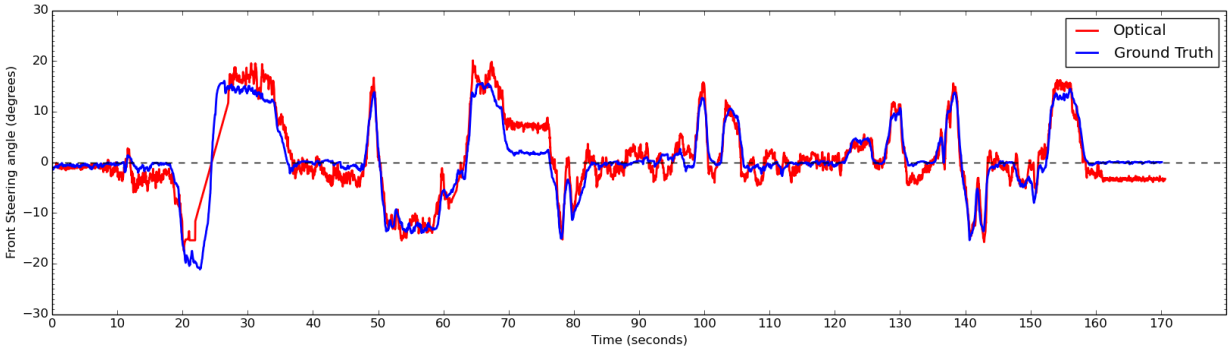

(a) Front Axle

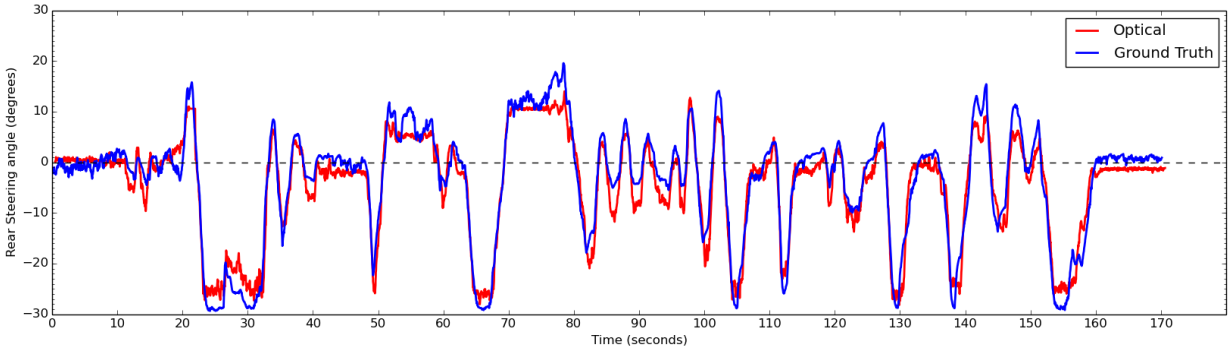

(b) Rear Axle

Figure 10: Estimated steering angle vs. Ground truth potentiometer data

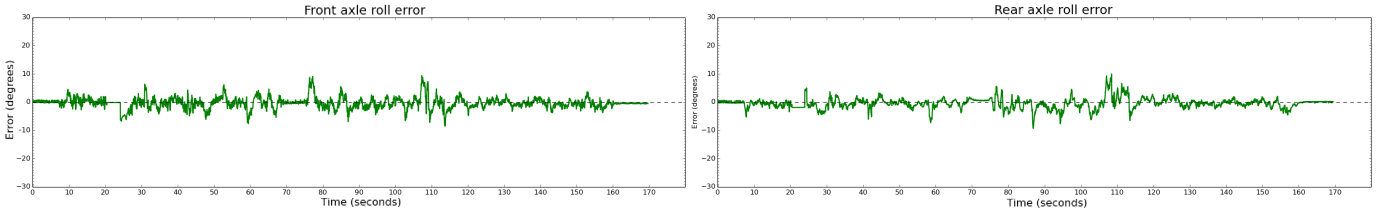

(a) Front Axle **(b)** Rear Axle

Figure 11: Estimated axle roll error with respect to Ground truth

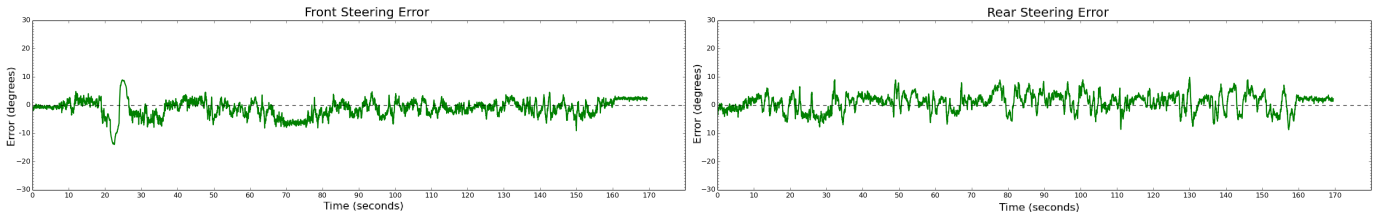

(a) Front Axle **(b)** Rear Axle

Figure 12: Estimated axle roll error with respect to Ground truth

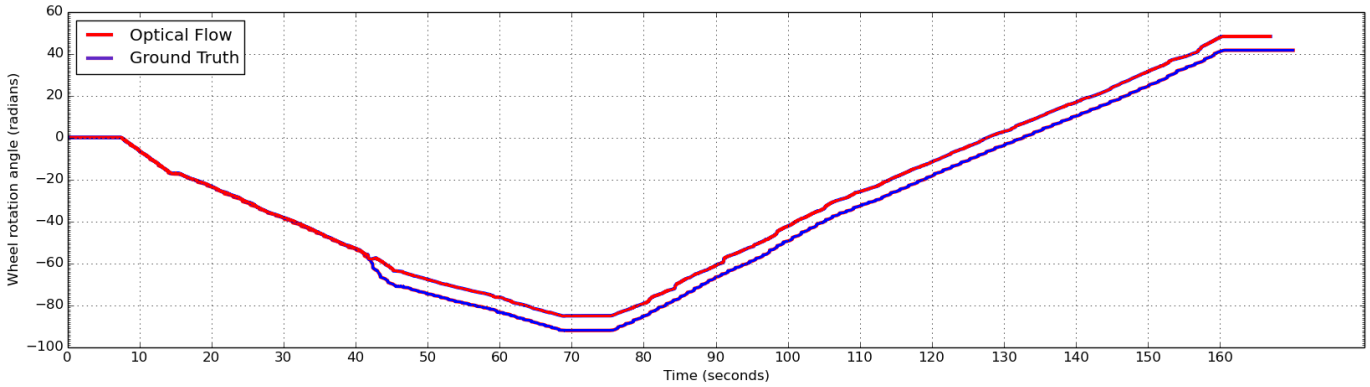

Figure 13: Estimated wheel rotation angle vs. Ground truth motor encoder data

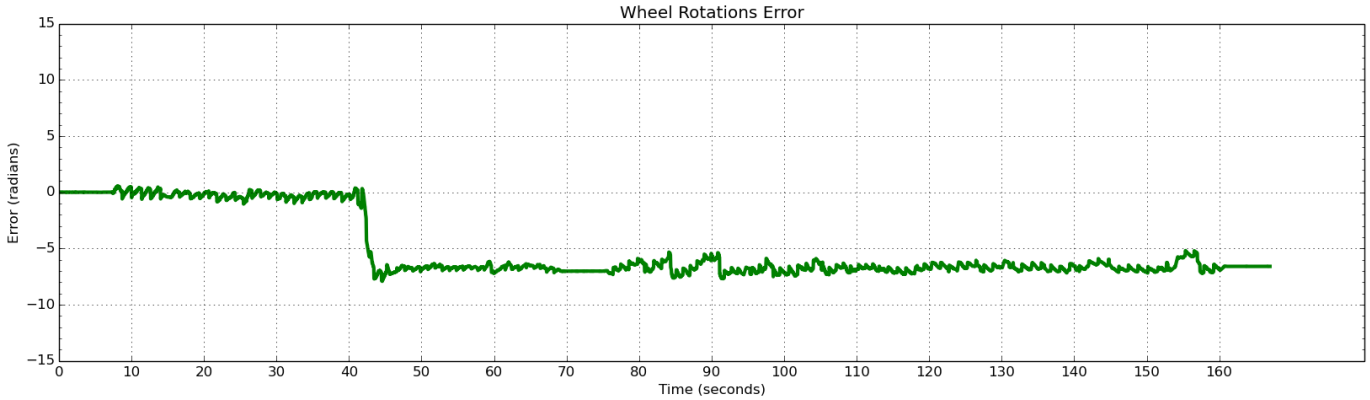

Figure 14: Drift in Estimated wheel rotation with respect to Ground truth data

Dataset	Duration (secs)	Axle Roll MAE		Steering Angle MAE		Maximum wheel rotation drift (rad)
		Front	Rear	Front	Rear	
Dataset 1	82	1.65°	1.51°	2.49°	1.98°	0.86
Dataset 2	170	1.41°	1.32°	2.40°	2.49°	3.70
Dataset 3	257	2.41°	1.97°	1.86°	2.55°	1.78

TABLE I: Mean Absolute Error (MAE) and Wheel rotation drift computed for three independent datasets

Ground truth data was recorded via three inertial measurement units for axle roll, two steering potentiometers for steering angle and a motor encoder for wheel angular rotation. The image sequences of three independent datasets, each traversing the span of the test-site, were used as inputs to estimate the 10 degrees-of-freedom.

Figures 9 to 14 are all plotted using **dataset 2**. Figure 9 and 10 illustrate the variations in ψ_F, ψ_R and λ_F, λ_R over time. The difference between the ground truth data and the estimated data is represented by Figure 11 and 12. The 2 methods correspond significantly well over all datasets, as seen in Table I, which shows the method's mean absolute error. Figure 13 plots the average of θ_{RL} and θ_{RR} compared against the rotations output from the rear motor encoder. Figure 14 shows the estimation error of the aforementioned graph.

The results are found to be highly satisfactory, and convey overall agreement between optical and conventional methods. Nevertheless, certain sources of error have been identified as below-

(i) Drop in frame rate results in the loss of continuous data. Frame rate is critical in the case of optical flow, where features are lost.
(ii) Excessive camera mount vibration results in noisy output.
(iii) Loss of marker positions due to self-shadowing causes gaps in kinematic state data.
(iv) Self-shadowing on wheels alter direction of optical flow vectors, giving erroneous angular rotation.

VII. Conclusion and Future Work

This paper introduces a novel approach to kinematic state estimation, and the results demonstrate high confidence in a single camera system. The method successfully precludes nine proprioceptive sensors, to achieve comparable kinematic state estimation with a single camera. As rovers minimalize and as reliability call for redundancy, state estimation via vision will be incorporated into planetary rovers. This is well-supported by the predicted increase in computational power in space robotics.

The method sets a precedent and provides a foundation for further work in vision-based kinematic state estimation. Future work would include developing an optimized version of the algorithm, capable of running on-board a planetary rover. Self-shadowing and abrupt changes in lighting affect the system adversely, and robustness to self-shadowing [12] and illumination-invariant tracking are rewarding avenues for further research. In addition, kinematic state estimation has no theoretical dependence on fiducial markers, and future work could focus on tracking inherent features of the robot. A downward-facing fisheye camera can also perform visual odometry, that tracks terrain features to give motion estimate. Future work will have visual odometry and optical kinematic estimation in conjunction to achieve unprecedented odometry with utmost simplicity. With two independent representations of motion, rover slip can also be detected.

Acknowledgments

The authors would like to thank the Robotics Institute Summer Scholars Program, Field Robotics Center and Carnegie Mellon for enabling and facilitating the above research.

References

[1] Richard Volpe. Navigation results from desert field tests of the rocky 7 mars rover prototype. *The International Journal of Robotics Research*, 18(7):669–683, 1999. doi: 10.1177/02783649922066493. URL http://ijr.sagepub.com/content/18/7/669.abstract.

[2] Michael Dille, Benjamin P Grocholsky, and Sanjiv Singh. Outdoor downward-facing optical flow odometry with commodity sensors. In *Proceedings Field & Service Robotics (FSR '09)*, July 2009.

[3] M.J. Roman. *Design and Analysis of a Four Wheeled Planetary Rover*. University of Oklahoma, 2005. URL https://books.google.com/books?id=d-hSNwAACAAJ.

[4] Ellery A. *Planetary Rovers: Robotic Exploration of the Solar System*. Springer Berlin Heidelberg, 2016. URL http://link.springer.com/book/10.1007%2F978-3-642-03259-2.

[5] Larry Matthies, Mark Maimone, Andrew Johnson, Yang Cheng, Reg Willson, Carlos Villalpando, Steve Goldberg, Andres Huertas, Andrew Stein, and Anelia Angelova. Computer vision on mars. *International Journal of Computer Vision*, 2007.

[6] T. M. Howard, A. Morfopoulos, J. Morrison, Y. Kuwata, C. Villalpando, L. Matthies, and M. McHenry. Enabling continuous planetary rover navigation through fpga stereo and visual odometry. In *Aerospace Conference, 2012 IEEE*, pages 1–9, March 2012. doi: 10.1109/AERO.2012.6187041.

[7] J. Balaram. Kinematic observers for articulated rovers. In *Robotics and Automation, 2000. Proceedings. ICRA '00. IEEE International Conference on*, volume 3, pages 2597–2604 vol.3, 2000. doi: 10.1109/ROBOT.2000.846419.

[8] M. Tarokh, G. McDermott, S. Hayati, and J. Hung. Kinematic modeling of a high mobility mars rover. In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, volume 2, pages 992–998 vol.2, 1999. doi: 10.1109/ROBOT.1999.772441.

[9] Pierre Lamon and Roland Siegwart. 3d position tracking in challenging terrain. *The International Journal of*

Robotics Research, 26(2):167–186, 2007. doi: 10.1177/0278364906075170. URL http://ijr.sagepub.com/content/26/2/167.abstract.

[10] J. Sturm, C. Plagemann, and W. Burgard. Unsupervised body scheme learning through self-perception. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 3328–3333, May 2008. doi: 10.1109/ROBOT.2008.4543718.

[11] Bruce D Lucas, Takeo Kanade, et al. An iterative image registration technique with an application to stereo vision. In *IJCAI*, volume 81, pages 674–679, 1981.

[12] N. Seegmiller and D. Wettergreen. Optical flow odometry with robustness to self-shadowing. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 613–618, Sept 2011. doi: 10.1109/IROS.2011.6094670.

[13] Eugene Fang . High-fidelity planetary route determination using computationally efficient monocular fisheye odometry and sun compass. Technical Report CMU-RI-TR-16-14, Robotics Institute, Pittsburgh, PA, May 2016.

[14] P. Backes, A. Diaz-Calderon, M. Robinson, M. Bajracharya, and D. Helmick. Automated rover positioning and instrument placement. In *2005 IEEE Aerospace Conference*, pages 60–71, March 2005. doi: 10.1109/AERO.2005.1559299.

[15] D. Scaramuzza, A. Martinelli, and R. Siegwart. A flexible technique for accurate omnidirectional camera calibration and structure from motion. In *Fourth IEEE International Conference on Computer Vision Systems (ICVS'06)*, pages 45–45, Jan 2006. doi: 10.1109/ICVS.2006.3.

[16] Davide Scaramuzza, Agostino Martinelli, and Roland Siegwart. A toolbox for easily calibrating omnidirectional cameras. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5695–5701. IEEE, 2006.

[17] Jianbo Shi and C. Tomasi. Good features to track. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*, pages 593–600, Jun 1994. doi: 10.1109/CVPR.1994.323794.

[18] Jean-Yves Bouguet. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel Corporation*, 5(1-10):4, 2001.