Visualizing and Analyzing Machine-soil Interactions using Computer Vision

• •

Krzysztof Skonieczny and Scott J. Moreland
Field Robotics Center, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, Pennsylvania 15213
e-mail: kskoniec@cmu.edu
Vivake M. Asnani and Colin M. Creager
NASA Glenn Research Center, Mail Stop 23-3, 21000 Brookpark Road, Cleveland, Ohio 44135
Hiroaki Inotsume and David S. Wettergreen
Field Robotics Center, Carnegie Mellon University, 5000 Forbes Avenue Pittsburgh, Pennsylvania 15213

This work presents an experimental method for visualizing and analyzing machine-soil interactions, namely the soil optical flow technique (SOFT). SOFT uses optical flow and clustering techniques to process images of soil interacting with a wheel or tool from photos taken through a glass wall of a soil bin. It produces results that are far richer than past approaches that utilized long-exposure photography. It achieves a performance comparable to particle image velocimetry or particle tracking velocimetry, but without the need for specialized measurement equipment or specially marked soil particles. The processing technique demonstrates robustness to different soil types. Ground-truth and cross-validation experiments exhibit subpixel accuracy in estimating soil motions. An example of an application of this technique for field robotics research is the detailed study of push-rolling for slope climbing and soft soil traverse. Push-rolling advances a vehicle by rolling a subset of its wheels while changing its wheelbase to keep the other wheels static and pushing against the ground. Experiments show that push-rolling achieves higher net traction than conventional rolling. Observing the two aspects of push-rolling (rolling and horizontal pushing) using SOFT shows that they result in entirely different forms of soil shearing ("grip failure" and "ground failure," respectively). SOFT also demonstrates how the direction of soil motion is more efficiently utilized for horizontal thrust by pushing than conventional rolling. Ongoing work utilizing SOFT has also demonstrated its potential use in studying excavation tool interactions, the effects of grousers on wheel efficiency, as well as a variety of other wheel-soil interactions. © 2014 Wiley Periodicals, Inc.

1. INTRODUCTION

Field robots, as well as other off-road vehicles and machines, are often required to traverse or work in granular soil. Wheels are commonly used in off-road applications that span construction, mining, agriculture, recreation, military operations, and planetary exploration. The wheels' interaction with the terrain governs performance; traction, slip, sinkage, and soil transport directly influence a machine's ability to climb slopes, traverse soft ground, and push tools through the soil.

The effect of wheel design parameters (geometry, surface features, compliance, control modes) on soil interaction (soil flow, thrust, and resistance) and overall performance (trafficability, efficiency) are not yet understood in sufficient detail, as evidenced by continued activity in the field of terramechanics. The goal of this work was to create a system that enabled visual analysis of the wheel-soil interactions, allowing the effects of design changes to be observed qualitatively and quantitatively. While the focus of this work was on evaluating wheel-soil interactions, the methods could

also be used to study other machine-soil interactions, such as those of tracks and excavation tools.

Observing subsurface soil motion near the machine-soil interface provides critical information about how soil is interacting with wheels and tools that cannot be gleaned from observing the soil surface. Normally, it is not possible to look through soil at subsurface particles. However, experiments with implements up against a glass-sided soil bin can provide insight into specific modes of soil shearing, the direction of soil flow (and the existence of forward flow regions, for example), and other phenomena.

Direct observation of soil motion through glass sidewalls has been utilized in soil mechanics and terramechanics research for over half a century (Bekker, 1948). A half-width test implement is placed up against the glass, and correspondingly half the payload weight is applied. When this half-width implement is driven along the glass wall, it simulates a full implement in unconstrained conditions with the same ground pressure. Wong concluded experimentally that as long as shear stress between glass and soil

is negligible, the glass surface acts as a plane of symmetry and the soil behaves as it would directly below an implement twice as wide (Wong, 1967; Wong & Reese, 1966).

One archetypical photographic method for observing soil motion uses long-exposure photos and distinguishes sharp and streaking soil grains as stationary and moving, respectively. Streaks in the photos also provide information about the directions of soil motion. This is the technique that was used by Bekker, Wong, and others, including Harrison (1973).

More recently, researchers in the fields of soil mechanics and geotechnical studies have applied imaging techniques adapted from fluid flow experiments, including particle image velocimetry (PIV) and particle tracking velocimetry (PTV) (Gnanamanickam et al., 2007; Murthy et al., 2009; White et al., 2003). These techniques, however, make use of specialized equipment such as high-speed cameras, pulsed lasers, or multiphase LED lamps (Cowen & Morismith, 1997; Gnanamanickam et al., 2007), and/or they require reference markers or specially colored sand (Gachet et al., 2003; White et al., 2003).

In addition to imposing stringent equipment and soil preparation requirements, prior techniques have also had to sacrifice data resolution and accuracy. PIV techniques divide images into a grid, and they estimate velocity just once per grid square (Gachet et al., 2003; White et al., 2003). With typical grid sizes on the order of 8×8 to 30×30 pixels, this results in approximately 1,000–3,000 data points per image. Optical flow techniques presented here calculate a data point at each pixel, resulting in closer to 1,000,000 data points per image. In PIV, there is also a tradeoff between resolution and accuracy. White achieves 95% error bounds of 0.02 pixels only if restricting measurements to large grids and thus fewer than 250 data points; motion at any of the pixels not at the center of a grid can at best be interpolated. To reach 2,500 data points with a finer grid spacing, the 95% error bound grows beyond 1 pixel and indeed is larger than the measured motion signal (White et al., 2003). The optical flow technique presented in this work achieves low subpixel errors (fewer than 0.3 pixels) while calculating a data point at every pixel, as will be shown in Section 3.

The objective of this work was to develop and evaluate an experimental technique that produces rich soil flow data using just a digital camera. A technique is sought that would not require special soil preparation, and that can accomplish the required processing using readily available software algorithms. This would make the study of machine-soil interactions both more data-rich than ever before and more widely accessible to researchers.

The work presents the soil optical flow technique (SOFT). This analytic technique is based on optical flow, a dense motion estimation algorithm (i.e., calculated at each pixel) developed in computer vision. Murthy did a preliminary study of applying optical flow to displaced soil, but it still required initial (sparse) estimates of the motion

collected using a high-speed camera and multiphase LED lamps (Gnanamanickam et al., 2007; Murthy et al., 2009). The optical flow algorithm applied here is robust enough to be used alone.

Previous work by the authors has already demonstrated preliminary results of SOFT's utility. It has been used to study how soil failure under wheels is modified when implementing a novel inching locomotion mode (Moreland et al., 2011). It has also been used to study the effects of grousers on wheel traction (Moreland et al., 2012; Skonieczny et al., 2012). This work develops SOFT in detail and characterizes its performance quantitatively.

The hardware and software implementation of SOFT is detailed in Section 2. Section 3 presents error analyses that quantify the performance of this method. Section 4 demonstrates applications of the technique to the design and operation of field robots. Section 5 reviews conclusions, and outlines ongoing and future work.

2. DESCRIPTION OF THE SOIL OPTICAL FLOW TECHNIQUE

This section describes both the software and hardware development of SOFT.

2.1. Hardware Setup

The experimental system, shown in Figure 1, drives a single wheel though granular terrain and allows for the soil particle reactions to be observed. The system consists of a glass-walled soil bin, a motor-driven wheel, a horizontally driven carriage, and a digital camera. Wheel rotation is velocity-controlled in coordination with the carriage travel to create a controlled slip rate between the wheel and the terrain. Note that slip rate is defined as

$$\text{Slip rate} = 1 - \frac{v_{\text{carriage}}}{r\omega} \qquad (1)$$

where v_{carriage} is the horizontal speed of the carriage, r is the wheel radius, and ω is the wheel's angular velocity. The wheel is connected to a vertical rail system, so its vertical travel is unconstrained. This allows for deadweight payload to be applied, and for the wheel to sink naturally into the terrain. The system may be alternatively configured to test other implements, such as tracks or excavation tools.

The motions of the test implement and soil are observed through a sheet of high hardness tempered glass that extends the depth of the soil bin. A half-width test implement is placed against the glass without normal pressure (no horizontal preload was applied to the wheel against the glass) to keep the interface friction low. Lack of scratching or marking suggests negligible friction. If lower shearing resistance is ever required, the inside of the glass can be lubricated with silicone wax (Gachet et al., 2003; Witney, 1968).

A digital camera with a 50 mm macrolens is used to image the soil in the region where it interfaces with the
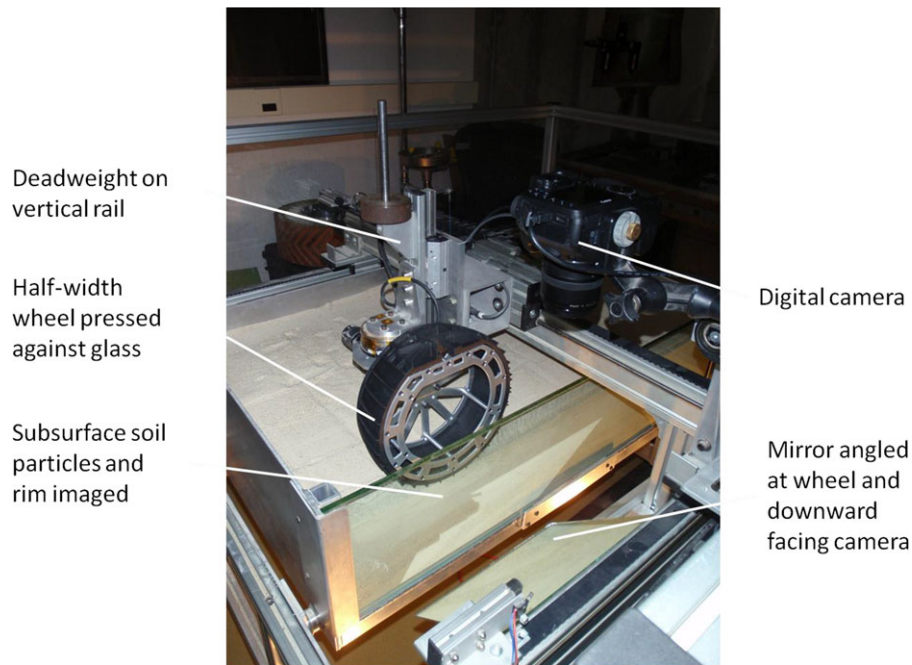
Figure 1. Experimental system for studying wheel-soil interactions using SOFT.

Table I. Soil parameters for GRC-1 and Fillite.

Soil type	GRC-1	Fillite
Particle size range	0.05–2 mm	0.005–0.5 mm
Median particle size	0.3 mm	0.1–0.2 mm
Bulk density	1,600–1,900 kg/m^3	350–450 kg/m^3

test implement. The camera travels with the implement in the horizontal direction as the carriage moves. The camera is pointing downward and views the region of interest through an angled mirror. The mirror arrangement is necessary for this particular setup, because there is not enough space for the camera to be mounted with a direct view. Two external halogen flood lights are placed approximately 1 m from the soil bin to provide constant illumination, high contrast, and reduced shadows. The lights are pointed at a high angle from the normal of the glass to avoid direct reflection into the camera.

Soil motion tests have been conducted on two distinct types of homogeneous soil to date, without needing to modify the technique or computations used. These soils include GRC-1, a relatively coarse lunar soil simulant (Oravec et al., 2010), and Fillite, commercial low-density aluminosilicate microspheres. Particle sizes and densities for GRC-1 and Fillite are given in Table I. Raw images from tests with these soils are shown in Figure 2.

Table II shows typical system parameters used for wheel testing.

2.2. Overview of Processing Technique

An overview of the SOFT processing technique is presented in Figure 3. The images from the test sequence are converted from color to black and white, and each sequential image pair is processed. Optical flow computes displacements between the image pairs at each pixel, in units of pixels; this step is described in detail in Section 2.3. Next, data are transformed to an inertial reference frame by subtracting the horizontal motion of the camera from the computed horizontal flow. Considering that most of the soil is static relative to the experimental apparatus, the median horizontal component of velocity of the soil moving through the camera's field of view is used to estimate the camera motion.

The displacements at each pixel, now in an inertial reference frame, are translated into polar coordinates (magnitude and direction) and are multiplied by frame rate and spatial pixel size to convert into soil velocity. A frame rate of 8 fps and a pixel size of 0.2 mm are typical settings for wheel tests (see Table I).

Clustering is used to distinguish the portion of the image that contains soil, as described in more detail in Section 2.4. Velocity magnitude and direction are graphically displayed over the top of the soil portion of the image, as shown in Figure 4. Further clustering on the soil velocity magnitude identifies the region of soil where significant motion is
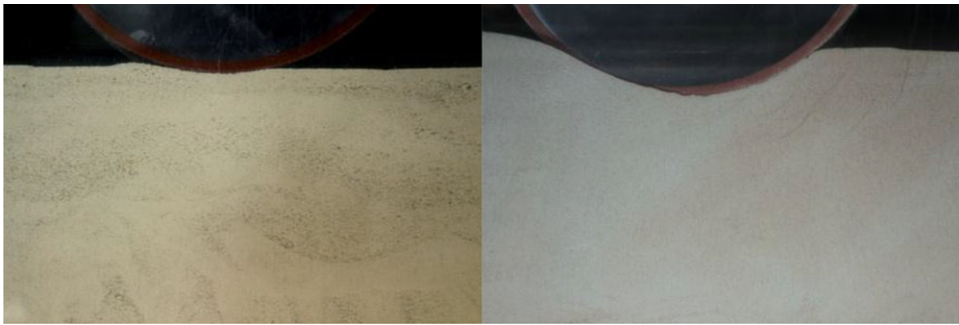
Figure 2. Soil motion tests set up in GRC-1 (left) and Fillite (right). The imaging technique has proven to be robust to changes in soil type, and it does not require specific preparation of the soil (with colored markers, etc.).

Table II. Typical system parameters for wheel testing.

Carriage speed	Up to 1 cm/s
Wheel radius	11.4 cm
Wheel angular velocity	Up to 0.087 rad/s
Camera acquisition rate	Up to eight frames per second (fps)
Field of view	12 cm × 24 cm
Resolution	600 × 1,200 pixels (pixel size = 0.2 mm)
Soil particle size	GRC-1: 0.05–2 mm
	Fillite: 0.005–0.5 mm

observed. An optional feature of SOFT is the demarcation of a shear interface around this region (i.e., between moving and static soil). Another optional feature is the filtering out of soil flow direction outside the region with significant

motion, as direction has little meaning when the velocity magnitude is near zero.

2.3. Optical Flow

Soil motion between subsequent images is estimated using optical flow. Optical flow is a dense motion estimation technique, meaning motion is calculated for each pixel of the image. The technique outputs two displacement fields, corresponding to horizontal and vertical motion, in units of pixels.

Optical flow solves for the displacement fields by attempting to satisfy constraints stemming from two assumptions: data conservation and spatial coherence. Data conservation assumes that each point visible in an image will look the same when present in another part of a subsequent image, while spatial coherence assumes that neighboring points can only move relative to one another according to some constraints. As these assumptions will not be strictly
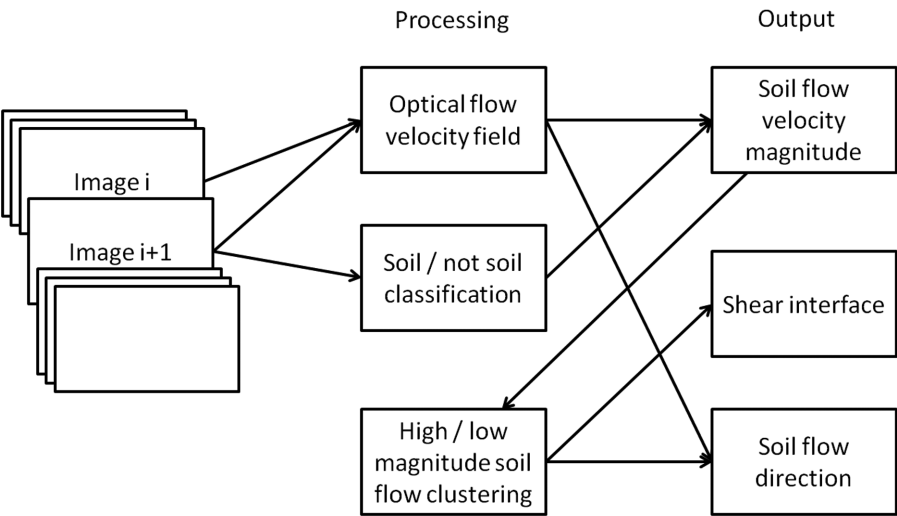
Figure 3. Image processing flowchart showing how optical flow and clustering lead to soil flow and shear interface output.
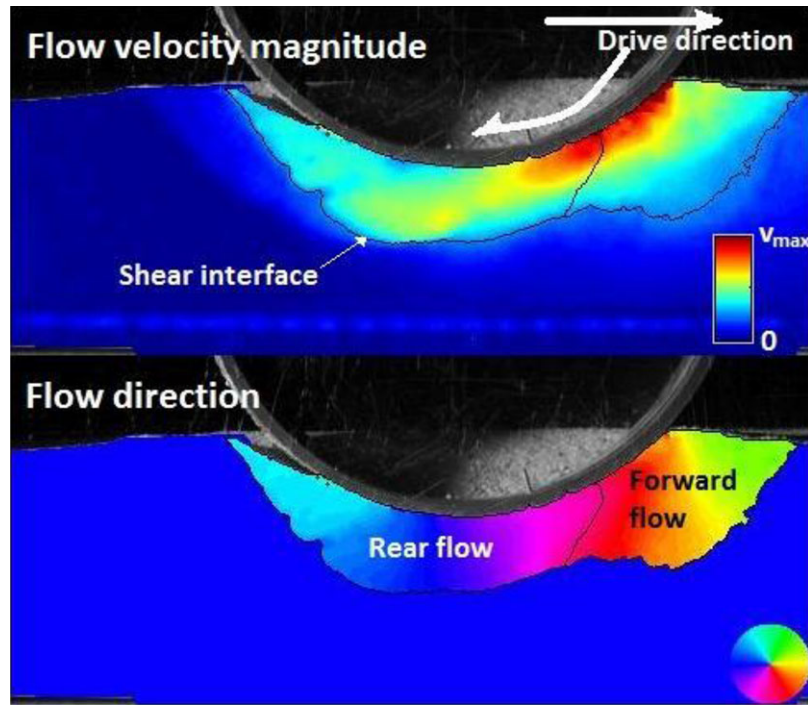
Figure 4. Sample processed output for a driven wheel experiment. Soil flow velocity magnitude (upper) is denoted from blue (static) to red (maximum). The shear interface defines the boundary between moving and stationary soil. Soil flow direction (lower) within the shear interface is denoted according to the color wheel on the bottom right.

satisfied for any general motion, optical flow minimizes an objective function that assigns weights to deviations from the different constraints.

The mathematical framework of optical flow described here is a summary of detailed formulations presented by Szeliski (2011) and Black and Anandan (1996) that were developed for general application. The cited references are valuable sources for investigating details beyond those included here.

The data conservation assumption underpinning optical flow states that the intensity of each point (each soil particle, for example) in a black and white image stays constant. This implies that

$$I(x + u(x, y), y + v(x, y), t) = I(x, y, t - 1) \quad (2)$$

where I is the intensity at pixel location (x, y) and time step t, and u and v denote the displacements in the x and y directions, respectively, at each pixel. What Eq. (2) states is that the intensity of the point at location (x, y) in the image indexed $t - 1$, when displaced by u and v, will look the same at its new location in the subsequent image. Assuming u and v are small [classically less than 1 pixel, though this can be increased to motions on the order of 10 pixels through course-to-fine Gaussian smoothing (Anandan, 1989)], the left-hand side of Eq. (2) can be replaced with its Taylor

series expansion, leading to

$$(\nabla I)^{\mathrm{T}} \mathbf{u} + I_t = 0, \quad (3)$$

where (∇I) is the gradient of the image intensity, $\mathbf{u} = [u, v]^{\mathrm{T}}$, and $I_t = I(x, y, t) - I(x, y, t - 1)$. Equation (3) is underconstrained (two unknowns with only one equation), so another constraint is typically introduced in the form of spatial coherence. Affine flow is assumed for the spatial coherence in this implementation, which allows for the region around each pixel to rotate, shear, as well as change aspect ratio and scale. This is appropriate for estimating soil motion, as soil particles naturally exhibit such types of motion relative to one another. The affine flow spatial constraint is written as

$$\mathbf{u}(x, y; \mathbf{a}) = \begin{bmatrix} u(x, y) \\ v(x, y) \end{bmatrix} = \begin{bmatrix} a_1 + a_2 x + a_3 y \\ a_4 + a_5 x + a_6 y \end{bmatrix} \quad (4)$$

where a_i are dimensionless coefficients (that must also be optimized).

The defined set of constraints in Eqs. (3) and (4), taken together, cannot be strictly satisfied for nontrivial motions, but they provide a starting point for setting up an error minimization to estimate $\mathbf{u}$. The exact form of this objective function is very cumbersome due to complexities of optimizing all the affine flow parameters; the interested reader

is encouraged to consult Black and Anandan's work directly (Black & Anandan, 1996).

Minimization is performed in several steps that fall into two higher-level stages: quadratic and nonconvex minimization. First, a quadratic error norm is used (note that a quadratic error norm does not require any parametrization); the minimization is thus a least-squares calculation. To avoid local minima while accounting for potentially large motions, the minimization is done at four different levels of scale (achieved by subjecting the image to varying degrees of Gaussian smoothing) in a course-to-fine refinement.

Once a least-squares estimate of the motion has been found, the solution is further refined with graduated non-convexity. This consists of two additional stages of minimization, now using a (nonconvex) robust error norm, such as a Lorentzian:

$$\rho(z; \sigma) = \log\left(1 + \frac{1}{2}\left(\frac{z}{\sigma}\right)^2\right) \tag{5}$$

where z is replaced by a constraint equation such as Eq. (3). A Lorentzian norm is less sensitive to outliers than a quadratic norm, which can lead to a better estimate of a pixel's dominant motion when there are multiple motions present in its vicinity. Each of the two successive stages of minimization with a Lorentzian reduces sensitivity to outliers, the second time by reducing the width parameter, σ. Both stages of minimization with the Lorentzian also calculate a two-scale course-to-fine refinement.

Black and Anandan's optical flow technique (including the key aspects described in this subsection) has been implemented by Sun (Sun et al., 2008), and is available open-source on the web.

2.3. Clustering

A technique called k-means clustering is used twice in the data processing: to find the dividing line between soil and not-soil as well as to locate the line between stationary and moving soil masses.

The k-means clustering algorithm is applied to a set of data to classify them into k categories. It is initialized by randomly selecting k data points and setting them as the centroids of the categories. Each data point is then assigned to whichever category has the closer centroid. The calculation iterates, alternating between category assignment and centroid update (i.e., recalculating the centroid for each newly populated category) steps, until convergence criteria are satisfied.

Soil versus not-soil clustering: The region of the image containing soil is distinguished from the rest of the image based on pixel intensity. In the work presented here, the test implements and background were generally darker than the soil being studied. To study the flow of particularly dark soil, a light-colored test implement and background could conversely be used. After initial clustering of pixel intensities using k-means (with $k = 2$), any gaps in the (single largest) "soil" region are filled (as the soil is assumed to remain a single connected mass). This is achieved with alternating binary dilation and erosion operations [for more information, see Gonzalez et al. (2009)].

High/low magnitude soil flow clustering: With soil motion expressed in the inertial frame, soil velocity magnitudes at each pixel are clustered into "stationary/low magnitude flow" and "high magnitude flow." After initial clustering, gaps within the single largest region (of high velocity magnitude) are filled. The boundary surrounding this high velocity magnitude region is interpreted as a soil-shear interface.

3. QUANTITATIVE EVALUATION OF OUTPUT

This section compares measured and processed soil flow to ground-truth motions as well as to output from another processing technique. The analyses demonstrate the technique's robustness and quantify its limitations.

3.1. Ground-truth Testing

Ground-truth testing was conducted using both GRC-1 and Fillite. See Section 2.1 for descriptions of both of these soils.

3.1.1. Controlled Rotation Tests

To quantify the error of soil flow output, soil was spread out flat covering a plate (Figure 5), which was then rotated by a very precise amount: 0.02944 ± 0.00003 rad.

The ground truth of this prescribed motion was compared to the soil flow output by processing pre- and post-motion photographs, as shown in Figure 6. The largest displacement, at the corners of the image, is 15.6 pixels, or 3.12 mm.

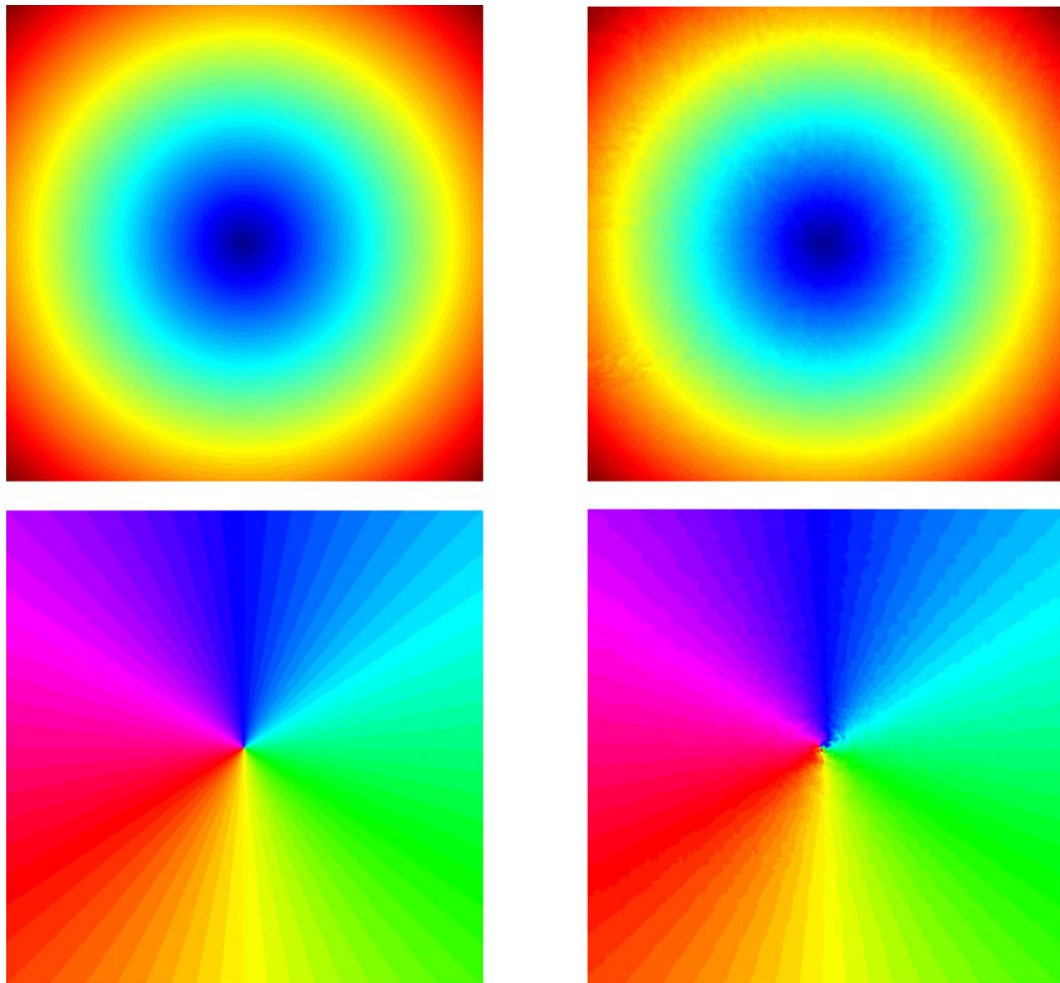Motion error at each pixel is measured by the Euclidean norm:

$$\text{Error} = \sqrt{(u - u_{ref})^2 + (v - v_{ref})^2} \tag{6}$$

This error is shown in Figure 7 for GRC-1 and Fillite, according to its distribution in the image. Figures 8 and 9 show the statistical distributions of the error magnitude for GRC-1 and Fillite, respectively. Errors are sub-pixel in magnitude throughout the image, with more of the higher magnitude errors observed near the edges where motion magnitude is also larger (as high as 15.6 pixels). Regression shows a weak correlation of error magnitude with motion magnitude, with a 0.01 pixel increase in error for every 1 pixel of additional motion and an R^2 of 0.25 for this correlation.

Figures 8 and 9 show confidence bounds based on multiple measurements output by SOFT. It shows that 95% of the errors in the ground-truth tests are less than or equal to 0.3 pixels. This means, for example, that motions of three pixels between subsequent images have a signal-to-noise
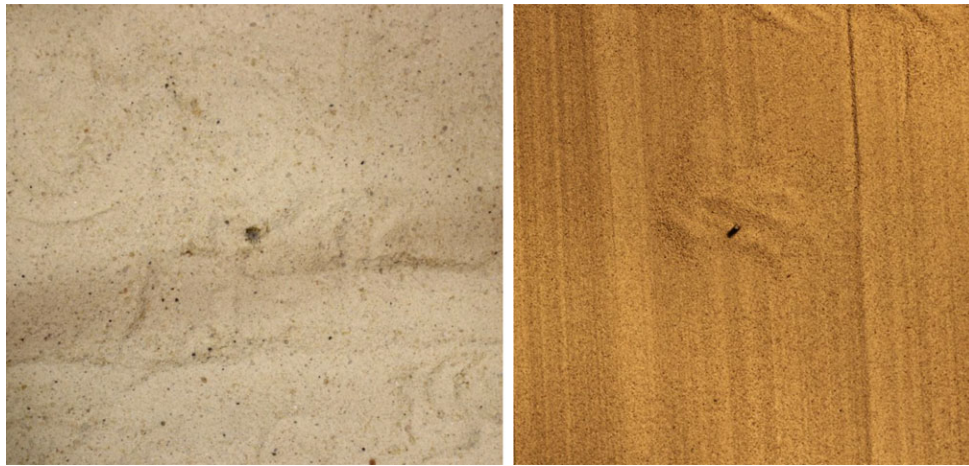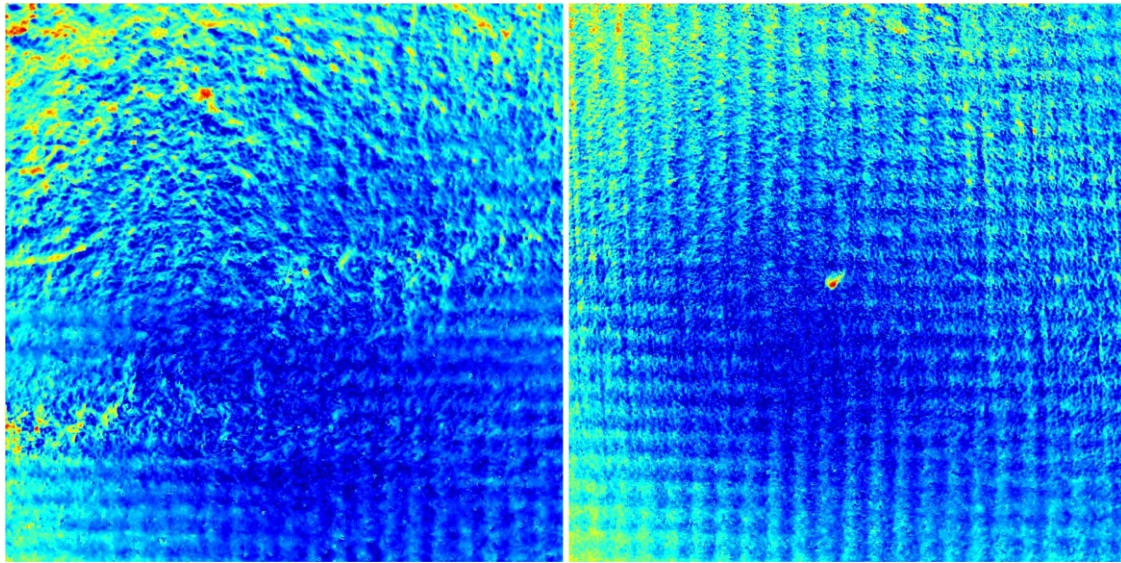
Figure 5. Test soils spread flat on a plate (GRC-1 left, Fillite right), viewed from directly above. A sharpened pin at the center marks the point of rotation for ground-truth testing.
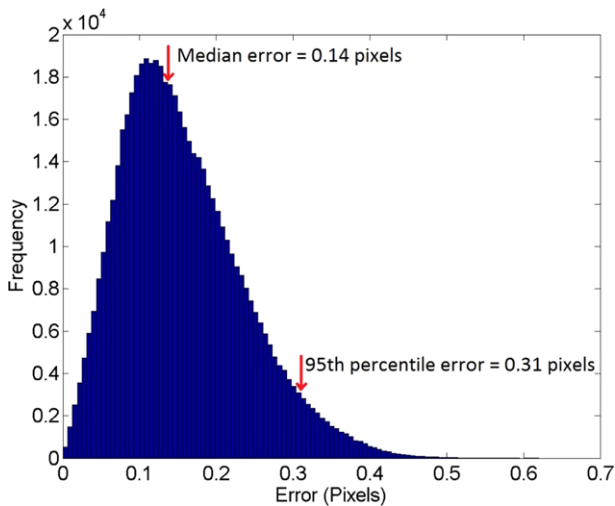

Figure 6. Ground truth (left) and measured (right) soil displacement magnitude (top) and direction (bottom), induced by a prescribed 0.02944 rad rotation (GRC-1 results shown).
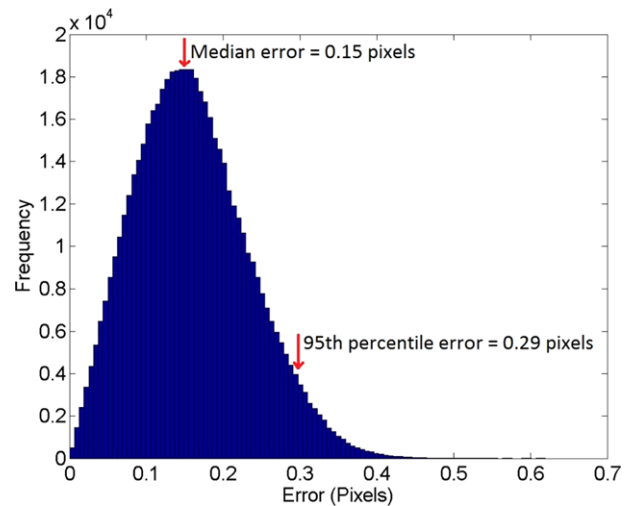
Figure 7. Motion error at each pixel of the image (GRC-1 left; Fillite right). Warmer colors indicate higher error, and are more prevalent at the edges where motions are also large.

Figure 8. Distribution of motion error by magnitude for GRC-1.

Figure 9. Distribution of motion error by magnitude for Fillite.

(S/N) ratio of at least 10 with 95% confidence, and, based on the median error of 0.15 pixels, a S/N of at least 20 with 50% confidence.

3.1.2. Constant Horizontal Speed Tests

Typical soil flow tests involve moving the test implement and camera horizontally at a constant speed (and, in the case of wheel tests, coordinated rotation of the wheel for constant slip, as described in Section 2.1). The majority of the soil captured in the image is not influenced by the test implement, and it remains static (see the large region of dark blue in Figure 10, for example). This static soil should thus move horizontally through the passing camera's field of view, at the speed at which the carriage is commanded (as discussed in Section 2.2). The velocity of the static soil in the camera's reference frame may deviate slightly from the commanded travel velocity due to error in the processing technique, error in the control of the horizontal motion, vibrations in the structure between the test implement and the camera, and/or inconsistencies in the camera's frame rate.
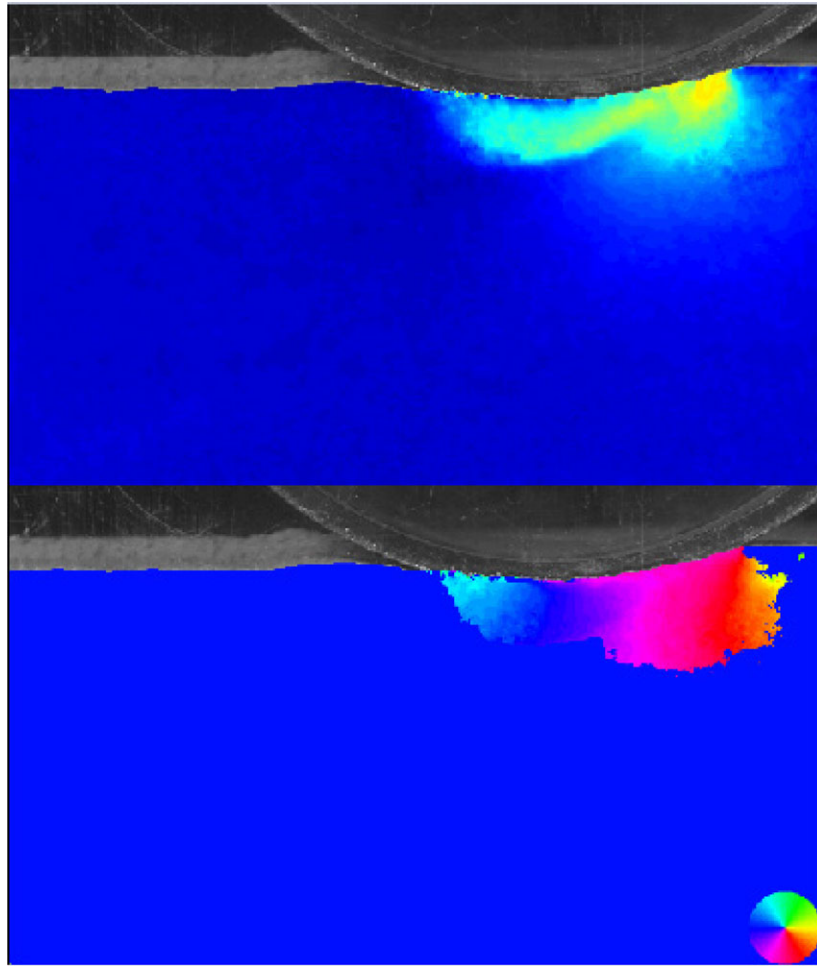
Figure 10. Sample output from a rolling wheel test at constant commanded horizontal speed (constant slip). The region colored dark blue in the velocity magnitude plot (top), covering the majority of the image, represents static soil unaffected by the test implement.

Figure 10 shows a sample of processed output from a rolling wheel test in GRC-1, commanded to a constant 20% slip rate. Results from a single frame pair are shown; the full test consists of 51 frames (and thus 50 frame pairs). The wheel is commanded at a tangential speed of 1 cm/s, and the carriage is advanced at 0.8 cm/s to maintain the desired 20% slip test condition. From the median horizontal soil flow of each of the 50 frame pairs of the test, the average horizontal velocity was estimated to be 0.79 cm/s with a 95% confidence range of ± 0.04 cm/s. This corresponds to 7.23 ± 0.34 pixels/frame. The variation in horizontal motion detected is thus similar to what was observed in the benchmark plate rotation test for GRC-1 (± 0.31 pixels/frame, as shown in Figure 8). This suggests that the additional potential sources of error are no more significant than the already small error in the optical flow itself.

3.2. Cross-validation with Feature Tracking

Ground-truth measurements of soil displacement can be made for simple motions, such as the rotation and horizontal flow described in Section 3.1, but they are not easy to obtain or enforce for more complex flows. When validation against ground truth is not feasible, it is useful to at least cross-validate against other techniques.

To gain insight into how SOFT handles complex soil flows, its output is compared to corresponding output calculated via scale-invariant feature tracking (SIFT) (Lowe, 1999). SIFT searches for robustly distinguishable features in an image, and it can then match those features in a subsequent image. Soil displacement can thus be calculated at key points between an image pair.

Cross-validation for optical flow and SIFT output was performed with a rolling wheel experiment. The wheel was
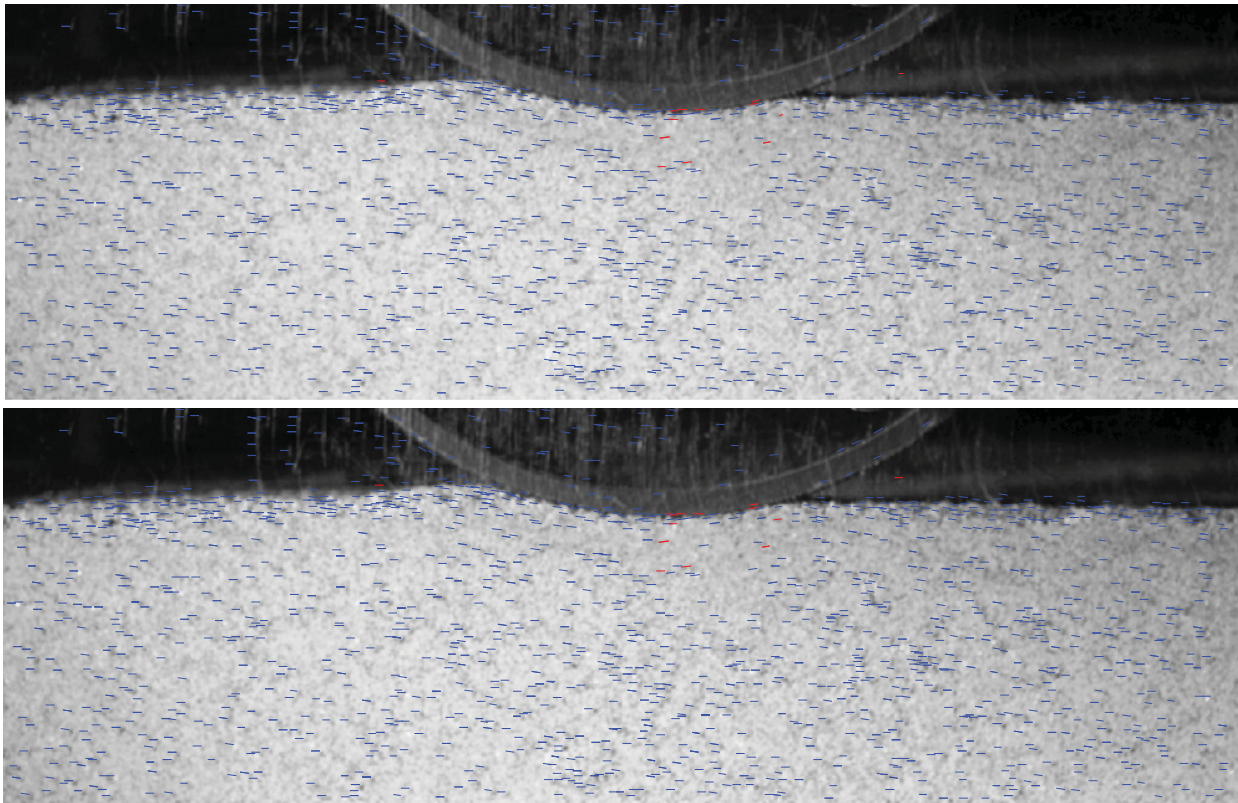
Figure 11. Soil flow captured using scale-invariant feature tracking (SIFT; top image) for cross-validation of optical flow (bottom). Motions at key points are indicated by line segments, shown blue when the two techniques agree within 0.56 pixels. Red lines indicate larger disparities, concentrated where the wheel interacts with the soil causing complex flow (see online for color).

driven at constant angular velocity in soil with a portion of particles dyed black. The black particles increased the number of features discernible by SIFT. For each SIFT feature, the corresponding optical flow (OF) motion is found by reading the u and v from the pixel nearest the feature. The cross-validation error between the motions computed by the two techniques is

$$\sqrt{(u_{OF} - u_{SIFT})^2 + (v_{OF} - v_{SIFT})^2} \qquad (7)$$

Sample output from both techniques for the driven wheel test is shown in Figure 11. Note that this output has not been converted to an inertial frame, so there is a dominant horizontal component to the motion. The top image shows motions of the tracked SIFT features between a pair of images, indicated by line segments. The bottom image shows the corresponding motion at the feature locations computed with optical flow.

For key points where the computed cross-validation error is less than 0.56 pixels (twice the 95% percentile error calculated from the ground-truth rotation test), the motion is displayed with blue line segments. When the error is larger, the motion is displayed in red. Note that these instances are rare, but they do tend to occur in regions of the soil affected by the test implement (where more complex flow occurs). Again, SIFT output is not ground-truth, so differences between optical flow and SIFT should not automatically be considered as errors on the optical flow side. However, these cross-validation tests do explicitly show that capturing soil motion for regions of complex flow is more challenging than for simple rotations and translations. Errors in these regions may thus be higher than those characterized in ground-truth testing. Further research is required to comprehensively quantify errors in regions of complex soil flow.

The fact that SOFT's optical flow output generally agrees with SIFT lends further confidence to the optical flow results. Optical flow is still highly preferable to SIFT, if choosing between the two, because it provides data at each pixel rather than a sparse subset. However, as discussed above, no individual technique has yet conclusively demonstrated superiority in capturing highly complex flows.
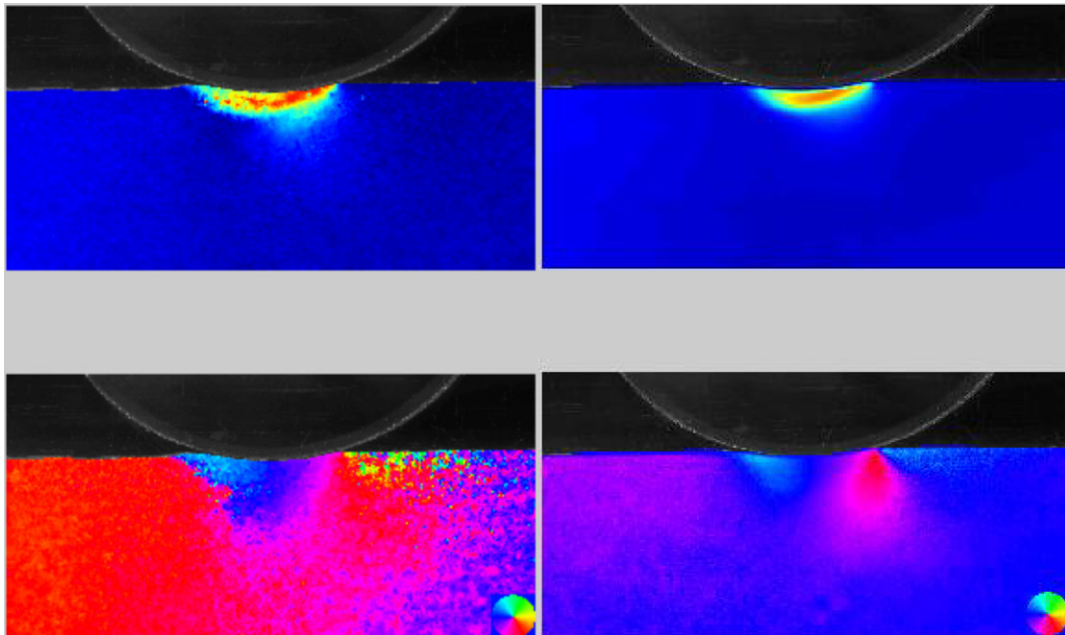
Figure 12. Soil flow calculated from a single frame pair (left), and averaged over 100 successive frame pairs (right) taken over the course of a rigid wheel test. The averaged data ignore transients caused by local soil irregularities, and they boost the signal-to-noise ratio.

3.3 Signal Averaging

It is often useful to discern average machine-soil interactions, ignoring transients caused by variability in the soil state, for example. To reduce the prevalence of transients in soil motion measurements, output can be averaged over successive frame pairs taken over the course of a test. Note that averaging is most appropriate for tests operating at steady state at their highest level (constant slip, sinkage, etc.). Conversely, averaging is not advised for tests with time-dependent soil response, as it can blur periodic events such as the responses to successive grouser passage.

Averaging can also boost the signal-to-noise ratio. In the ideal case, S/N is multiplied by the square root of the number of samples over which the signal is averaged. However, this assumes error is random and uncorrelated with the signal. As Figure 7 shows, this is not quite the case. Faint grid and spiral patterns are discernible in the error plots. Also, the error magnitude is (weakly) correlated to the magnitude of the signal, as discussed in Section 3.1.1. Averaging still boosts S/N, but not quite by the maximum amount predicted for the ideal case.

Figure 12 shows relatively noisy soil flow computed from a single frame pair, as well as the smoother flow computed by averaging 100 frames from the same test. In addition to smoothing the results for any (non-time-varying) test, averaging can actually reveal meaningful results for tests in which motion is too small to discern from a single frame pair. One such example, for a compli-

ant wheel that induced very little soil motion, is shown in Figure 13.

4. APPLICATIONS TO FIELD ROBOTICS

As mentioned in the introduction to this article, field robots (as well as other off-road vehicles and machines) are often required to traverse or work in granular soil. Wheels are commonly used in off-road applications that span construction, mining, agriculture, recreation, military operations, and planetary exploration. Agricultural and earthworking robots must also interact with granular soils using buckets, plows, and other tools.

This section describes application of the soil motion visualization and analysis technique to a field robot capable of a novel form of push-rolling locomotion. It also demonstrates the technique's applicability to field robot implements other than wheels.

4.1. Soil Motion Analysis for a "Push-rolling" Field Robot

Advancing a vehicle by changing its wheelbase, while rolling a subset of its wheels and keeping the others static to push against the ground, is referred to as push-rolling (Bekker, 1956). Scarab is a robot with independent suspension of its side-frames, and it achieves push-rolling by raising and lowering its body while expanding and contracting, like an inchworm (Scarab's push-rolling locomotion is thus
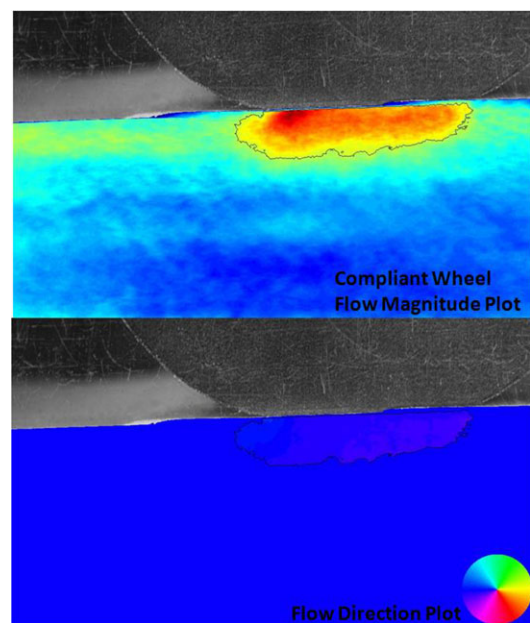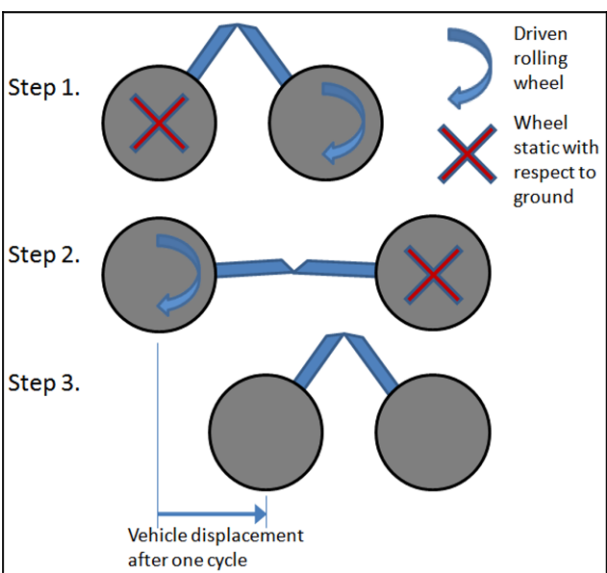
Figure 13. The soil motion induced by a compliant wheel is an order of magnitude smaller than that induced by a rigid wheel (shown in Figure 10, for example). To make meaningful observations, signal averaging (over 100 frames, in this case) is required.


Figure 15. Simplified Scarab kinematics demonstrating how push-rolling is achieved by rolling a single pair of wheels while side-frames expand or contract (Moreland et al., 2011).

sometimes called "inching") (Wettergreen et al., 2010). Photos of Scarab performing push-rolling in the field are shown in Figure 14.

The motion cycle commanded to achieve push-rolling with Scarab is shown in Figure 15. First, the angle between the side-arms is increased, lowering the body and expanding the wheelbase. During this expansion, the front wheels roll forward while the rear wheels remain static relative to the ground. Once the limit of body expansion is reached, the body raises and the wheel base contracts while the rear wheels roll forward and the front wheels remain

static with respect to the ground. To implement this motion on Scarab, the "static" wheels must actually counter-rotate in synchrony with the expanding or contracting side frame; if they were to be locked they would rotate in the soil as the side frame moves and thereby break their static contact with the soil.

The concept of push-rolling is not unique, and it has been investigated in the past. At the U.S. Army Land Locomotion Laboratory (Czako et al., 1963), a segmented vehicle with the ability to push-roll was introduced. It was determined through theoretical analysis that keeping one axle stationary and propelling the other forward decreases net resistance on the vehicle as a whole. It was surmised that the thrust generated by the stationary wheels would assist the rolling wheels and allow them to overcome greater
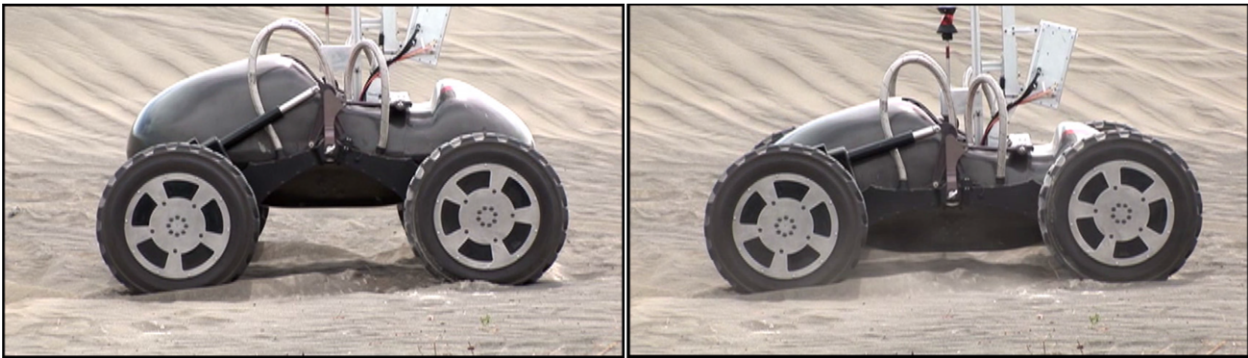
Figure 14. Scarab performing push-rolling during field tests.
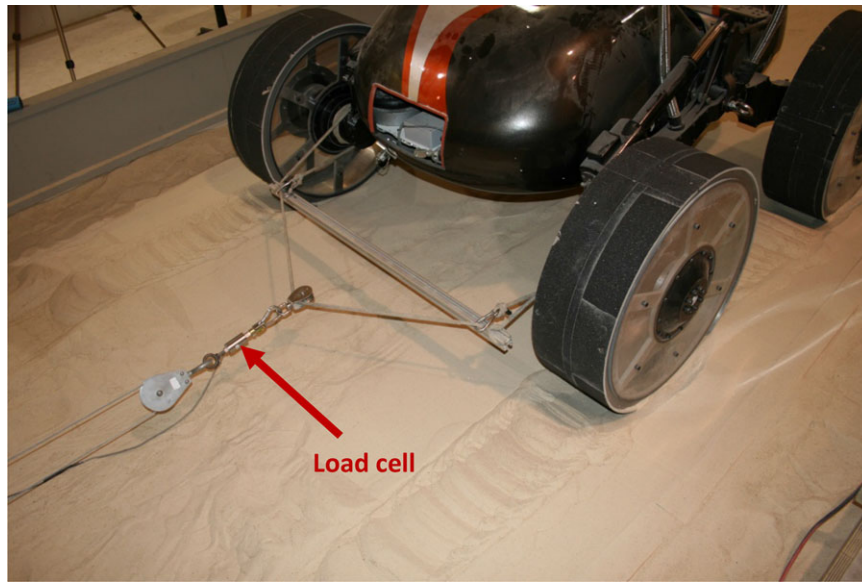
Figure 16.   Drawbar pull force being applied to Scarab through a load cell at rear wheel hubs (Creager et al., 2012).

rolling resistance. However, the possibility was not considered that the static wheels may in fact also generate higher thrust forces leading to additional utility for doing external work (e.g., climbing or towing).

This research [see also Creager et al. (2012) and Moreland et al. (2011)] provides a better understanding of the thrust benefit of push-rolling through full-scale robot tests in the laboratory as well as soil flow visualization and analysis using SOFT.

4.1.1.   Experimental Comparison between Rolling and Push-rolling

A series of experiments were conducted in the SLOPE laboratory at NASA GRC to quantitatively compare the performance of rolling to push-roll locomotion. One major challenge for off-road vehicles is traversing steep slopes in soft soil. The drawbar pull test can be used to predict a vehicle's slope-climbing performance (Freitag et al., 1970). Drawbar pull force is essentially the external force that a vehicle can generate, which is indirectly related to its ability to climb a slope. The test consists of driving a vehicle across prepared terrain while applying resistance opposite to the direction of travel (the drawbar pull force) and observing how the vehicle's driving performance is affected. To apply a drawbar pull force, a cable is attached to a hitch at the rear wheel hubs of Scarab (see Figure 16). The cable's tension is measured by a load cell and controlled (kept constant) using a magnetic clutch connected to the cable's spooling drum.

Travel reduction was chosen as the metric for evaluating these two modes. It can be defined as the reduction in vehicle velocity for a given drawbar pull force relative to the vehicle's self-propelled velocity (no drawbar pull force), normalized to the self-propelled velocity (assuming constant wheel rotational velocity):

$$TR = \frac{V_{sp} - V_{actual}}{V_{sp}} \times 100\% \tag{8}$$

TR = travel reduction (%),
V_{sp} = self-propelled velocity of vehicle (m/s or m/cycle),
V_{actual} = velocity of vehicle for given drawbar pull force (m/s or m/cycle),

$$= \frac{d(t) - d(t-1)}{\Delta t} \tag{9}$$

where d is the distance from total station (m), and t is the time (s).

For the case of rolling, distance measurements are taken with a laser total station at 1 Hz. For inching, because the vehicle velocity is not constant throughout each cycle, the velocities are calculated as distance per cycle (i.e., using one total station data point from the start of each cycle).

Figure 17 shows how drawbar pull affects travel reduction for various wheels (rigid versus pneumatic) and locomotion modes (rolling versus inching/push-rolling). The values near the top of the figure, during motion with high travel reduction, indicate the maximum drawbar pull achievable for each case. Push-rolling (inching) achieves at least 30% higher DP than that achievable with conventional rolling (0.33 versus 0.24 for rigid wheels, and 0.37 versus 0.28 for pneumatic tires).
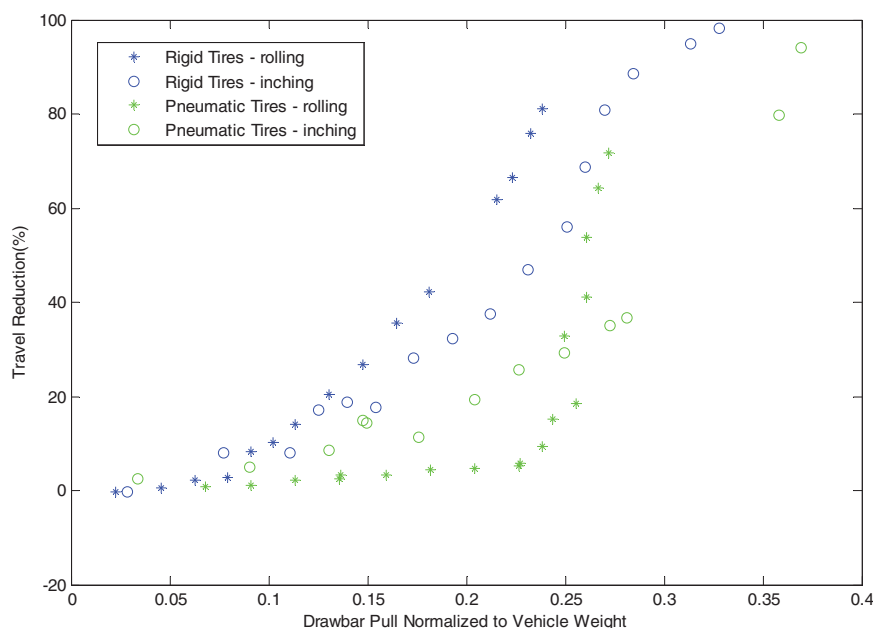
Figure 17. Comparison of rolling versus inching for compliant and rigid wheels in loose GRC-1 (Creager et al., 2012). For rolling, the load was increased approximately every 0.5 m; for inching, the load was held constant for at least two inching cycles, and then increased.

### 4.1.2. Understanding Soil Response using Visual Analysis

Two types of wheel travel were reproduced in the glass-walled soil bin described in Section 2. For the first type of travel, conventional rolling, the wheel and carriage were driven at speeds corresponding to slip and positive drawbar pull. This creates a noticeable amount of shearing in the terrain below the wheel. For the second type of motion, pushing, the wheel rotational speed was held at zero (braked) while the carriage was moved horizontally a short distance. This produced a situation similar to push-rolling, where the stationary wheels generate thrust by pushing off of the soil in a horizontal direction.

Examples of the soil shearing patterns for the two cases can be seen in Figure 18 (particle velocity magnitude) and Figure 19 (particle velocity direction). In both cases, the simulated motion of the full vehicle is to the right; the rolling wheel is driving clockwise and the braked/pushed wheel moves to the left slightly to push off of the terrain (note that in push-rolling, the static wheel is pushed opposite to the direction of net vehicle motion). Large differences in the direction and shape of the shearing patterns are evident. The soil beneath the rolling wheel appears to follow the shape of the wheel in a direction tangential to the wheel rotation. The failure occurs close to the wheel-terrain contact and is fairly uniform. This type of soil response was defined by Bekker as "grip" failure (Bekker, 1960). The braked/pushed wheel, however, produces a much different response in the soil.

The soil displacement occurs as a bulk quantity moving at once in a single direction. Bekker describes this soil response as "ground failure" or "general shear failure" (Bekker, 1960).

For these specific examples, the magnitude of the particle velocities cannot be used for comparison because different amounts of thrust are being generated here. However, the direction of the soil particles being displaced provides valuable insight. For rolling, only a certain portion of the particle motion is in the direction of thrust. Much energy is spent moving the soil vertically. The constant changing of particle direction requires more wheel motion to compact the soil and generate thrust. The braked/pushed wheel, on the other hand, creates soil displacement almost exclusively in the direction of thrust. Soil compaction is produced more quickly because the soil particles can only move in one direction. Due to this lack of vertical motion, pushing or pulling a braked wheel appears to be a much more efficient method of generating large amounts of thrust. This difference in soil failure could help explain the observed differences in performance for rolling versus inching. Adding to previous theories that claimed inching and other methods of push-roll locomotion exhibited more DP capability over rolling by decreasing the wheel resistance (Czako et al., 1963), the image analysis shown here indicates that an entirely different mode of soil response occurs when inching, producing an increase in thrust force as well.
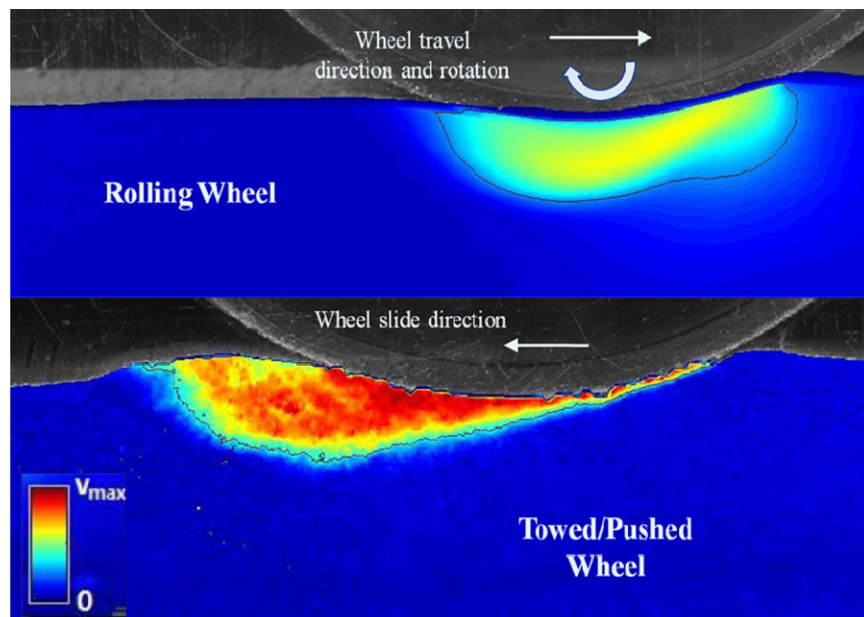
Figure 18. Soil velocity response to rolling wheel versus pushed wheel; color (see online) indicates the magnitude of soil motion.
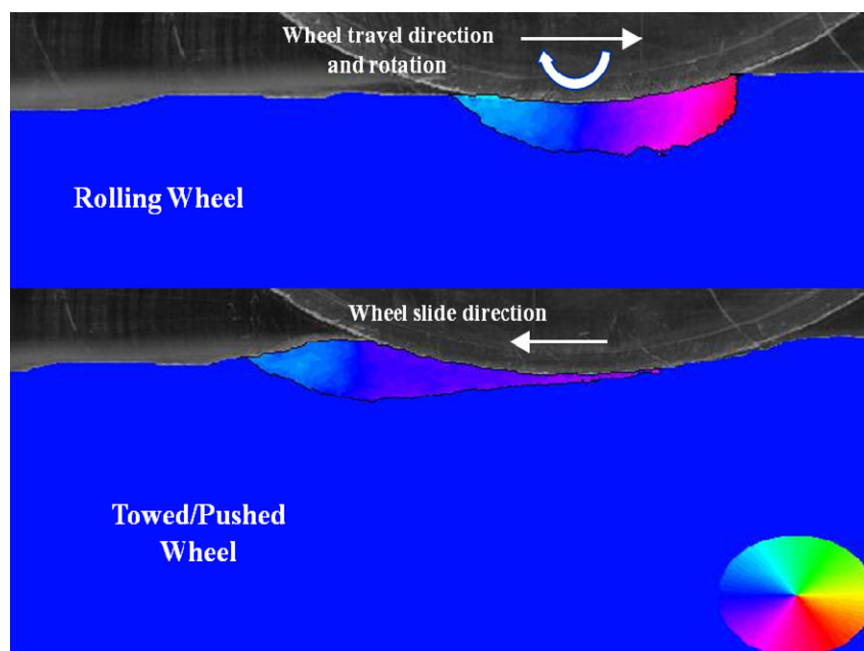

Figure 19. Soil directional response to rolling wheel versus pushed wheel; color (see online) indicates the direction of soil motion.

4.2 Soil Flow during Excavation

Excavation tools are another important type of robot-soil interaction. A deeper understanding of soil flow during excavation can lead to more efficient earthworking and agricultural machines. Agui (Agui & Wilkinson, 2010) describes the evolution of a pile of granular soil at the leading edge of a bucket cutting horizontally, based on measurements of the soil surface. Figure 20 shows how the SOFT images motion below the soil surface. The bucket in this analysis has half-scale dimensions compared to that described by Agui, and
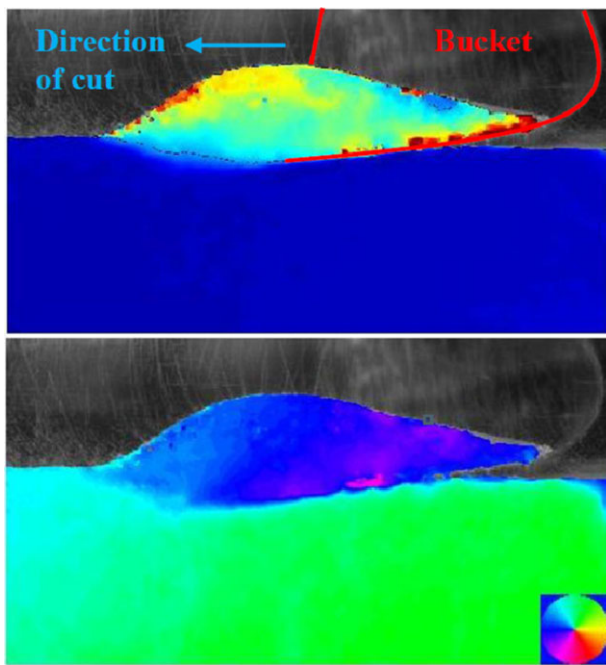
Figure 20. Soil is pushed forward as it is excavated with a bucket. Soil at the fore of the pile is also pushed up, while soil near the rear of the pile has a downward component to its flow.

it is commanded at half-scaled depth and forward speed, but with the same cut angle.

All soil in the pile is pushed forward (to the left in the image) with the bucket. Soil at the fore of the pile is also pushed up (light blue), while soil near the rear of the pile has a downward component (magenta) to its flow.

Additional applications of soil imaging to excavation tool analysis have been demonstrated in other ongoing work by the authors (Diaz Lankenau et al., 2012).

5. CONCLUSIONS AND FUTURE WORK

The work presents a novel experimental method for visualizing and analyzing machine-soil interactions, namely the soil optical flow technique (SOFT). It produces results that are far richer than past approaches used in terramechanics, with data at each image pixel. It achieves accuracy comparable to approaches used in geotechnical engineering, but without the need for specialized measurement equipment or specially marked soil particles. Its utility has been demonstrated on two different soil types (fine low-density Fillite and the relatively coarse GRC-1), both used without any modification of the soils or the processing technique.

SOFT uses computer vision techniques to process images of soil interacting with a wheel or tool. The critical processing step is optical flow, which estimates a data-rich flow field; each pixel of the image is a data point. The constraints imposed on the optical flow processing are particularly relevant to soil motions. Affine flow captures shearing and rotating motions as well as changes in aspect ratio, while an error norm that is robust to outliers (a Lorentzian norm) allows estimates with multiple distinct flows in a single proximate region.

The k-means clustering technique is effectively employed to classify regions of the image as soil/not-soil, and to identify shear interfaces between regions of significant soil flow and static soil.

A total of 95% of errors in SOFT's optical flow output are approximately 0.3 pixels or smaller (60 μm or smaller for the test conditions). The median error is approximately 0.15 pixels (30 μm). These errors were quantified with ground-truth experiments in which a flat plate covered in GRC-1 and Fillite soils was rotated by a precisely measured angle. The error is approximately constant, with just a weak correlation between signal and noise. In cases of relatively low signal-to-noise ratio (S/N), it is possible to boost it by signal averaging over several frame pairs capturing multiple instances of the same motion.

Additional ground-truth tests where a wheel is driven at constant slip and constant horizontal speed show that estimated bulk horizontal motion is constant within ±0.34 pixels (95% confidence bounds). This error is not much larger than the error induced by the optical flow technique itself, suggesting that errors in rig velocity control, vibrations in the experimental apparatus between camera and test implement, and inconsistencies in camera frame rate are not dominant sources of error in the overall experimental procedure.

Cross-validation of SOFT's optical flow output with scale-invariant feature tracking (SIFT) shows general agreement. The vast majority of points where motion is estimated by both techniques agree within 0.56 pixels (twice the error identified for the optical flow output). The few points where the two techniques disagree by more than 0.56 pixels are all found in the region with the most complex soil flow. This suggests that errors in these complex regions may be higher than those quantified by the (necessarily less complex) ground-truth experiments. Additional research is required to comprehensively quantify soil motion errors in regions of complex flow.

Using SOFT to observe the two aspects of a novel form of locomotion for field robots, termed push-rolling (conventional rolling as well as pushing horizontally with a static wheel) leads to a deeper understanding of the phenomenon than was possible before.

Tools for excavation are another important regime, in addition to wheels, of robot-soil interactions. The SOFT technique demonstrates utility in visualizing and analyzing this type of interaction. All soil in a pile is pushed forward when a simple bucket executes a horizontal cut. Additionally, soil at the fore of the pile is also pushed upward, while

soil near the rear of the pile has a downward component to its flow.

Ongoing and future work utilizing SOFT will lead to further discoveries related to wheel-soil interactions. Studying grousers has identified the strongly periodic effects they induce in soil. SOFT has engendered the discovery of periodic resistive forward flow (Moreland et al., 2012), leading to the development of a quantitative equation for predicting appropriate grouser spacing (Skonieczny et al., 2012). Other ongoing experimental campaigns include studying the effects of varying slip in different soil types, of varying wheel diameter, and of varying normal load on the wheel.

REFERENCES

Agui, J. H., & Wilkinson, R. A. (2010). Granular flow and dynamics of lunar simulants in excavating implements. In Earth and Space, (pp. 84–94).

Anandan, P. (1989). A computational framework and an algorithm for the measurement of visual motion. International Journal of Computer Vision, 2(3), 283–310.

Bekker, M. G. (1948). Photographic method of determining the soil action beneath footings. In Proceedings of the 2nd International Conference on Soil Mechanics and Foundation Engineering (pp. 193–194).

Bekker, M. G. (1956). Theory of land locomotion. Ann Arbor, MI: The University of Michigan Press.

Bekker, M. G. (1960). Off-the-road locomotion. Ann Arbor, MI: The University of Michigan Press.

Black, M. J., & Anandan, P. (1996). The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. Computer Vision and Image Understanding, 63(1), 75–104.

Cowen, E. A., & Monismith, S. G. (1997). A hybrid digital particle tracking velocimetry technique. Experiments in Fluids, 22(3), 199–211.

Creager, C., Moreland, S., Skonieczny, K., Johnson, K., Asnani, V., & Gilligan, R. (2012). Benefit of 'push-pull' locomotion for planetary rover mobility. In Earth and Space (pp. 11–20).

Czako, T. F., Janosi, Z. J., & Liston, R. A. (1963). An analysis of multi-element inching vehicles. Center Line, MI: U.S. Army Tank-automotive Center, Land Locomotion Laboratory.

Diaz Lankenau, G. F., Skonieczny, K., Whittaker, W. L., & Wettergreen, D. S. (2012). Effect of bucket-wheel scale on excavation forces and soil motion. Journal of Terramechanics, 49(6), 341–348.

Freitag, D. R., Green, A. J., & Melzer, K. J. (1970). Performance evaluation of wheels for lunar vehicles. Vicksburg, MS: U.S. Army Engineer Waterways Experiment Station.

Gachet, P., Klubertanz, G., Vulliet, L., & Laloui, L. (2003). Interfacial behavior of unsaturated soil with small-scale models and use of image processing techniques. Geotechnical Testing Journal, 26(1), 12–21.

Gnanamanickam, E. P., Lee, S., Sullivan, J. P., & Chandrasekar, S. (2007). Direct measurement of large-strain deformation field in machining. In Proceedings of the SEM Annual Conference and Exposition (Vol. 2, pp. 1080–1087).

Gonzalez, R. C., Woods, R. E., & Eddins, S. L. (2009). Digital image processing using MATLAB (2nd ed.). Knoxville, TN: Gatesmark Publishing.

Harrison, W. L. (1973). Soil failure under inclined loads—II. Journal of Terramechanics, 10(1), 11–50.

Lowe, D. G. (1999). Object recognition from local scale-invariant features. In Proceedings of the Seventh IEEE International Conference on Computer Vision (Vol. 2, pp. 1150–1157).

Moreland, S. J., Skonieczny, K., & Wettergreen, D. S. (2011). Inching locomotion for planetary rover mobility. In IEEE Aerospace (pp. 1–6).

Moreland, S., Skonieczny, K., Inotsume, H., & Wettergreen, D. (2012). Soil behavior of wheels with grousers for planetary rovers. In IEEE Aerospace (pp. 1–6).

Murthy, T. G., Gnanamanickam, E. P. Saldana, C., & Chandrasekar, S. (2009). Deformation field in indentation of granular materials. In Proceedings of the 6th International Conference on Micromechanics of Granular Media (pp. 263–266).

Oravec, H., Asnani, V., & Zeng, X. (2010). Design and characterization of GRC-1: A soil for lunar terramechanics testing in Earth-ambient conditions. Journal of Terramechanics, 47(6), 361–377.

Skonieczny, K., Moreland, S. J., & Wettergreen, D. S. (2012). A grouser spacing equation for determining appropriate geometry of planetary rover wheels. In IEEE Intelligent Robots and Systems (IROS) (pp. 5065–5070).

Sun, D., Roth, S., Lewis, J. P., & Black, M. J. (2008). Learning optical flow. In Computer Vision–ECCV 2008 (pp. 83–97).

Szeliski, R. (2011). Computer vision: Algorithms and applications. New York, NY: Springer.

Wettergreen, D., Moreland, S., Skonieczny, K., Jonak, D., Kohanbash, D., & Teza, J. (2010). Design and field experimentation of a prototype Lunar prospector. The International Journal of Robotics Research, 29(12), 1550–1564.

White, D. J., Take, W. A., & Bolton, M. D. (2003). Soil deformation measurement using particle image velocimetry (PIV) and photogrammetry. Geotechnique, 53(7), 619–631.

Witney, B. D. (1968). The determination of soil particle movement in two dimensional failure. Journal of Terramechanics, 5(1), 39–52.

Wong, J. Y. (1967). Behaviour of soil beneath rigid wheels. Journal of Agricultural Engineering Research, 12(4), 257–269.

Wong, J. Y., & Reece, A. R. (1966). Soil failure beneath rigid wheels. In Proceedings of the 2nd International Conference of the International Society for Terrain-vehicle Systems (pp. 425–445).