Proceedings of the 1991 IEEE
International Conference on Robotics and Automation
Sacramento, California - April 1991

Laser Rangefinder Calibration for a Walking Robot

Eric Krotkov
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

In rugged terrain, walking robots that select footholds can be more mobile and more energy efficient than machines that roll on wheels or crawl on tracks. To achieve these footholds requires calibration of the terrain sensors with respect to the walking mechanism. We present an implemented technique to calibrate scanning laser rangefinders to legged robots. The procedure accommodates two scanners, one manufactured by Erim and the other by Perceptron, and two walkers, a one-legged robot and a six-legged robot. The technique acquires two sets of corresponding three-dimensional points and identifies the rigid transformation that maps one onto the other with least squared error, i.e., it solves the absolute orientation problem. We report experimental results with the two different scanners and vehicles. For the Erim and the one-legged robot, the technique achieves an accuracy of 6–12 cm with a precision no lower than 2–5 cm. For the Perceptron and the Ambler, the accuracy is 2–7 cm with a precision no lower than 2–5 cm. These results have proven to be satisfactory for constructing terrain maps and using them to select footholds during our rough terrain walking experiments.

1 Introduction

In order to act autonomously and intelligently, mobile robots must be able to sense their environment, and to relate the sensor readings to their actions. For example, threading a needle requires coordination of the eye, or whatever senses the relative positions of needle and thread, and the hand, or whatever acts on them. Similarly, walking requires coordination of the eye, which senses where on the terrain to place the foot, and the leg, or whatever supports and propels the robot. This paper presents a technique to establish such "leg-eye" coordination for a six-legged robot (Figure 1, the Ambler planetary rover [1]), and for a one-legged robot (Figure 2).

Recently, researchers have questioned the need for calibration, and techniques to avoid it have gained favor. In the case of the Ambler, it is likely that we too could survive without calibration; because the machine is big, rugged, and heavy, many objects that are obstacles *before* the Ambler steps on them are *planar* afterwards. But for missions to distant, rugged regions like planetary surfaces, Antarctica, and the ocean floor, survival is not enough. *Energy-efficient* locomotion is essential. With a calibrated sensor, the Ambler can select where to step, and thus can prevent spending significant fractions of the total power budget on stumbling rather than productive advance. This ability to select footholds is central to the fundamental advantages of high mobility and energy efficiency that walkers enjoy over rolling and crawling machines. Achieving those footholds requires calibration.

Other walking robots face the same requirements. However, walkers that rely on a human operator to designate footholds, such as the Adaptive Suspension Vehicle [6], do not require an automated solution.

In this paper we present a single calibration procedure that works for multiple scenarios; the same code calibrates the Erim scanner with respect to a one-legged robot, and calibrates the Perceptron scanner with respect to the six-legged Ambler. In Section 2 we define the problem. Next, we describe in detail how to acquire two three-dimensional point sets, one in a vehicle-centered reference frame (Section 3), the other in a sensor-centered reference frame (Section 4). Then we show how to identify the rigid transformation that best relates the two point sets, i.e., we present a solution to the absolute orientation problem. In Section 6 we report experimental results on accuracy, precision, and execution time. We conclude by discussing possible improvements and extensions.

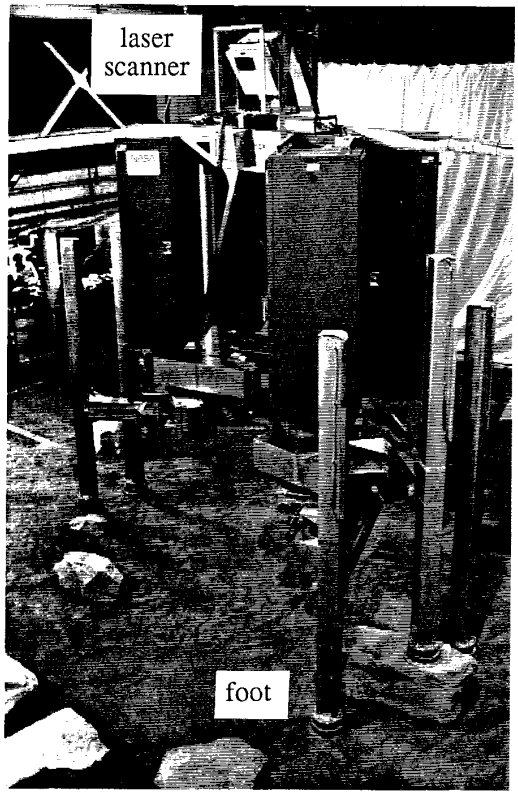

Figure 1: Six-legged robot
The figure shows the Ambler, and the Perceptron laser scanner mounted on the bridge between the two leg stacks.

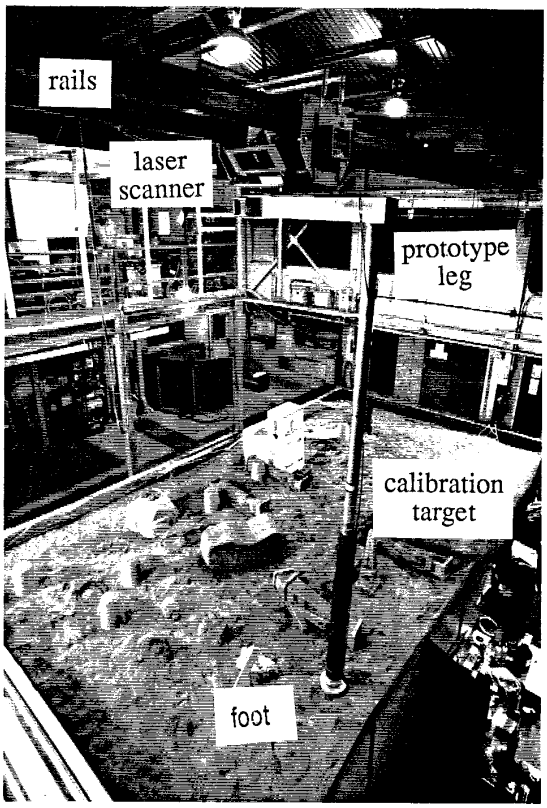

Figure 2: One-legged robot
The figure shows the prototype leg, the calibration target on the upper leg, and the Erim laser scanner mounted above the leg.

2 Problem Definition

The overall problem is to identify the rigid transformation relating a vehicle-centered reference frame to a sensor-centered reference frame. The origin of the *scanner frame* $\mathcal{S}$ is attached to the scanner and lies somewhere nearby it. The origin of the *body frame* $\mathcal{B}$ is attached to the walking robot.

We attach a number T of targets to the legs. Then, we move the legs to a number L of different stations. At each, we identify the position $\vec{r}_B$ of each target in the body frame (by reading joint positions and using known kinematics, see Section 3), and we identify the position $\vec{r}_S$ of each target in the scanner frame (by image analysis, see Section 4). After acquiring a sufficient number of pairs of measurements, we seek the rotation $\mathbf{R}$ and translation $\vec{t}$ that refer a vector in $\mathcal{S}$ to $\mathcal{B}$:

$$\vec{r}_{Bi} = \mathbf{R}\vec{r}_{Si} + \vec{t}\ , \quad 1 \le i \le L \times T\ , \qquad (1)$$

where $\vec{t} = [t_x, t_y, t_z]^T$ is the translation vector relating the two origins, and $\mathbf{R}$ is a 3×3 rotation matrix ($\det \mathbf{R} = +1, \mathbf{R}^T\mathbf{R} = \mathbf{I}$).

In practice, it is unlikely that $\mathbf{R}$ and $\vec{t}$ exist that satisfy Equation 1, because measurements are not exact and may be contaminated by noise. Instead, we seek $\mathbf{R}$ and $\vec{t}$ that best satisfy Equation 1 in the least-squares sense: Find $\mathbf{R}$ and $\vec{t}$ minimizing the sum of squares of errors

$$E = \sum_{i=1}^{L\times T} \|\vec{e}_i\|^2\ , \qquad (2)$$

where $\|\vec{x}\|^2 = \vec{x} \cdot \vec{x}$ is the square of the length of the vector $\vec{x}$, and the error of the i^{th} pair of measurements is $\vec{e}_i = \vec{r}_{Bi} - \mathbf{R}\vec{r}_{Si} - \vec{t}$.

This problem is closely related to a number of other problems that arise in photogrammetry and computer

vision. Given the pairs of measurements, the problem is equivalent to the *absolute orientation* problem in photogrammetry, and to the *exterior orientation* part of the *camera calibration* problem in computer vision (see Chapter 13 of [3], and references therein).

3 Target in Body Frame

For the one-legged prototype, the origin of $\mathcal{B}$ coincides with the shoulder joint. Calibration targets (pieces of reflective tape) are attached to the upper leg at fixed and manually measurable positions.

For the six-legged Ambler, the origin of $\mathcal{B}$ lies at the center of the downward-facing surface of the structural bridge that connects the two leg stacks and supports the scanner. Calibration targets (pieces of brown paper) are attached to the top of the vertical links.

In both cases, the joint angles and link lengths are known, so elementary kinematics suffice to determine the target coordinates with respect to $\mathcal{B}$.

4 Target in Sensor Frame

We consider two scanning laser rangefinders, one manufactured by Erim, and the other by Perceptron. The devices digitize two images: a *range image*, with pixel values proportional to object distance modulo a known constant; and a *reflectance image*, with pixel values proportional to reflected energy.

The Erim acquires data in 64×256 pixel images at a rate of 2 Hz [7]. The scanner digitizes to 8 bits over approximately 20 meters, which provides a nominal range resolution of 7.62 cm. The measurements cover 80 degrees in the horizontal direction (azimuth) and 30 degrees in the vertical direction (elevation).

The Perceptron acquires data in 256×256 pixel images at a rate of 2 Hz. The scanner digitizes to 12 bits over approximately 40 meters, which provides a nominal range resolution of 0.98 cm. The measurements cover 60 degrees in azimuth and 60 degrees in elevation.

Given a pair of reflectance and range images $\alpha(u,v)$ and $\beta(u,v)$, as in Figure 3, the task is to compute the image coordinates of the target. The four following steps perform this task [5].

1. Register the range and reflectance images.

2. Locate the leg in image space.

2a. Threshold the range image (Figure 4), removing pixels with ranges that are either too close or too far

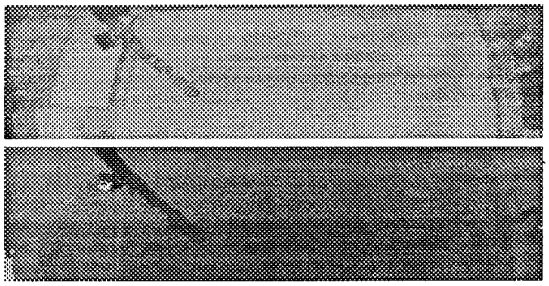

Figure 3: Erim images of prototype leg
The leg appears in the left-hand side of the reflectance (top) and range images as a tapered cylinder. The bump on its lower left side is a cable reel; the cable appears faintly below the reel. In the reflectance image, the calibration target appears as a bright band below the reel.

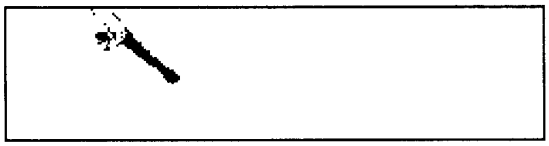

Figure 4: Thresholded Erim range image of leg
The output contains regions that do not belong to the leg, and is noisy.

to lie on the upper leg in the vicinity of the target(s):

$$\beta_{thresh}(u,v) = \begin{cases} 1 & \text{if } I_{near} \le \beta(u,v) \le I_{far}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

2b. Filter the thresholded range image by first shrinking and then growing, in order to eliminate regions that are too small to be the target (Figure 5). To eliminate small regions, threshold (*shrink*) the output of the Grassfire transform as follows:

$$\beta_{shrink}(u,v) = G(\beta_{thresh}(u,v), K_{grass}) \ , \quad (4)$$

where G is the forward Grassfire transform [5]. To restore (*grow*) the target region to its original size, apply the reverse Grassfire transform as follows:

$$\beta_{leg}(u,v) = G^{-1}(\beta_{shrink}(u,v), K_{grass}) \ . \quad (5)$$

3. Identify target pixels on the leg to be those that both belong to the leg, and exceed an intensity threshold that is characteristic of the calibration target.

4. Compute the centroid $(\bar{u}, \bar{v})$ of the target pixels.

The method for Perceptron imagery is similar. For brevity, we refer readers interested in the differences to

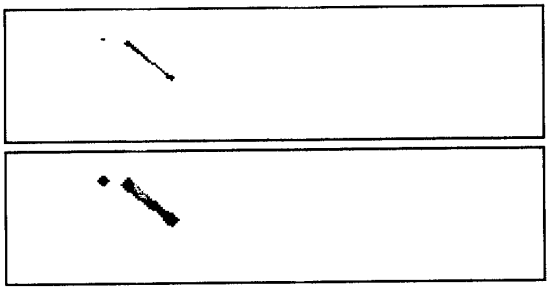

Figure 5: Grassfire, forward (top) and reverse.

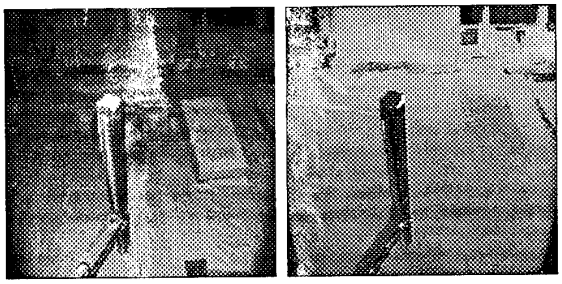

Figure 6: Perceptron images of Ambler leg
The contrast of the reflectance (left) and range images has been enhanced. The extensional link of the leg appears in the bottom left, and the vertical link appears in the center. In the center of the reflectance image, the calibration target is visible as a white patch at the top of the vertical link.

[5], and illustrate typical inputs in Figure 6 and intermediate results in Figure 7.

The completed image analysis yields row $\bar{u}$, column $\bar{v}$, and sensed range $\beta(\bar{u},\bar{v})$ of the target. The transformation to spherical-polar coordinates [5] is

$$\phi = \bar{u}\Delta_\phi + \phi_0 \ , \ \theta = \bar{v}\Delta_\theta + \theta_0 \ , \ \rho = k\beta(\bar{u},\bar{v}) \ , \quad (6)$$

where Δ_ϕ and Δ_θ are angular increments, ϕ_0 and θ_0 are initial orientations, and k is the scanner range resolution in meters/bit. The transformation to Cartesian coordinates is given by

$$x = \rho\sin\theta \ , \ y = \rho\cos\theta\cos\phi \ , \ z = \rho\cos\theta\sin\phi \ . \quad (7)$$

5 Solution

There are a number of approaches to solving the least-squares absolute orientation problem, which we review elsewhere [5]. We implement the technique in [2] (related to [4]), which is an exact closed-form solution

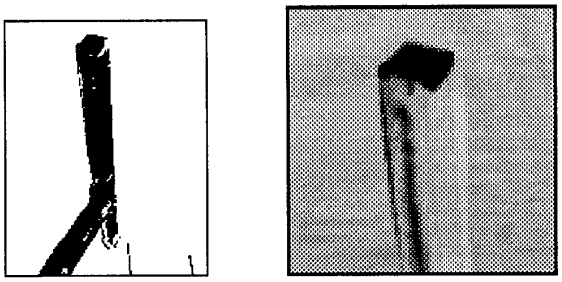

Figure 7: Perceptron range images of leg
The left panel illustrates the result of rangethresholding. This typical result includes pixels that do not belong to the leg, and does not include some pixels that do belong to the leg. The right panel shows the result of intensity thresholding, prior to executing the Grassfire transform.

that uses unit quaternions to represent the rotation. The quaternion representation affords two advantages: simplicity — it is simpler to enforce a unit norm for a quaternion than it is to ensure that a matrix is orthonormal; and closed-form solution — no iteration is required.

The solution for the desired quaternion is the eigenvector associated with the smallest eigenvalue of a symmetric matrix, whose elements are combinations of sums of products of corresponding coordinates from $\vec{r}_S$ and $\vec{r}_B$. For convenience in the physical interpretation of the solution rotation, we do not use the solution quaternion $\dot{q}$ itself. Instead, we first express $\dot{q}_{min}$ as an orthonormal matrix $\mathbf{R}$, and then parameterize $\mathbf{R}$ under the roll-pitch-yaw convention.

6 Results

To state the experimental procedure, let i and j be counters, B and S be sets of target position vectors, L be the number of leg stations, and T be the number of targets on each leg.

1. $i \leftarrow 1$, $B \leftarrow \emptyset$, $S \leftarrow \emptyset$
2. while $i \leq L$ do
   (a) Move leg to station i
   (b) Compute $\vec{r}_{Bij}$, $1 \leq j \leq T$, as in Section 3
   (c) $B \leftarrow B \cup \{\vec{r}_{Bij}\}$, $1 \leq j \leq T$
   (d) Acquire reflectance and range images
   (e) Compute $\vec{r}_{Sij}$, $1 \leq j \leq T$, as in Section 4

(f) $S \leftarrow S \cup \{\vec{r}_{Sij}\}$, $1 \leq j \leq T$

(g) $i \leftarrow i + 1$

3. Compute $\mathbf{R}$ and $\vec{t}$ from B and S, as in Section 5

We have automated this procedure entirely, and have executed it hundreds of times, perhaps a thousand times. Over the course of these trials, we have tested the procedure under a wide variety of conditions, including 50°F differences in temperature, ambient illuminations ranging from bright sunlight to night-time fluorescent lighting, and background materials ranging from sand to people to heavy machinery. The time required to execute the calibration procedure with 10–20 points totals approximately 15 minutes for either scanner. The procedure spends roughly 70 percent of this time moving the legs, and roughly 20 percent acquiring and filtering images. Computing the rigid motion parameters accounts for most of the remaining 10 percent.

6.1 Accuracy

One basis for evaluating the accuracy of the computed parameter values is the distance ${d_i}^2 \equiv \vec{e}_i \cdot \vec{e}_i$ between corresponding points after applying the transformation, where $\vec{e}_i = \vec{r}_{Bi} - \mathbf{R}\vec{r}_{Si} - \vec{t}$. Figure 8 shows $\bar{d}$ (the mean of the distribution of the d_i) for ten Erim trials. We observe that the mean varies from 5.5 to 11.2 cm, and that it varies significantly between data sets of different sizes, and between data sets of the same size. This range of values and variations is typical of other trials. We conclude that the accuracy of the calibration procedure for the Erim is 6-12 cm.

Figure 8 also shows $\bar{d}$ for ten Perceptron trials. The results show that the mean values range from 1.8 to 6.7 cm, and vary significantly between data sets. This range of values and variations is typical of other trials. We conclude that the accuracy of the calibration procedure for the Perceptron is 2–7 cm, or two times greater than for the Erim. This difference in accuracy may be due to better leg position sensing with the Ambler, or superior accuracy on the part of the Perceptron scanner, or both.

6.2 Precision

To evaluate the precision of the computed rigid motion parameters, we execute the calibration procedure several consecutive times, and observe the difference in the computed parameters. Between data sets, we do not move the scanner or alter any settings.

Trial	N	$\bar{d}$ (cm)	s_D (cm)	Trial	N	$\bar{d}$ (cm)	s_D (cm)
1	8	5.5	2.4	11	12	2.5	1.0
2	8	5.7	2.4	12	12	2.2	1.3
3	8	6.4	1.5	13	12	6.7	4.3
4	8	5.8	1.9	14	12	2.9	1.6
5	18	9.6	4.3	15	12	6.7	3.8
6	18	10.1	3.1	16	12	6.0	3.2
7	18	8.6	3.8	17	12	4.0	3.0
8	50	11.2	3.6	18	12	2.7	1.5
9	50	8.9	2.9	19	12	1.8	1.9
10	50	10.7	4.3	20	12	2.9	1.5

Figure 8: Statistics of error distributions
Trials 1–10 use the Erim, trials 11-20 use the Perceptron.

Trial	Δt_x (cm)	Δt_y (cm)	Δt_z (cm)	$\Delta\phi$ (deg)	$\Delta\theta$ (deg)	$\Delta\psi$ (deg)
A,B	16.1	7.3	-6.5	0.4	2.0	-0.8
B,C	-3.4	10.6	-0.2	0.1	-0.5	1.4
A,C	12.7	-3.2	-6.7	0.6	1.5	0.6
D,E	7.4	3.5	2.1	1.6	0.8	-0.7
E,F	-3.8	-3.6	-5.7	-0.7	-0.5	1.6
D,F	3.6	-0.1	-3.6	0.8	0.3	0.7

Figure 9: Variation of the computed parameters
Trials A–C use the Erim with $N = 18$. Trials D–F use the Perceptron with $N = 12$.

Figure 9 shows by how much the estimated parameters change between three data sets for each sensor. Some parameters change by significant amounts; the largest observed differences are 16.1 cm and 2.0° for the Erim and 7.4 cm and 1.6° for the Perceptron.

The variations in the rigid motion parameters are not due to numerical instability of our computations, because we observe condition numbers between 1 and 120 for the symmetric matrix mentioned in Section 5. The variations in the rigid motion parameters may be due to poor leg position sensing, or poor image analysis, or both. If poor leg position sensing reported target locations imprecisely, then the rigid motion parameters would vary, even with perfect image analysis. Similarly, the parameters would vary if poor image analysis reported target locations imprecisely, even with infinitely repeatable leg position sensing. For the Erim trials, both factors are conflated, making it difficult to

identify their relative magnitudes. For the Perceptron trials, sensing the position of the stiff Ambler legs is highly repeatable, so the dominant cause of poor calibration precision is image acquisition and analysis. We observe the range variations to correlate with ambient temperature; the higher the temperature, the greater the variations. We have not been able to eliminate the variations, but can achieve acceptable precision by conducting calibration during cool conditions, where inferior to 75° is our heuristic measure of cool.

To assess the precision of the points transformed into $\mathcal{B}$, we evaluate the standard deviation s_D of the distribution of the d_i. This statistic quantifies the scatter of the points mapped into $\mathcal{B}$. Figure 8 shows s_D for ten Erim trials. We observe that the values of s_D range from 1.5 to 4.3 cm, and that they vary significantly between data sets of different sizes, and between data sets of the same size. This range of values is somewhat better than for other trials, where it is not uncommon to observe standard deviations of 10 cm. We conclude that the precision of the calibration procedure for the Erim is no better than 2–5 cm. Figure 8 shows s_D for ten Perceptron trials. The results show that the precision ranges from 1.0 to 4.3, which is commensurate to that of the Erim, and several times worse than the resolution of the Perceptron. As in the case of the Erim, the reported range of values is somewhat better than for other trials.

7 Discussion

In this paper we have presented an implemented technique to calibrate scanning laser rangefinders to a vehicle-centered reference frame. We reported results for two different sensors and two different vehicles. For the Erim and the prototype leg, the procedure achieves an accuracy of 6–12 cm with a precision no lower than 2–5 cm. For the Perceptron and the Ambler, the accuracy is 2–7 cm with a precision no lower than 2–5 cm. These results have proven to be satisfactory for constructing terrain maps and using them to select footholds during our rough terrain walking experiments.

We have also successfully calibrated the Perceptron scanner to a fixed reference frame. The techniques and results are very similar to those reported for the vehicle-centered reference frames, so we have not treated them separately. We note that this success provides further evidence for the generality of the approach.

One promising direction for future research is toward more comprehensive sensor calibration. This would involve identification of intrinsic sensor parameters such as the relationship between range grey level and absolute distance, the mirror starting angles, and the angular increments, in addition to the six rigid motion parameters.

Acknowledgements

This research was sponsored by NASA under Grant NAGW-1175. We would like to acknowledge contributions by the entire CMU Planetary Rover research group. We thank R. Hoffman and P. Balakumar for acquiring and analyzing data, K. Arakawa for finding targets in Perceptron images, M. Hebert for exact solutions and helpful discussions, M. Blackwell for scanner interfaces, N. Harding for designing the Perceptron mount, G. Roston for assistance with survey instruments, and C. Tomasi and A. Garvin for commenting on a draft of this paper.

References

[1] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: An Autonomous Rover for Planetary Exploration. *IEEE Computer*, 22(6):18–26, June 1989.

[2] O. Faugeras and M. Hebert. The Representation, Recognition and Locating of 3D Shapes from Range Data. *International Journal of Robotics Research*, 5(3):27–52, Fall 1986.

[3] B. Horn. *Robot Vision*. MIT Press, Cambridge, Massachusetts, 1986.

[4] B. Horn. Closed-form Solution of Absolute Orientation using Unit Quaternions. *Journal of the Optical Society of America A*, 4:629–642, 1987.

[5] E. Krotkov. Laser Rangefinder Calibration for a Walking Robot. Technical Report CMU-RI-90-30, Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, December 1990.

[6] S. Song and K. Waldron. *Machines that Walk: The Adaptive Suspension Vehicle*. MIT Press, Cambridge, Massachusetts, 1989.

[7] D. Zuk, F. Pont, R. Franklin, and V. Larrowe. A System for Autonomous Land Navigation. Technical Report IR-85-540, Environmental Research Institute of Michigan, Ann Arbor, Michigan, 1985.