

Relative Localization from Image Sequences

Brian L. Stuart
Innovation/Scanning Technology
FedEx Corporate Services
2847 Business Park Dr.
Memphis, TN 38118
blstuart@fedex.com

ABSTRACT

This paper presents a technique for vehicle localization. The technique is based on processing images taken by one or more downward pointing cameras. Motion represented by shift between successive images allows us to directly measure the velocity at two points on either side of the axis of motion. From these possibly different velocities, we can infer the vehicle speed and turning radius and from those, the vehicle orientation, position, velocity and acceleration vectors.

Categories and Subject Descriptors

L.4.8 [Image Processing and Computer Vision]: Scene Analysis—*motion, tracking*; I.2.9 [Artificial Intelligence]: Robotics—*autonomous vehicles*

General Terms

Algorithms, Localization

1. BACKGROUND

In order to make navigational decisions, an autonomous robot must know its position and orientation. Borenstein, et al[1] present a good survey of common techniques in the field of localization. Most techniques can be categorized according to whether they provide absolute position and orientation or relative position and orientation.

Common examples of absolute positioning systems include the Global Positioning System (GPS), magnetic compasses and landmark identification. All of these require external references. These external references are subject to interference. For example, vehicles traveling cross-country almost always experience intermittent loss of GPS signals. Environments with significant steel structures (such as warehouses and airports) will interfere both with GPS signals and with magnetic compasses. Landmark-based localization, on the other hand, requires that the robot have prior knowledge of the location and appearance of landmarks. In some applications, this is simply not possible. In others, anything

which removes, defaces or obscures a landmark interferes with localization.

The most common forms of relative positioning are wheel odometry and inertial navigation. Just as an automobile counts axle rotations to measure distance traveled, many robotic systems measure axle or wheel rotations to determine position. Steering angle can be included to determine orientation. More sophisticated implementations measure the rotation of all wheels to determine both orientation and position. Incremental changes in wheel position are integrated to give a total distance traveled. Because we are integrating once, a constant error in the data will yield a distance error that grows linearly with time. However, wheel odometry is subject to unobserved interactions with the environment. The most common of these is wheel slip. Just because a wheel turns a certain angle about a known axis does not mean that the vehicle necessarily moves accordingly. Sometimes the wheels will lose traction and at other times external forces will cause motion not induced by the wheels.

Inertial navigation systems are based on measurements from accelerometers or gyroscopes. While they are not subject to the environmental interactions that can affect wheel odometry, they are not without a downside. Since we must integrate acceleration twice to get position, any constant error will manifest itself as a position error that grows with the square of time. Gyroscopes, on the other hand, suffer from drift over time even when the vehicle is still.

These techniques are rarely used in isolation. For example, in a constrained environment such as a factory, we may provide landmarks which can update a vehicle's absolute position when encountered but use relative positioning systems between landmarks. Similarly, in cross-country applications, we often use GPS to give an absolute position when possible. However, during loss of GPS signal, wheel odometry and inertial navigation devices are often used to augment GPS. It is quite common for the various sources of data to be combined using Kalman filters[4].

The remainder of this paper describes a new technique that shows the linear error growth of odometry without being subject to the same environmental effects. It is also not subject to the drift that affects gyroscopic systems.

2. MOTION INFERENCE

Our technique is a form of odometry based on processing images of the ground taken as the vehicle moves. Thus it is immune to issues like wheel slippage. This is similar to the technique used by optical mice as they track movement

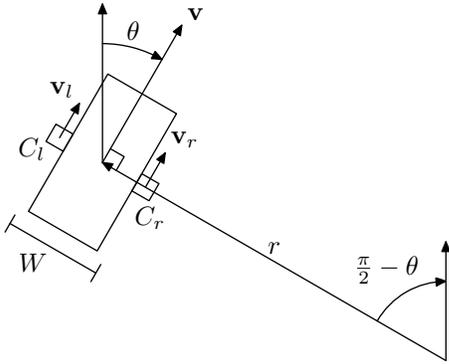


Figure 1: Vehicle Geometry

over textured surfaces. Suppose we have two images I_1 and I_2 taken Δt seconds apart. If we correlate the images and find that I_2 is shifted from I_1 by Δx pixels in the horizontal direction and Δy pixels in the vertical, then we can compute an instantaneous approximation to the vehicle velocity by

$$v = \frac{s\sqrt{(\Delta x)^2 + (\Delta y)^2}}{\Delta t} \quad (1)$$

where s is a scaling factor that relates a distance in pixels to a distance in the physical world. In the experiments reported here, it is calibrated in meters/pixel.

This computation of velocity gives us only the magnitude which would be sufficient if the vehicle always traveled in a straight line. In order to determine the direction of the velocity, however, we must compute the change in orientation. We base this calculation on the assumption that between successive images the vehicle moves in a circular arc with radius r , as illustrated in Figure 1. Before deriving an expression for r , we need to point out a simplifying observation. The vehicle's orientation is always at right angles to the radius of the turn. Consequently, the change in orientation angle is always equal to the change in angle about the center of the turn arc. This observation allows us to speak solely in terms of the change in orientation angle throughout this paper. Furthermore, throughout this derivation, we talk about a turn to the right. It should be noted that the results still follow in the case of a left-hand turn if we take $r < 0$ to indicate the left-hand direction. Referring again to Figure 1, the distance traveled and the length of the arc, l , are related by

$$l = vt = r\theta$$

where v is the magnitude of the velocity vector, \mathbf{v} . Thus

$$\theta = \frac{1}{r}vt$$

and the change in orientation is given by

$$\frac{d\theta}{dt} = \frac{1}{r}v. \quad (2)$$

We position two cameras (labeled C_l and C_r in Figure 1) separated by a distance W such that they lie along the turning radius and such that the center between them coincides with the center of the vehicle. Since the change in orientation of the cameras is the same as that of the vehicle, the velocity

seen by the left-hand camera, C_l , is given by

$$v_l = \frac{1}{r}v \left(r + \frac{W}{2} \right) = \left(1 + \frac{W}{2r} \right) v$$

and that seen by the right-hand camera, C_r , is given by

$$v_r = \left(1 - \frac{W}{2r} \right) v.$$

Solving for v and setting the two expressions equal gives us

$$v = \frac{v_l}{\left(1 + \frac{W}{2r} \right)} = \frac{v_r}{\left(1 - \frac{W}{2r} \right)}.$$

Solving for r gives

$$r = \frac{W}{2} \left(\frac{v_l + v_r}{v_l - v_r} \right).$$

Working backwards, we find that the vehicle velocity is just the mean of the left and right camera velocities,

$$v = \frac{v_l + v_r}{2}. \quad (3)$$

From the turning radius, we can calculate the orientation angle at time T by integrating (2)

$$\theta = \int_0^T \frac{1}{r} v dt \quad (4)$$

$$= \int_0^T \frac{2}{W} \left(\frac{v_l - v_r}{v_l + v_r} \right) \frac{v_l + v_r}{2} dt \quad (5)$$

$$= \int_0^T \frac{1}{W} (v_l - v_r) dt. \quad (6)$$

This gives us the velocity vector

$$\mathbf{v} = (v \cos \theta, v \sin \theta). \quad (7)$$

Similarly, we obtain the position vector by integrating the velocity components

$$\mathbf{p} = \left(\int_0^T v \cos \theta dt, \int_0^T v \sin \theta dt \right). \quad (8)$$

The final component of the state vector is the acceleration which is given by the derivative of the velocity

$$\mathbf{a} = \left(\frac{d}{dt} v \cos \theta, \frac{d}{dt} v \sin \theta \right). \quad (9)$$

3. IMPLEMENTATION

In Section 2, we work from the model of two cameras, one on either side of the vehicle. Early experimental work followed this model. However, we later migrated to a model where a single camera is used. The camera is positioned at the front of the vehicle looking at the ground just in front of the vehicle as it moves. This arrangement is shown in the picture in Figure 2.

The camera we use is an Axis 2100, which is capable of streaming JPEG compressed images at a rate of approximately 10 per second. It is configured to send images which are 320 pixels wide by 240 pixels tall. Each image is decompressed into a one byte-per-pixel greyscale representation.

After decompression, we apply a correction to remove the barrel (fisheye) distortion. We use the method of Perš and



Figure 2: Test Vehicle

Kovačić[5]. For each pixel in the corrected image, we determine its distance from the center, r_l , and compute the corresponding position in the original (distorted) image by

$$r' = f \ln \left(\frac{r_l}{f} + \sqrt{1 + \frac{r_l^2}{f^2}} \right)$$

where f is the focal length of the lens in the same units as r_l . We use bilinear interpolation to find the pixel value in the new image. Figure 3 shows an example of such an image after distortion removal.

Next we extract two sub-images, each 100×100 pixels. The two images are positioned at the extreme left and right edges of the original image and are centered vertically. These two sub-images serve as the left and right cameras.

In order to find the velocity magnitude, we need to determine the relative shift between successive images on each side. We use the orientation correlation technique of Fitch, et al[2] to find this shift. Their technique involves creating a new pair of images in which each position is a complex number giving the direction of the gradient at that position of the original image. For image I_1 , we compute

$$g_1(x, y) = \text{sgn} \left(\frac{\partial I_1(x, y)}{\partial x} + i \frac{\partial I_1(x, y)}{\partial y} \right)$$

where

$$\text{sgn}(x) = \begin{cases} 0 & \text{if } |x| = 0 \\ \frac{x}{|x|} & \text{otherwise} \end{cases}$$

and where we approximate the partial derivatives by

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x-1, y) - f(x+1, y)}{2}.$$

Likewise, we compute g_2 from image I_2 . The correlation is then computed quickly using Fast Fourier Transforms

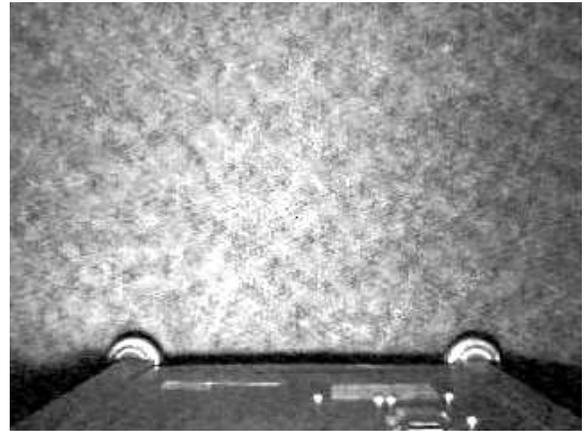


Figure 3: Example Image After Distortion Correction

(FFTs). If $G_1(x, y)$ is the FFT of $g_1(x, y)$ and $G_2(x, y)$ is the FFT of $g_2(x, y)$, then

$$\Re\{\text{IFFT}(G_1(x, y)G_2^*(x, y))\}$$

gives what Fitch, et al call the orientation correlation matching surface. The location of the peak in this surface gives the shift amount in the x and y directions. (We use the FFTW[3] library to compute the FFTs and inverse FFTs.) From the resulting shifts, we compute the camera velocities by Equation (1). Then, we compute the velocity magnitude by Equation (3), the orientation by Equation (6) and the other state vector components using Equations (7)–(8). In numerically evaluating the integrals of Equations (6) and (8), we use trapezoid rule integration. (We do not compute the acceleration in the implementation reported here.)

4. CALIBRATION

There are two physical parameters in our model that must be calibrated. The first is the parameter s in Equation (1), and the second is the camera separation parameter W in Equation (6). In principle, both of these can be measured directly from static images. However, we can get greater precision by moving the vehicle in a known trajectory and using the results to calculate backwards to the correct parameter values.

To complicate matters, the previous two sections assume ideal camera placement and orientation. In a real configuration, however, there are errors in placement and orientation. Errors in absolute placement of the camera do not affect the results. The fact that the camera in our experiments is placed forward of the vehicle's center of gravity only means that we track the position of the camera rather than the position of the vehicle. However, because the relative position and orientation of the camera to the vehicle is fixed, a simple transformation gives us the orientation and position of the vehicle.

Errors in camera orientation do affect the results. Consider first an error in the orientation of the camera about the axis parallel to the direction of motion. If the camera is rotated about this axis, the line from the camera to the center of its imaging area will not be normal to the ground. As a result, there will be a perspective effect that differentiates the left sub-image from the right. For example, if

in the second column is the total odometry $d = \sum v\Delta t$. The final orientation θ is given in the third column and the error in orientation is shown in the fourth column. This error is expressed as a percentage of the total change in orientation, i.e. $|\frac{\theta-2\pi}{3\pi}|$. Similarly, the final position is reported in the fifth column. The sixth column shows the error as a percentage of d .

The last two rows of the table summarize the results. In the next to last row, we present the means of the error columns. These represent the amount of error that can be expected on average for a path that is about 85 meters long with six right-angle turns. In the last row, we present the means of the final odometry values and compute the errors on those means. We must be careful in interpreting these data. The 0.2% error in the average orientation represents the orientation error we would expect if we ran the ten trial laps consecutively. This figure suggests that the orientation error is random, with a mean close to 0. On the other hand, the average (x, y) position, with its 2.0% error, does not represent the error we would expect after ten consecutive laps. Instead it suggests the error that we would get if we used an orientation reference such as a magnetic compass to correct the orientation error once per lap.

6. PERFORMANCE ISSUES

As reported in the previous section, we moved at approximately 0.30m/s during testing. This speed is well below the maximum speed for this technique. Our calibration results show that each camera pixel covers approximately 0.0026m. In the calculation of image shift, we interpret shifts greater than one-half of the image to be negative shifts. Using 100×100 pixel images, this limits the motion to no more than 50 pixels, or 0.13m, per image. Since we capture approximately 10 images per second, we can travel no faster than 1.3m/s, or 4.68k/h. Allowing a factor of two margin for error, we could expect the technique to work up to 0.65m/s or 2.43k/h. To support faster speeds, we would need to increase the frame rate, increase the number of meters per pixel or increase the size of each frame.

The computational load imposed by this technique is primarily dependent on the size of the images. If we use an $n \times n$ image, then each image has n^2 pixels. Because the FFT is an $O(n \log n)$ computation and since the FFTs dominate the computation, the tracking algorithm has $O(n^2 \log n)$ complexity. While we haven't carried out extensive studies on computational load, we have observed that using 100×100 pixel images we used approximately 20% of the CPU resources on a Compaq laptop computer using a 2.4GHz Pentium IV processor running Linux. Similarly, in earlier experiments with two cameras and 240×240 pixel images, the CPU was loaded to the point where images had to be dropped at times. While such a load is significantly higher than that of other localization techniques, it is still well within the capabilities of modern hardware.

It would be appropriate to compare the accuracy of image-based localization with other techniques. We have not yet done so experimentally. However, we can draw some rough comparisons to other reported results. In [1], the most applicable comparison is to the reported characteristics of the ENV-O5S Gyrostar inertial navigation system. They report

a positional drift of between 1 and 8cm/s and an orientation drift of $0.25^\circ/s$. This orientation drift over our 85m course at 0.3m/s would result in an error of 1.24rad or 13%. However, Borenstein, et al also report on a drift compensation that reduces the error rate to $0.0125^\circ/s$ which would result in a 0.7% error. It is difficult to directly determine the implications for the positional drift on our circuit. However, over the approximately 283 seconds that it takes to complete the course, the inertial navigation system would accumulate between 2.83m and 22.64m of drift which could represent as much as 3.3% to 26.6% positional error.

7. CONCLUSION

We have presented a method for tracking a vehicle's motion in order to augment absolute positioning systems. The technique can be used to fill in the gaps during GPS outages and during times when magnetic compasses do not provide reliable data. Initial results look promising.

Several additional avenues for further research are apparent. First, we would like to investigate the sources of error with an eye toward improving the system's accuracy and toward creating a more automatic calibration process. At this point, there is little data on how the error relates to vehicle speed and to turning radius. A second area of development aims at advancing the motion inference technique to support full-sized road vehicles. The camera used in these experiments substantially limits the speed at which the vehicle can travel. The third area of further research is combining this vision-based technique with other localization techniques. Finally, we seek to adapt the technique to allow motion inference from images taken from a forward-looking camera. This would allow us to piggy-back on infrastructure already in place for other purposes (e.g. obstacle detection and landmark identification). Imaging in a forward direction would also place the imaged ground farther from the camera resulting in a larger meters per pixel scaling factor which, in turn, would allow for higher velocities.

8. REFERENCES

- [1] J. Borenstein, H. Everett, L. Feng, and D. Wehe. Mobile robot positioning: Sensors and techniques. *Journal of Robotic Systems*, 14(4):231–249, 1997.
- [2] A. Fitch, A. Kadyrov, W. Christmas, and K. J. Orientation correlation. In P. Rosin and D. Marshall, editors, *British Machine Vision Conference*, volume 1, pages 133–142, 2002.
- [3] M. Frigo and S. G. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proc. 1998 IEEE Intl. Conf. Acoustics Speech and Signal Processing*, volume 3, pages 1381–1384. IEEE, 1998.
- [4] R. Negenborn. Robot localization and Kalman filters. Master's thesis, Utrecht University, 2003.
- [5] J. Perš and S. Kovačič. Nonparametric, model-based radial lens distortion correction using tilted camera assumption. In H. Wildenauer and W. Kropatsch, editors, *Proceedings of the Computer Vision Winter Workshop 2002, Bad Aussee, Austria*, pages 286–295, 2002.