A Tool for Range Sensing and Environment Discovery for the Blind

D. Yuan and R. Manduchi
Department of Computer Engineering
University of California, Santa Cruz
Santa Cruz, CA 95064

Abstract

This paper describes the development of a hand-held environment discovery tool for the blind. The final device will be composed of a laser-based range sensor and of an onboard processor. As the user swings the hand-held system around, he/she will receive local range information by means of a tactile interface. In addition, the time profile of the range will be analyzed by the onboard processor to detect environmental features that are critical for mobility, such as curbs, steps and drop-offs. In our current implementation, range is collected by a short-baseline triangulation system formed by a point laser and a miniaturized camera, producing readings at frame rate. An Extended Kalman filter is used to track the range data and detect environmental features of interest.

1. Introduction

There are an estimated 200,000 totally blind individuals in the US [10]. Approximately five times this number are legally blind, but they still have some usable vision left which is not correctable by standard eyeglasses. They experience difficulty performing visual tasks, because of reduced acuity or field of view[1].

Losing one's vision can be devastating for a number of reasons. In particular, one's ability to move about independently may be seriously affected. It has been estimated that more than 30% of the blind population do not ambulate autonomously outdoors [8]. Given the importance of mobility in one's daily life, technology that enables visually impaired persons to walk with confidence even in unfamiliar environments has a very significant social impact.

Currently, the most used mobility device for the blind is the long cane. The cane is economical, reliable and long–lasting; it allows one to extend touch and to "preview" the lower portion of the space in front of oneself. Not all visually impaired individuals, however, are willing or able to use the long cane. Indeed, the rigid cane being an "invasive" tool, it is ill–suited to social gatherings, on public or private transportation, and in congested areas where it may trap pedestrians.

Electronic Travel Aids (ETA) have been commercially available for more than thirty years. The most common and economical class of ETAs on the market includes simple hand–held devices that can be pointed toward any direction to give "clear path" information. An example is the Mowat ultrasound sensor [26], the most successful ETA to date [17]. ETAs may also be mounted on wheelchairs (such as the Wheelchair Pathfinder [6]) or embedded in long canes (such as the Nurion Laser Cane [4]). Other proposed mobility tools include the Sonic Pathfinder [5] and the SonicGuide [16].

An alternative family of devices aim to provide rich 3–D snapshots to the blind user. The proponents of this approach face the formidable challenge of encoding complex geometric and topological information into acoustic or tactile stimuli. A partial solution is given by the *sonification* technique, whereby spatialized sounds are synthetically generated to provide distance, azimuth and elevation cues to the listener [19, 7].[2]. These devices, however, are not practical as mobility aids, both because of the reduced spatial resolution, and because of the difficulty of interpretation of the sonified 3–D data.

In contrast to the sonification approach, the device concept described in this paper is inspired by the "long cane paradigm" for environment exploration[3]. Instead of using a physical cane, the visually impaired user will scan the scene with a hand–held non–contact range sensing device, integrating spatial information into a "mental image" of the scene. We believe that this approach has several advantages with respect to the sonification paradigm. One advantage is the smaller sensing and computational complexity (the system deals only with one–dimensional data) which translates into a more compact and economical device. Moreover, the

[1] Legal blindness refers to clinically measured visual acuity of 20/200 in the better eye with best correction, or visual field of 20° or less [10].

[2] One such system producing "soundscapes" has been recently commercialized under the brand name "The vOICe" [27].

[3] A system based on the "dog guide paradigm" was developed by Ulrich and Borenstein [23]. In this case, a blind user pushes forward a lightweight wheeled device (the GuideCane) equipped with ultrasound sensors. The GuideCane steers around obstacles, guiding the user in its path through the steering force felt in the handle.

information acquired by the sensor and processed by the onboard computer can be presented to the user through a simple haptic interface, which is less annoying and easier to interpret than complex sound structures. Note that the position and pointing direction of the device are "known" by the user via proprioception [6] (in the same way that a blind user "feels" the position and orientation of the cane he or she is handling). Finally, perception through exploratory movements appears to be a natural procedure for environment discovery. This is in accordance with Gibson's principle that perception and action are a cycle, whereby people act in order to learn their surrounding, and use what they learn to guide their actions [12, 13].

Measuring the distance to surfaces in the scene may help the user to take decisions about where to move, but is not sufficient for safe deambulation. Indeed, important features such as curb, steps, or drop–offs, cannot be easily detected by static range measurements. If, however, the user is allowed to swing the sensor around (for example, by a pivoting motion around an horizontal axis), the stream of measured range data may contain enough information to enable detection of such structures. As an example, the time profile of the range as the user pivots the device in an upward motion in front of different environmental features of interest is shown in Figure 1 for an ideal case (perfect pivot at constant angular rate around the optical center of the sensor, no range error)

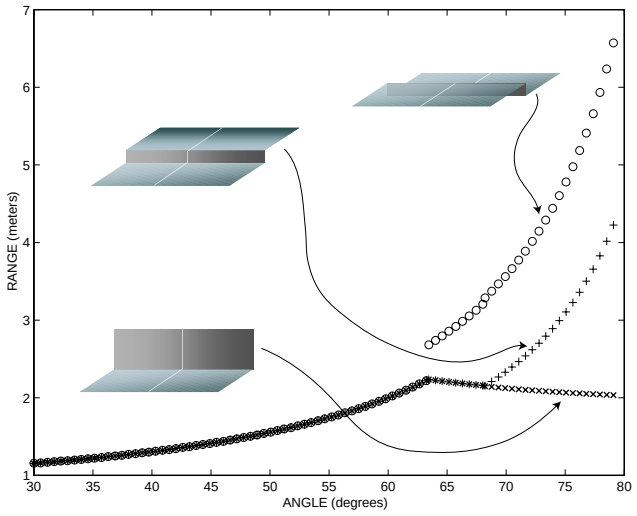

Figure 1: The range as a function of the pitch angle as the device is pivoted in an upward motion in front of a wall, a step, and a drop–off. Ideal measurements are supposed to be taken at a regular angular period of 0.6°. The horizontal rotation axis is at 1 meter of height; the step is 20 cm high. The time profile of the range is represented by "$\times$" for the wall case, "$+$" for the step case, and by "$\circ$" for the drop–off case.

It is perhaps unrealistic to assume that the detection of these features (which, as seen in Figure 1, are manifested by differential singularities of the range profile) can be performed by a human, especially given the inherent latency of tactile interfaces. Instead, we propose to analyze the range data by an onboard processor *before* the data is presented to the user. The detected features can then be communicated by special haptic primitives (such as a small jolt).

Thus, the device would in many aspects work as a "virtual white cane", enabling the visually impaired user to "feel" the outside environment by non–invasive sensing. On the technical side, tools developed in the robotic field for the analysis of range data from a rotating lidar could be used for the detection of environmental features, as disscussed in this paper.

The proposed system is logically divided in three parts: range sensing; onboard range processing; and tactile interface. This paper reports progress in the first two components. In particular, we describe the design and implementation of the range sensor in Section 2, and preliminary algorithms for the detection of environmental features from the time profile of the range in Section 3. Section 4 has the conclusions and future research directions.

2. The Range Sensor

The range sensor for our "virtual white cane" should provide measurements between 0.5 m (right beyond the reach of one's arm) and 3–4 m, with a resolution of a few centimeters at the maximum distance. Since we are only interested in point measurements, active triangulation systems are well suited to these requirements. Lidars, which are used extensively (and successfully) for robotic navigation, are too expensive and bulky for this application.

Ideally, the range sensor could be simply composed by a point laser paired with a linear CCD. This scheme, however, requires that its two components be perfectly aligned to ensure that the laser return will hit the array. Sophisticated mechanical calibration and a sturdy casing would thus be needed, which contribute to increasing the cost of the instrument. The requirement for precise alignment is relaxed with either one of these two schemes: a striper laser [20] coupled with a linear array, or a point laser coupled with a matrix array. The first solution has the advantage of a higher frame rate, and is inherently less expensive. However, given the spread of the laser stripe, less power is reflected back to the sensor, and reliable return detection may be difficult. A popular technique to increase the detectability of the laser return is to modulate the light power synchronously with the photoreceptor reading rate. For example, with a simple on–off modulation, frame differencing can be used for reliable detection in static scenes. Note, however, that in the case of highly dynamic scenes (or when the device is moved

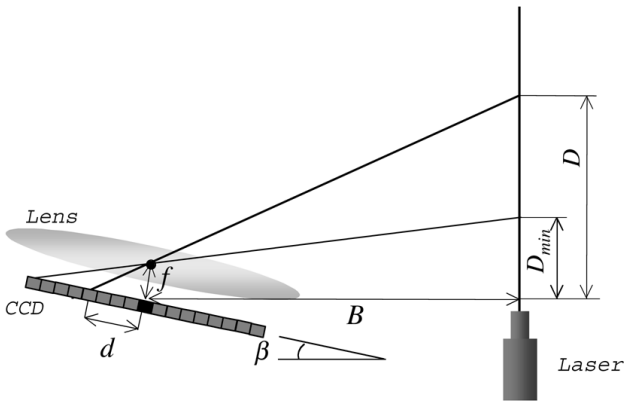

Figure 2: The layout of our laser triangulation system.

around, as in our case), frame differencing may yield noisy results, and more complex processing may be needed. In addition, since two frames are used for detection, the effective reading rate is one half the actual frame rate.

For our prototype, we opted for the second proposed solution (a point laser matched with a matrix CCD). In particular, we used a PointGrey Dragonfly greyscale Firewire camera [21], with a resolution of 1024×768 pixels, frame rate of 15 frames/s, and mounting a 6 mm microlens. The laser is an inexpensive module, with nominal divergence of 2 mrads and centerband at 650 nm. This laser is of class II, although we expect that a class I laser should work as well. Note that the sensor is expected to be used mostly for measurements at the ground level (much like a white cane), thereby reducing the opportunity of undesirable eye contact with the laser beam. An optical filter sheet was glued on top of the camera lens to reduce the effect of ambient light outside the laser's bandwidth. Details about the physical layout, detection algorithm, and performances, are presented next.

2.1. Physical Layout

Figure 2 shows the side view of the system, seen from a direction orthogonal to the plane identified by the laser beam and the focal center of the lens. Neglecting radial distortion, elementary geometry shows that the relationship between the distance D to a surface and the displacement d of the laser return on the linear sensor is:

$$D = \frac{fB\cos\beta - dB\sin\beta + df}{d\cos\beta + f\sin\beta} \tag{1}$$

where B is the baseline, f is the focal length , and β is the vergence angle. The baseline, the vergence, the focal length, the size of the detector array and the number of pixels in it determine the minimum measurable distance D_{min} as well as the range resolution $1/\Delta D$ at a certain maximum distance D_{max}. For a given vergence angle β, the minimum measurable distance sets an upper bound on the baseline:

$$B \leq \frac{D_{min}\left(h\cos\beta + 2f\sin\beta\right) - hf}{2f\cos\beta - h\sin\beta} \tag{2}$$

where h is the width of the CCD. The range resolution at a given distance D, corresponding to the inverse of the interval of distances subtended by one pixel, is:

$$1/\Delta D = \frac{f^2\sin\beta - fB - B\bar{d}\sin\beta\cos\beta - \bar{d}D\cos^2\beta + \bar{d}f\cos\beta}{\bar{d}\left(B\sin\beta + D\cos\beta - f\right)^2} \tag{3}$$

where $\bar{d}$ is the width of a pixel ($\bar{d} = 4.65\mu$m in our camera). Note that the resolution can be improved by subpixel accuracy (Section 2.2.3). Equations (2) and (3) determine constraints on B, f and β that must be satisfied given D_{min} and the resolution at D_{max}. In addition, we require that the laser return falls on the CCD for all distances $D \geq D_{min}$, even if we are mostly interested in the $[D_{min}, D_{max}]$ interval. This requirement is necessary to reduce the risk of false laser return detection. The constraint translates into:

$$\beta \leq \arctan\frac{h}{2f} \tag{4}$$

For the system design, we imposed that $D_{min} = 0.5$ m and $\Delta D \leq 100$ mm at $D_{max} = 3$ m before subpixel interpolation. The following parameters, used for our prototype, satisfy our specifications: $B = 80$ mm; $\beta = 11.1°$; $f = 6$ mm. A small value for the baseline was chosen to ensure a small form factor of the device.

2.2. Laser Return Detection

2.2.1. Determining the Epipolar Line

The epipolar line is the intersection of the image plane with the plane identified by the laser beam and the focal center of the system (epipolar plane) [15]. The laser return will always impinge on the epipolar line, so this is where detection is performed. As hinted earlier, the position of the epipolar line is very sensitive to mechanical misalignment of laser and camera, and therefore frequent re–calibration is necessary. We have implemented a simple self–calibration procedure that determines the epipolar line every time the system is turned on. A number of measurements are taken at different distances, each requiring an image–wide search for detection. From this pool of measurements, the RANSAC [15] algorithm is used to eliminate outliers. Finally, the epipolar line is determined by Total Least Squares fitting [14]. The Total Least Squares criterion, which minimizes the sum of the squared distances of the measurements to the fitting line, is the preferred tool for this type of problems (see e.g. [25].

This calibration procedure takes only a few seconds, and simply requires that the user moves the device around so that measurements at different distances are collected.

2.2.2. Light Compensation

For correct detection of the laser return, it is imperative that image saturation be avoided. Saturation may destroy the brightness profile information that allows the system to discriminate between the laser return and other sources of light in the scene. Since the device is meant to be moved around, the light conditions will continuously change, and fast adaptive light compensation is in order. As the microlens does not have iris control, the only available control parameters are the integration time and the gain of the camera. Our algorithm computes in each frame the number of pixels with brightness level between 252 and 255 (let's call it N_1), and the number of pixels with values between 248 and 251 (N_2). Ideally, $N_1 = 0$ and $N_2 > 0$, meaning that no pixel is saturated and yet the (almost) complete dynamic range is used. Based on the values of N_1 and N_2 computed for the n–th frame, the gain and integration time are modified for the (n+1)–th frame using a simple heuristic. Note that light compensation needs to be performed only in a neighborhood of the epipolar line.

2.2.3. Detection

The laser return is usually one of the brightest spots in the image, thanks in part to the optical filter in front of the camera lens and to the light compensation mechanism that reduces the chances of saturation. However, other image points may be brighter than the laser return. Typically, this occurs in correspondence of specular reflection from shiny surfaces, and/or when the laser is reflected by a dark surface. Hence, the brightness profile over a whole region must be analyzed for robust detection. A classic technique for waveform detection in a signal is matched filtering. Suppose the profile $g(x)$ of the brightness of the laser spot is known, and let $l(x)$ by the measured brightness. Then, the short–time matched filter is defined as follows:

$$F(x) = \frac{\int_{-T}^{T} g(y) l(x+y) dy}{\sqrt{\int_{x-T}^{x+T} l(y)^2 dy \int_{-T}^{T} g(y)^2 dy}} \qquad (5)$$

where the integral is computed on the epipolar line[4], and T is the half–size of the analysis window (it is assumed that the support $[-T_g, T_g]$ of $g(x)$ is small compared to the interval $[-T, T]$). The locations of the maxima of $F(x)$ are candidate positions of the laser return.

The brightness profile in the image in correspondence of the laser spot when reflected by a planar surface is well modelled by a Gaussian function. We found that a simple approximation to $g(x)$ by a box filter (that is, a binary function) works rather well. The advantage of the box filter is that it can be implemented efficiently as a running sum, requiring only $2(2T_g + 1)$ sums and one division per pixel for the implementation of the matched filter.

Unfortunately, the size of the laser spot on the image is not constant, as simple geometric reasoning shows. Assume for simplicity that the vergence angle is null and that the beam section at the output of the laser is negligible, and let α be the beam divergence. Suppose the intersection with the epipolar plane of a visible surface at distance D (as measured along the direction of the laser beam) forms an angle ϕ with the baseline. Then, the width of the projected spot on the image plane is

$$\Delta d = 2f \tan \tfrac{\alpha}{2} + \frac{fB}{D} \left(1 + \tan \frac{\alpha}{2} \tan \phi\right) \left(\frac{\cos\left(\frac{\alpha}{2} + \phi\right)}{\cos\left(\frac{\alpha}{2} - \phi\right)} - 1 \right) \qquad (6)$$

When the surface is parallel to the baseline ($\phi = 0$), the spot size is independent of the distance D. However, for slanted surfaces, Δd is a function of ϕ and D. In addition, we observed that light diffusion in certain material (e.g., human skin) determines larger spots than expected. For these reasons, we use three different matched filters, set at different scales. The parameters of these filters are: $T_g = 2, T = 5$; $T_g = 4, T = 8$; $T_g = 6, T = 11$, where T_g and T (defined above) are measured in pixels. The matched filter processing is only performed in the neighborhood of the pixels that passed the brightness criterion. In our implementation, we select the 10% brightest pixels in the small stripe around the epipolar line considered for search. All three matched filters $g_i(x)$ are applied to these pixels, and the pixel $\bar{x}$ that achieves the maximum value of $\max_i F_i(x)$ is the winner.

To improve the range resolution, the output of the matched filter is interpolated for subpixel accuracy. More precisely, the filter $g_i(x)$ that yields the maximum value for $\bar{x}$ is considered. A quadratic kernel is used for linear interpolation over a segment $[-N, N]$ of pixels around $\bar{x}$ (the size of this segment is chosen by thresholding $F_i(x)$ at 1/10 of its maximum). The maximizer of the interpolant can be obtained in closed form as

$$\bar{x}_{int} = \frac{\sum_{y=\bar{x}-N}^{\bar{x}+N} y l(y)}{\sum_{y=\bar{x}-N}^{\bar{x}+N} l(y)} \qquad (7)$$

In our current implementation, light compensation and laser spot detection are performed by a 2GHz laptop connected to the camera via Firewire. Range measurements are produced at frame rate (15 Hz) using a fraction of CPU time. To make the system portable, we are considering porting the code to a PDA; we expected that the computational weight of detection is well within the capabilities of a hand–held computer.

[4] We actually search in a set of ± 10 lines around the estimated epipolar line, to account for possible small calibration errors.

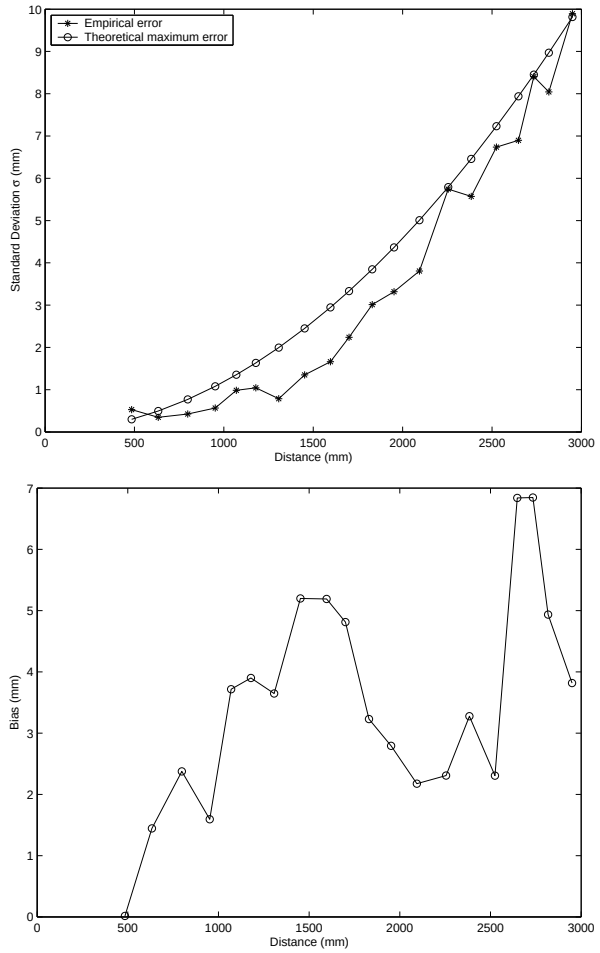

Figure 3: Top: The standard deviation of the empirical measurement noise as a function of distance, with subpixel interpolation, plotted together with the theoretical standard deviation $\Delta D/\sqrt{12}$. Bottom: The magnitude of themeasurement bias.

To experimentally assess the performance of the system, we took a number of measurements at different distances, and plotted our results in Figure 3.

The top part of Figure 3 shows the standard deviation of the measurement error as a function of distance. It is instructive to compare these results with the theoretical standard deviation of the error due to quantization, which, under the "high resolution" hypothesis [11] is equal to $\Delta D/\sqrt{12}$, where ΔD is derived in (3). This theoretical standard deviation is plotted together with the measured standard deviation (using subpixel interpolation and averaging over more than 500 measurements for each distance) in Figure 3. Note from the figure that subpixel interpolation effectively reduces the measurement variance with respect to the theoretical non–interpolated case. The magnitude of the error bias (averaged over all tests) is shown in the bottom part of Figure 3. This small bias (less than 1% of the real value) is probably due to inaccurate system calibration. Since we are mostly interested in the differential properties of the time profile of the measured range (as discussed in he next section), this bias is negligible for our purposes.

3. Environmental Feature Detection

An important function of the proposed device is the analysis of the the stream of range data acquired as the user moves the instrument around. The purpose of this analysis is the detection of features that are important for safe deambulation, such as curbs, steps and drop–offs. In general, these features are revealed by differential singularities of the corresponding range profile (such as cuspids, or discontinuities – see Figure 1).

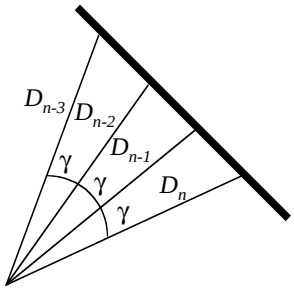

Figure 4: The measurement model for a planar surface, assuming constant angular velocity.

Assuming a simple motion model for the sensor, environmental features can be detected by a suitable tracking algorithm. Our approach was inspired by the work of Adams [2, 1], who considered the problem of detecting range discontinuities (corresponding to the edges of planar surfaces) from data produced by a rotating lidar. In this case it is possible to model the dynamics of the range data readings in front of a surface (see Figure 4) as follows:

$$D_{n+1} = \frac{D_n D_{n-1}}{2D_{n-1}\cos\gamma - D_n} \tag{8}$$

where D_n represents the distance to the surface at the n-th reading time, and γ is the angle (assumed constant) spanned by the lidar between two consecutive measurements. In [2], an Extended Kalman Filter (EKF) based on this model to detected range discontinuities by means of a validation gate equation [3].

Since indoor and urban environments are for the most part composed by planar structures, a similar detection technique could be used. A movement similar to that of a rotating lidar can be obtained by pivoting the device around an horizontal axis coincident with one's wrist. This is not dissimilar to the motion used to manouver a white cane in the

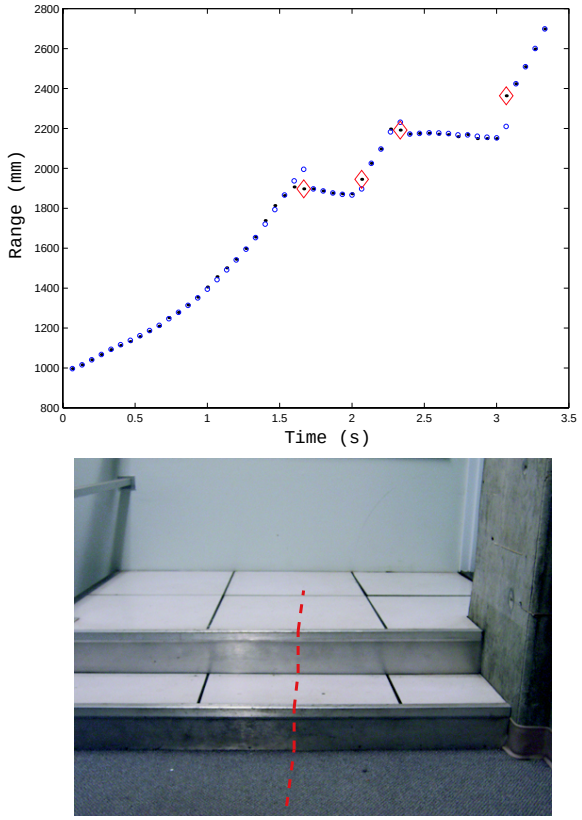

Figure 5: The time profile of range acquired as the device was pivoted in an upward motion in front of the scene shown in the bottom figure (the dashed red line indicates the trace of the laser). The sensor was at a height of about 50 cm from the floor. Black dots: original measurements. Blue circles: filtered range. Red diamonds: detected features.

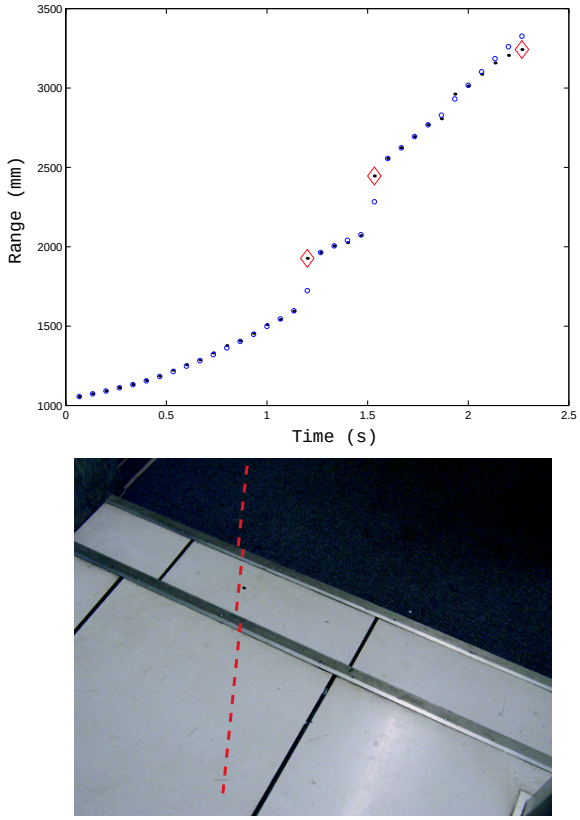

Figure 6: See caption of Figure 5.

so–called two–point touch technique [24]. Whereas the angular velocity during a pivot may be assumed to be approximately constant, its actual value is unknown, and therefore equation (8) cannot be used directly. A possible solution is to equip the device with a MEMS attitude sensor [9], which could provide direct measurements of the angular velocity. Another solution, which we implemented in our prototype, is to manipulate Equation (8) to obtain a model that does not require direct knowledge (or observation) of γ. This is achieved by considering an augmented state, composed by three consecutive values of range. Indeed, it is easy to show that, assuming constant angular velocity and a planar surface (see Figure 4), the following identity holds:

$$D_{n+1} = \frac{D_n}{\frac{D_{n-1}}{D_{n-2}} + \frac{D_{n-1}}{D_n} - \frac{D_n}{D_{n-1}}} \tag{9}$$

Thus, a system model can be defined with the following state vector at time n:

$$\mathbf{x}(n) = [D_{n-2}, D_{n-1}, D_n]$$

We assume a noiseless state update model, defined by

$$\begin{aligned} x_1(n+1) &= x_2(n); \\ x_2(n+1) &= x_3(n); \\ x_3(n+1) &= \frac{x_3(n)}{\frac{x_2(n)}{x_1(n)} + \frac{x_2(n)}{x_3(n)} - \frac{x_3(n)}{x_2(n)}} \end{aligned} \tag{10}$$

As observed in [2], a noisy state update model could be useful to account for uneven or not perfectly planar surfaces (or perhaps to compensate for involuntary translation of the device in addition to pivoting). We plan to investigate this further in future research.

The observation model is simply

$$z(n) = x_3(n) + w(n) \tag{11}$$

where $w(n) \sim N(b_n, \sigma_n^2)$ is the observation noise, b_n is the bias between real distance and its observation, and σ_n^2 is the variance. b_n and σ_n are both functions of distance. Since we are only concerned with the difference between consecutive observations, the influence of bias can be neglected. Therefore, without subpixel interpolation, the noise, as a function of distance, is better modeled by a uniform random variable with variance equal to $\Delta D^2/12$, where ΔD

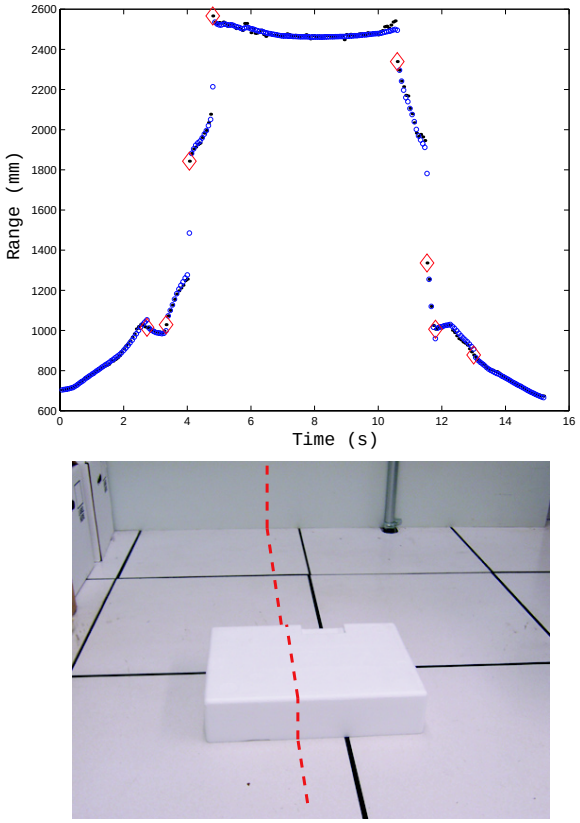

Figure 7: See caption of Figure 5. In this case, a downward pivot followed the upward motion.

is the inverse of the resolution as defined in Section 2.1. We will use this value as the basis for the variance of the noise $w(n)$. More precisely, we will assume that , when the current estimate of the distance is $\bar{D}_n$, the variance of the observation noise is equal to

$$\sigma^2(\bar{D}_n) = \Delta D(\bar{D}_n)^2/12 \qquad (12)$$

where $\Delta D(\bar{D}_n)$ is the inverse of the resolution as in (3)(see Figure 3 and related discussion).

State prediction at time $n+1$ is obtained using an Extended Kalman Filter, which requires linearization of the update equation (10) about the previous prediction. As in [2], features that correspond to the edge of a planar patch are detected by thresholding the normalized innovation, which is equal to the square of the innovation (the difference between the observation and the predicted observation) divided by the variance of the innovation. This *validation gate* approach is very effective at finding outliers, that is, measurements that do not satisfy the model [3]. When a trail of consecutive measurements are detected as outliers, only the first one is labeled as a feature. Every time there is a detection, the state is re–initialized.

Some results of range tracking and feature prediction are shown in Figures 6–7 for scenes containing both steps and drop–offs. In all cases the instrument was pivoted upwards around an horizontal axis; for the case of Figure 7, a downward pivot followed the upward one. The time profiles of the range should be compared with the theoretical curves shown in Figure 1. Although these are only preliminary results, obtained in controlled situations and with relatively slow angular velocity (approximately 7°/s), it is seen that the detection of features corresponding to the edges of planar patches is rather reliable. More experiments will be need to assess the performances of the system in different conditions, and to fine–tune the expression for the error variance (12) as well as the validation gate.

4. Conclusions and Future Work

We have described a prototype sensor built in our laboratory that produces reliable local range measurements using active triangulation. This sensor will be part of a hand–held mobility device for use by the visually impaired. In addition, we have presented an algorithm, based on Extended Kalman Filter tracking, that can detect environmental features corresponding to curbs, steps, and drop–offs. Preliminary experiments in laboratory environment have given promising results.

We are currently exploring several research directions to improve the performances of the current system. For what concerns the sensor, we will consider implementing a layout with a laser striper coupled with a linear detector, to be tested against the current point laser – matrix CCD scheme. Since linear detectors provide a much higher frame rate than their matrix counterpart, this would enable reliable feature detection when the device is pivoted at a high angular velocity. We are also considering more sophisticated algorithms for feature detection based on Multiple Hypothesis Tracking [18, 22], whereby a bank of Kalman filters are used, each representing a different planar model. We expect that this approach will provide more robust feature detection, as well as structural information necessary to discriminate between different geometric features.

References

[1] M.D. Adams. On–line gradient based surface discontinuity detection for outdoor scanning range sensors. In *Proc. 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '01)*, pages 1726–31, Maui, HI, 2001.

[2] M.D. Adams and A. Kerstens. Tracking naturally occurring indoor features in 2–D and 3–D with lidar range/amplitude data. *The International Journal of Robotics Research*, pages 907–23, 1998.

[3] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.

[4] J.M. Benjamin, N.A. Ali, and A.F. Schepsis. A laser cane for the blind. In *Proc. San Diego Biomedical Symposium*, volume 12, pages 53–7, 1973.

[5] D. Bissitt and A.D. Heyes. An application of biofeedback in the rehabilitation of the blind. *Applied Ergonomics*, 11(1):31–3, 1980.

[6] B.B. Blasch, W.R. Wiener, and R.L. Welsh. *Foundations of Orientation and Mobility*. AFB Press, 1997. Second Edition.

[7] S.A. Brewster. Using non–speech sounds to provide navigation. *ACM Trans. on Computer–Human Interaction*, 5(3):224–259, 1995.

[8] D.D. Clark-Carter, A.D. Heyes, and C.I. Howarth. The effect of non–visual preview upon the walking speed of visually impaired people. *Ergonomics*, 29(12):1575–81, 1986.

[9] D.Gebre-Egziabher, G.H.Elkaim, J.D. Powell, and B.W. Parkinson. A gyro-free quaternion-based attitude determination system suitable for implementation using low-cost sensors. In *Proc. IEEE Position Location and Navigations Symposium (IEEE PLANS)*, pages 185–192, San Diego, 2000.

[10] American Foundation for the Blind. http://www.afb.org.

[11] A. Gersho and R.M. Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publisher, 1992.

[12] J.J. Gibson. *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston, MA, 1966.

[13] J.J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA, 1979.

[14] P. De Groen. An introduction to Total Least Squares. *Nieuw Archief voor Wiskunde*, 14(12):237–254, 1996.

[15] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, 2000.

[16] L. Kay. A sonar aid to enhance spatial perception of the blind: Engineering design and evaluation. *Radio and Electronic Engineer*, 44(1):605–27, 1974.

[17] R. Ludt. Western Blind Rehabilitation Center. Personal communication.

[18] D.T. Maggil. Optimal adaptive estimation of sampled stochastic processes. *IEEE Transactions on Automatic Control*, AC-10(4):431–439, 1979.

[19] P.B.L. Meijer. An experimental system for auditory image representations. *IEEE Transactions on Biomedical Engineering*, 38(2):112–121, February 1992.

[20] C. Mertz, J. Kozar, J.R. Miller, and C. Thorpe. Eye-safe laser line striper for outside use. In *IV 2002, IEEE Intelligent Vehicle Symposium*, 2002.

[21] Inc. Point Grey Research. http://www.ptgrey.com/.

[22] S.I. Roumeliotis, G.S. Sukhatme, and G.A. Bekey. Sensor fault detection and identification in a mobile robot. In *Proc. 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1383–1388, Victoria,Canada, 1998.

[23] I. Ulrich and J. Borenstein. The GuideCane – Applying mobile robot technologies to assist the visually impaired. *IEEE Trans. on Systems, Man, and Cybernetics – A*, 31(3):131–6, March 2001.

[24] R. Wall. Biomechanical substrates of the two–point touch cane technique: A review of research. *AFB Journal of Visual Impairment and Blindness*, 96(2), 2002.

[25] J. Weber and J. Malik. Robust computation of optical flow in a multi-scale differential framework. *International Journal of Computer Vision*, 14(1):5–19, 1995.

[26] Wormald International Sensory Aids, 6140 Horseshoe Bar Rd., Loomis, CA 95650.

[27] http://www.seeingwithsound.com/.