

Detecting water hazards for autonomous off-road navigation

Larry Matthies*, Paolo Bellutta, Mike McHenry
Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

ABSTRACT

Detecting water hazards for autonomous, off-road navigation of unmanned ground vehicles is a largely unexplored problem. In this paper, we catalog environmental variables that affect the difficulty of this problem, including day vs. night operation, whether the water reflects sky or other terrain features, the size of the water body, and other factors. We briefly survey sensors that are applicable to detecting water hazards in each of these conditions. We then present analyses and results for water detection for four specific sensor cases: (1) using color image classification to recognize sky reflections in water during the day, (2) using lidar to detect the presence of water bodies and to measure their depth, (3) using short-wave infrared (SWIR) imagery to detect water bodies, as well as snow and ice, and (4) using mid-wave infrared (MWIR) imagery to recognize water bodies at night. For color imagery, we demonstrate solid results with a classifier that runs at nearly video rate on a 433 MHz processor. For lidar, we present a detailed propagation analysis that shows the limits of water body detection and depth estimation as a function of lookahead distance, water depth, and lidar wavelength. For SWIR and MWIR, we present sample imagery from a variety of data collections that illustrate the potential of these sensors. These results demonstrate significant progress on this problem.

1. INTRODUCTION

Perception systems for autonomous navigation of unmanned ground vehicles (UGVs) must detect all manner of potential navigation hazards. Until quite recently, UGV perception research had focused largely on perception of 3-D terrain geometry with range sensors [1], perception of roads [2], and to a lesser degree on perceiving other traffic [3]. Research effort is now growing on recognizing vegetation to improve the efficiency of off-road navigation [4,5]. However, there still has been very little work on detecting bodies of water that could be navigation hazards or on estimating the depth of potential water hazards. This paper presents first results on this problem, using a variety of sensors including color imagery, lidar, short-wave infrared imagery, and thermal imagery.

We start by cataloging environmental variables that affect the difficulty of the water detection problem and by distinguishing conditions under which different sensors apply. The first variable is day versus night operation. Daylight obviously enables the use of reflected solar illumination, whereas night operation requires thermal imagery or active sensors, particularly lidar. A second variable is whether the body of water reflects the sky or other aspects of the terrain, such as trees, hills, or buildings. As we will show, water reflecting the sky during the day is a relatively easy case. At the other extreme, bodies of water completely roofed over with tree canopy may be particularly difficult. A third variable is the size of the water body; this matters because size will affect the thermal contrast between the water body and surrounding terrain, which matters for night operation. Other variables include surface wave state, water depth, water turbidity, and presence or absence of scum or other material on the water surface. Among other consequences, these variables affect the ability to measure water depth by optical means.

A wide variety of sensors are applicable to this problem. During the day, color image classification can be used to recognize water by its reflection of the sky; in off-road, open terrain, this is a fairly reliable, easy-to-compute signature. When still water reflects other structure, it may be possible to recognize the mirror-image structure of the scene. Moreover, stereo vision often will match on the reflections and give range to the reflected objects. This will produce anomalously large range measurements, which can be interpreted to reveal water or at least suspicious terrain of some kind. The robotics research community has recognized for some time that lidar tends not to give returns from water bodies and that this might be used to detect water bodies; however, this phenomenon has not been analyzed in detail. Our exploration of lidar/water phenomenology pointed out two other possibilities: (1) appropriate lidar wavelengths might enable sensing the depth of water bodies and (2) passive imaging at appropriate near-infrared (NIR) or short-wave infrared (SWIR) wavelengths may reveal water hazards due to the strong absorption of light at those wavelengths. SWIR is also applicable to detecting snow and ice for the same reason. Thermal infrared imagery is applicable because large bodies of water will generally be cooler than the surrounding terrain during the day and warmer at night; this is

* lhm@robotics.jpl.nasa.gov; phone 1-818-354-3722; fax 1-818-393-4085

especially relevant at night, when visible and SWIR signatures are unavailable. Finally, polarization of reflections from water is a well-known phenomenon in everyday life and has potential use for robotics. This will be complicated by the fact that light incident on water will already be partly polarized in context-dependent ways. Skylight is polarized as a function of the sun position and the clarity of the atmosphere [6, 7]. Light that reflects one more times off other terrain structure before hitting the water may also be polarized, potentially in randomly varying orientations.

Our work on this problem has progressed from the most common sensors and environmental conditions to less common sensors and conditions. In section 2, we present relatively mature results for day-time water detection with color imagery, including a discussion of how to use range data from stereo to augment color and how to interpolate range data from around the edge of the water body to place it in an onboard, local map. This has all been implemented in real-time on a robotic vehicle. In section 3, we present an in-depth analysis of lidar propagation for detecting water and sensing water depth; we have not yet incorporated this analysis into an onboard perception system. Section 4 explores the possibility of using SWIR attenuation for detecting water, snow, and ice. The results are very promising for snow and ice; for water, the technique works for relatively low angles of incidence, but surface reflections dominate at high angles of incidence. Section 5 uses imagery from 24-hour mid-wave thermal infrared (MWIR) observations of a reservoir to illustrate the potential for using thermal infrared imagery for water detection at night. Initial results are promising, but more work is needed to complete a robust algorithm. Section 6 summarizes what aspects of the overall water detection problem are now reasonably well solved and where open problems remain.

2. COLOR IMAGERY

In our experience, for many, if not most off-road conditions, the reflections of sky in water are easily discriminated from other terrain by their color and brightness, regardless of whether the sky is clear, partly cloudy, or completely overcast and regardless of the surface wave state of the water. We demonstrate this in section 2.1. After water bodies are detected in an image, they must still be placed in a local map to be useful for route planning; however, currently available sensors do not provide range data to the water surface itself, so this must be inferred from range data around the boundary of the water body. Section 2.2 shows initial results from a simple algorithm for interpolating range over the surface of a water body. When other terrain structure is reflected in the water surface, it is difficult to recognize this with color alone. In section 2.3, we show that stereo can produce range data to such reflections and suggest how to use such range data to correctly label the reflections as water.

2.1 Detecting Sky Reflections

Figure 1 shows six images from a data collection obtained under the Army Research Lab "Demo III" program at Fort Indiantown Gap, PA, that include a mixture of off-road terrain and water bodies. We used these images to evaluate the brightness [$I = (R+G+B)/3$] and color saturation [$1 - \min(R,G,B)/I$] of the terrain, water, and sky for manually segmented images; Table 1 shows the resulting means and standard deviations of these regions. In many of these images, the sky fully saturates the dynamic range of the imagery, which is why the sky mean brightness is so close to 255 and the standard deviation is so low. In our experience, it is common for sky to saturate the dynamic range of outdoor imagery. Note that the sky brightness nevertheless is still two and half times higher than the mean brightness of the terrain. The average brightness of the water, where it reflects the sky, is mid-way between that of the sky and the terrain.

Under the Demo III program, we had already developed a color image classifier based on supervised classification with a mixture of Gaussians model that runs at 40 ms/frame on a 433 MHz PowerPC; therefore, the simplest and most uniform approach for us to exploit the results in Table 1 was to train this classifier on water regions in RGB color space. We tested this under the DARPA Perceptor program with a 24-hour image sequence collected at the Chatfield Reservoir in Denver. Sky conditions for this data set ranged from partly cloudy to completely overcast; water surface conditions ranged from very still to completely rippled. Figure 2 shows results for one case of still water and one case of rippled water. For still water, the classifier correctly labelled essentially all pixels that reflected the sky, but labelled pixels reflecting vegetation as vegetation. For rippled water, essentially all pixels that were water were correctly labelled. At the time this data was processed, we were not concerned about correctly labeling sky regions as such, so there is some misclassification of sky as water; however, this can be corrected easily by using knowledge of the camera attitude to find the horizon line and to correct the labels of pixels above the horizon.



Figure 1: Images from Fort Indiantown Gap used for class statistics in Table 1.

	Brightness mean	Brightness std. dev.	Saturation mean	Saturation std. dev.
Sky	253.4	4.0	0.0067	0.014
Water	176.0	37.6	0.21	0.091
Rest of image	106.2	27.0	0.29	0.10

Table 1: Brightness and color saturation region statistics for imagery in Figure 1.

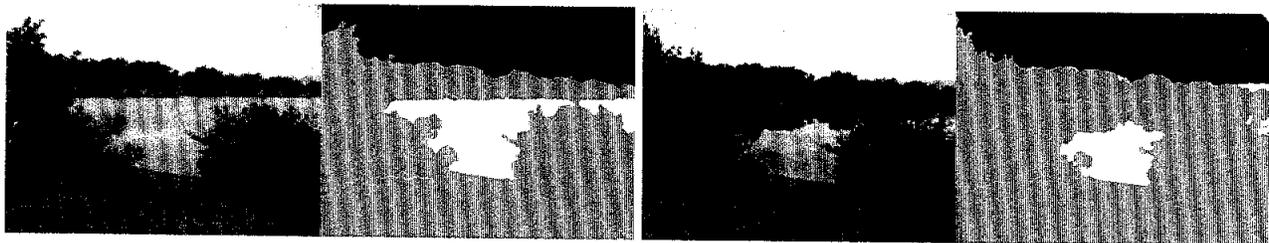


Figure 2: Color classification results for rippled (left) and still (right) water, from a sequence of images collected every 10-15 minutes over a 24 hour period at the Chatfield Reservoir in Denver. White: water; brown: soil; green: vegetation; blue: other. In the rippled case, all of the water is correctly labelled. In the still case, water reflecting the sky is correctly labelled, but water reflecting trees is labelled as vegetation.

2.2 Placing Water Bodies in the Map

To place water regions in the onboard map, we interpolate range data over the water from range data around its boundaries. This is subject to potential elevation errors if the water edge is occluded by raised objects, such as the bushes in Figure 2. The elevations of border pixels can be analyzed in a variety of ways to attempt to overcome this problem; for example, by taking the minimum elevation over all pixels bordering the bottom edge of the water region in the image or by more sophisticated analysis of the histogram of border elevations. Once an elevation is determined for the water surface, it is straightforward to interpolate range data for each pixel up each column of the image over the whole water region, starting at the bottom of the water region in each column. Figure 3 shows results of such interpolation for a still water case.

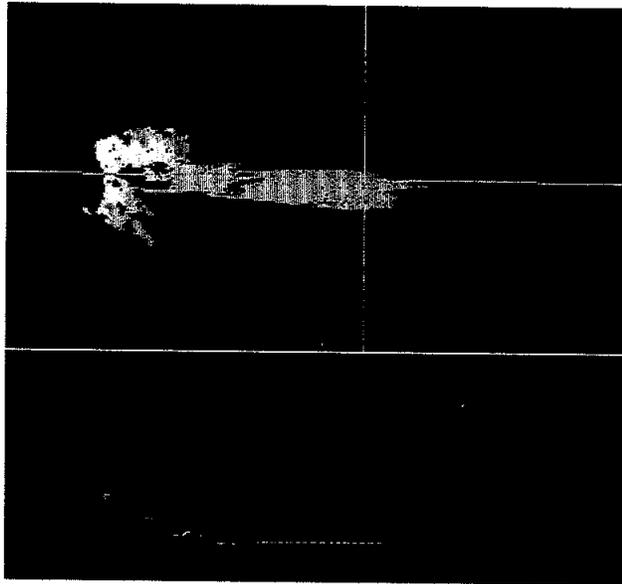


Figure 3: Results of interpolating range data for the water body from Figure 2. Top: false color elevation plot, viewed from above; camera vantage point was from the left. Reddish-brown is lowest; yellow, green, blue are progressively higher elevations. The water is at the lowest elevation. Bottom: height vs. range plot through the horizontal cross-hair in the top view. The water surface is the level portion on the right.

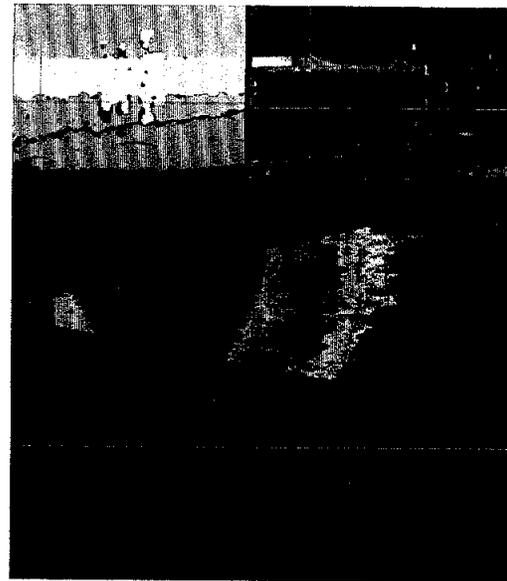


Figure 4: Pond of still water with stereo matching on the reflections. Upper left: false color range image; reddish is closest, green is furthest. Reddish portion is the near shore; yellow portion contains the far shore and its reflection. Upper right: left image of the stereo pair. Center: elevation plot. The large empty area is where the water really lies. Bottom: height vs. range plot for the cross-hair line that is vertical in the intensity image and near horizontal in the elevation plot. Note that the range to the reflections roughly match the range to the actual objects and plots much below the true ground level.

2.3 Using stereo to label terrain reflections as water

As Figure 2 shows, where the water surface is rippled its color and brightness will often allow it to be labelled as water; however, where the surface is still this is more difficult. Since reflections on still water are often distinct enough to enable stereo matching, the resulting range data provides a promising approach to correctly labeling such regions. Figure 4 shows a distinctive case, where the entire surface of a pond reflects soil, trees, and man-made structures on the opposite shore. In the height vs. range plot in the bottom of the figure, it is clear that the range data on the reflections is below the nominal ground surface, that it wraps backward in range, and that there is a large gap in the range data in front of it. These features are characteristic of water reflections or, in less extreme examples, of negative obstacles. In this example there is no sky reflection, but if there were sky reflection in the foreground, then the initial segment of labelled water region would even more strongly indicate that the following range data anomaly was due to a reflection on more water. Finally, a reflection hypothesis based on range data can be checked by mirror-image correlation of the intensity image. We are in the process of implementing and testing these recognition techniques for reflections.

3. LADAR

It has been observed anecdotally for several years that water bodies tend not to produce returns for UGV-mounted ladar. This has generally been attributed to specular reflection at the air-water interface. However, deeper analysis shows that the situation is not so simple, because some of the ladar energy penetrates the interface and can even produce a range measurement to the bottom of the water body, depending on angle of incidence, ladar wavelength, and attenuation in the water column (Figure 5). In this section, we present a propagation model for the entire optical path to elucidate these factors; this enables us to predict conditions under which ladar can be used to detect water bodies and to measure their depth.

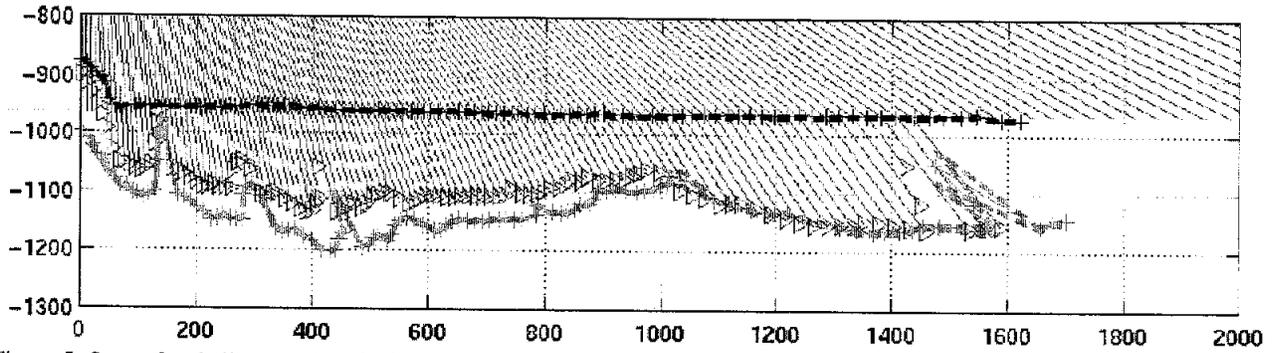


Figure 5: Scan of a shallow pond (units in mm) made with a SICK ladar with its scanning plane oriented vertically. The ladar was about 95 cm above the water level; the water was 15 to 20 cm deep. Red lines trace the path of each pulse in the scan. The blue line marks the water surface, as determined by holding a meter-stick at water level in the path of the scanner; the raised bit of blue line on the left was the experimenter's hand. The green symbols mark positions of raw 3-D measurements. The red triangles correct these for refraction at the air-water interface and for the speed of light in water, given knowledge of the location of the water surface; the red lines trace the path of each pulse. The water bottom was detected until the incidence angle at the air-water interface was about 56 degrees; at greater incidence angles, no return was recorded by the ladar. The ladar wavelength is 905 nm.

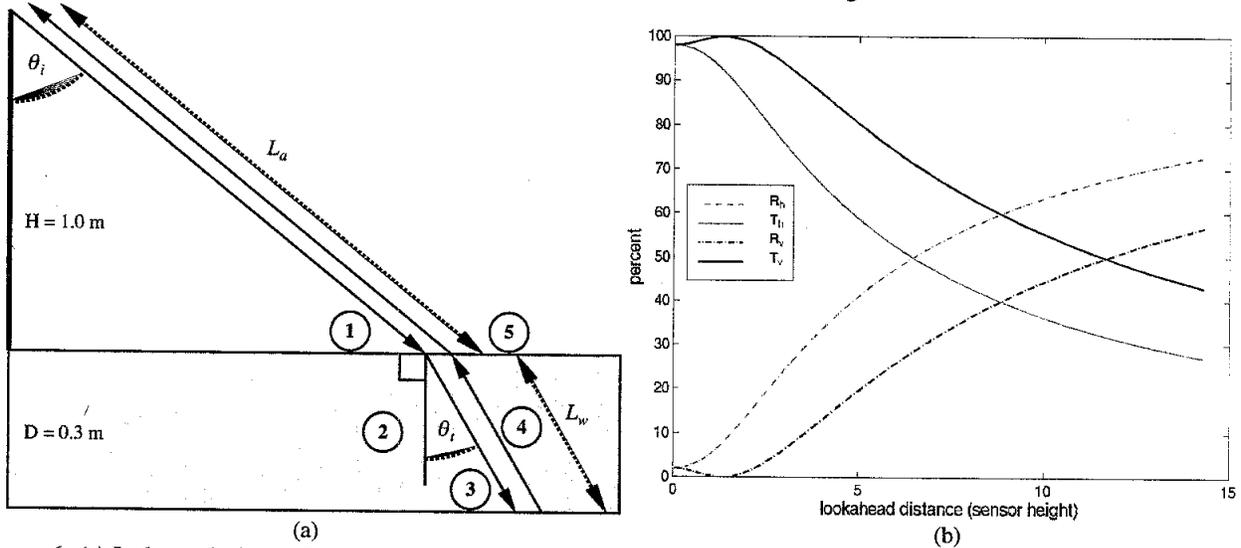


Figure 6: (a) Ladar optical path for sensing a water body, for an example with the sensor 1.0 m above a water body 0.3 m deep. Taking atmospheric attenuation as negligible at short range, the principal losses occur due to (1) reflection at the first air-water interface, (2) attenuation through the water column, (3) diffuse reflection on the water bottom, (4) attenuation back through the water column, and (5) reflection at the second air-water interface. L_a and L_w denote the one-way path length through the air and water. (b) Reflected and transmitted power percentages through the first air-water interface for horizontal and vertical polarizations, as a function of lookahead distance in multiples of the sensor height (eg. interpret as meters for 1 m sensor height).

Figure 6a illustrates the optical path to be analyzed and shows the locations where principal light losses occur. Assuming that atmospheric attenuation is negligible at short range, the first loss is due to reflection at the air-water interface. The transmitted and reflected power for horizontal and vertical polarizations are given by the standard Fresnel equations [8] as:

$$\begin{aligned}
 T_v &= \left(\frac{n_w \cos \theta_t}{n_a \cos \theta_i} \right)^2 \\
 T_h &= \left(\frac{n_w \cos \theta_t}{n_a \cos \theta_i} \right)^2 \\
 R_v &= 1 - T_v \\
 R_h &= 1 - T_h
 \end{aligned} \tag{1}$$

where θ_i and θ_t are the incident and transmitted angles, $\theta_t = \sin^{-1}(1/n_w \sin\theta_i)$ by Snell's law, n_a and n_w are the indices of refraction on each side of the interface (1.0 and 1.33 for air and fresh water), and the amplitude coefficients t_v and t_h are:

$$t_v = \frac{2 \sin \theta_i \cos \theta_t}{\sin(\theta_i + \theta_t) \cos(\theta_i - \theta_t)}$$

$$t_h = \frac{2 \sin \theta_i \cos \theta_i}{\sin(\theta_i + \theta_t)}$$

The power transmitted through the first air-water interface is:

$$\Phi_1 = \Phi_{0h} T_h + \Phi_{0v} T_v \text{ Watts (W)} \quad (2)$$

where Φ_{0h} and Φ_{0v} are the incident power in horizontal and vertical polarizations, respectively. Figure 6b illustrates the power transmitted into the water as a function of lookahead distance for both polarizations. The subscripts here and below on Φ_i relate directly to the numbered losses in Figure 6a.

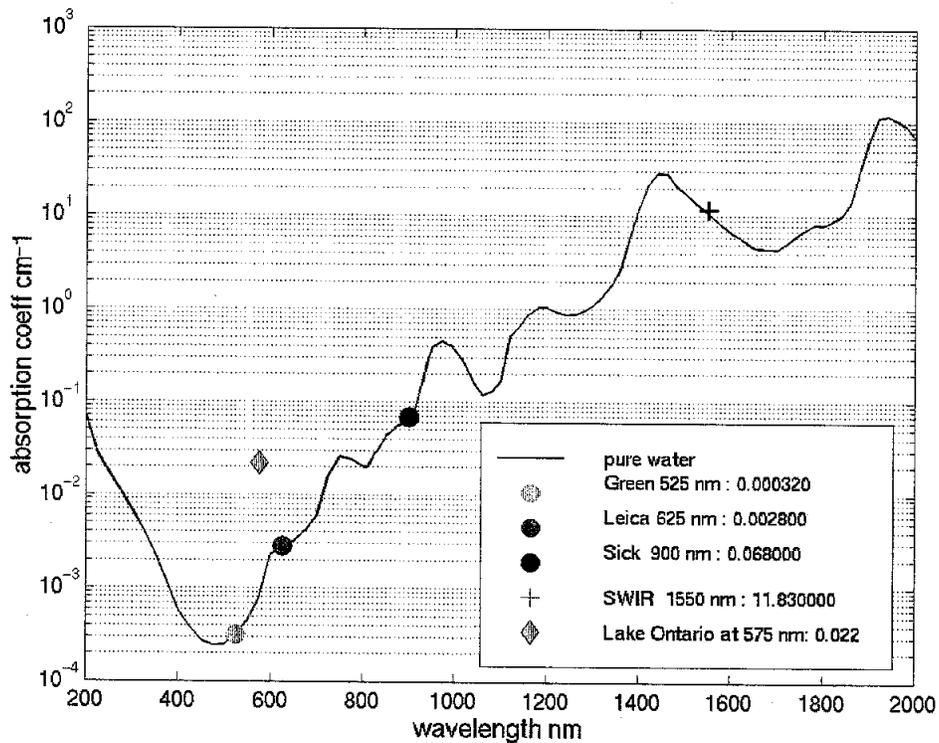


Figure 7: Absorption coefficient for pure water vs. wavelength, with specific values noted for the SICK lidar, a surveying laser range finder by Leica, a green laser pointer, and a common SWIR laser wavelength. For natural water bodies, data from [9] for coastal Lake Ontario shows the minimum total attenuation occurring around 575 nm, at the value shown in the legend.

Light propagation through water is subject to exponential attenuation according to [9]:

$$\Phi = \Phi_0 e^{-c(\lambda)x} \quad (3)$$

where c is the wavelength-dependent total attenuation coefficient and x is the path length in the water. c is further comprised of an absorption term a and a scattering term b , such that $c = a + b$. Figure 7 shows the absorption coefficient for pure water from 200 to 2000 nm. There is an order of magnitude jump from the minimum in the blue-green region to red, another order of magnitude to the near infrared (eg. the SICK wavelength), and two more orders of magnitude to the region around 1550 nm, which is a wavelength of strong interest for improving the maximum range and the eye safety of UGV ladars. This shows that a lidar operating in the green region will have far better propagation in water than the SICK, whereas ladars at 1550 nm will have far worse propagation. For example, based only on the pure water absorption coefficients in Figure 7, 99% of the initial power is lost after 144, 16, 0.68, and 0.004 m at 525, 625, 900, and 1550 nm,

respectively. Losses in natural waters will be greater. Figure 7 also shows a total attenuation coefficient given in [9] for coastal Lake Ontario; we show this to illustrate the difference between absorption in pure water and absorption plus scattering in natural water. We suspect that this represents rather murky water, since a very similar value is given in [10] for 514 nm in turbid harbor sea water. For relatively clear, natural water, the total attenuation shown in [10] at 514 nm is an order of magnitude less, which is roughly what we show for pure water absorption at the red wavelength of the Leica in Figure 7. Thus, for natural water, we estimate that the total attenuation coefficient will be roughly in the range of 0.0028 to 0.022 per cm for the wavelength of best penetration.

To describe the rest of the path losses, the path length in the water is $x = L_w = D/\cos\theta_i$, where D is the water depth, so the power reaching the water bottom is $\Phi_2 = \Phi_1 e^{c(\lambda)L_w}$. Assuming Lambertian diffuse reflectance on the bottom with an albedo of ρ , the power reflected toward the receiver aperture is:

$$\Phi_3 = \Delta\Omega_w \frac{\rho}{\pi} \Phi_2 \cos\theta_i$$

where ρ/π models the Lambertian reflectance, $\cos\theta_i$ models the foreshortening of the surface patch in the direction of the receiver, and $\Delta\Omega_w$ is the solid angle of the receiver as seen from the surface patch on the bottom. We assume that the field of view of the receiver is large enough to include the entire footprint of the beam, so the size of the footprint cancels out of the derivation. Deriving $\Delta\Omega_w$ requires accounting for the beam refraction at the water-air interface on the return path and for the beam expansion that happens as a result. The beam expansion is modelled by Straubel's invariant, which gives the solid angle in the air as [3]:

$$\Delta\Omega_a = \frac{n_w^2 \cos\theta_t}{\cos\theta_i} \Delta\Omega_w$$

Using this and considerable algebra yields:

$$\Delta\Omega_w = \frac{A_r \cos\theta_i}{\cos\theta_i [n_w L_a + L_w \cos\theta_i / \cos\theta_t]^2}$$

where A_r is the area of the receiver aperture. Next, attenuation on the return path through the water column gives $\Phi_4 = \Phi_3 e^{c(\lambda)L_w}$. Assuming that the bottom reflection and the path through the water produce a largely unpolarized beam, the transition through the water-air interface gives the final, attenuated beam power reaching the receiver as:

$$\Phi_5 = \Phi_4 \frac{T_h + T_v}{2} \text{ Watts}$$

To use all of this, we need to instantiate values for the bottom albedo, the area of the receiver aperture, and the total beam attenuation coefficient. For illustration, we use an albedo of 0.3 and an aperture diameter of 1.0 inch (2.54 cm), which is close to the foreshortened diameter of the scanning mirror on the SICK. For attenuation coefficients, we look at four cases: three use the pure water absorption for 525, 625, and 900 nm from Figure 7 and one uses the total attenuation coefficient shown for 575 nm in murky water. Figure 8a plots the ratio of the return power to the output power as a function of lookahead distance for an example with a sensor 1 m high and water 30 cm deep. For the SICK wavelength, the return power is vanishingly small by a lookahead of less than 2 m; this is consistent with the empirical results in Figure 5. Even visible wavelengths are attenuated to almost nothing by about 7 m of lookahead. The figure does not show data for the 1550 nm case because it is so strongly attenuated by even one cm of water to be too small to plot on the scale of this graph.

The ultimate question we would like to answer is: for a given lookahead, what is the maximum water depth that yields a range measurement from the water bottom? To answer this, we need to know the minimum measurable return power for the lidar. We used experimental data like that in Figure 5 to estimate this for the SICK, then applied it to the examples plotted in Figure 8a. Figure 8b shows the results. For all cases, there is no return for lookaheads beyond about 6.5 m, even with negligible water depth. For completely dry ground, the simulation shows that the return signal would be too weak to register at a lookahead of 11 m. In tests with the real SICK lidar, we got no data for dry horizontal surfaces beyond 15 to 20 m with this sensor height; given various parameter uncertainties in our simulation, this is reasonable agreement. For the goal of estimating water depth, we argued earlier that the range of performance to be expected for natural water for the best case wavelength (green) is bounded roughly by total absorption coefficients of 0.0028 to 0.022 per cm. These correspond to the curves in Figure 8 for coastal Lake Ontario (murky) and for pure water at 625 nm (clear). At the murky end of the scale, a depth of 0.5 m should be measurable at a lookahead of about 2 m for a sensor 1 m high. This is enough to be useful for some vehicles, but there are vehicles with greater fording depth than this. At the

clear end of the scale, depths of over 1.5 m are penetrable out to over 4 m of lookahead, which is plenty for most vehicles. Note that these results all represent a simulation of a flat water surface; performance with waves will differ and will require additional work to model. Compared to Figure 5, the curve in Figure 8b for the SICK is somewhat optimistic; this is most likely because the natural pond water in Figure 5 had a larger total attenuation coefficient than the pure water absorption coefficient we used in Figure 8b (we have not found published data for total attenuation at this wavelength). To increase the measurable depth, the key requirement is to reduce the angle of incidence of the beam on the water, either by mounting the sensor higher on the UGV or, more radically, by putting the sensor on an air vehicle.

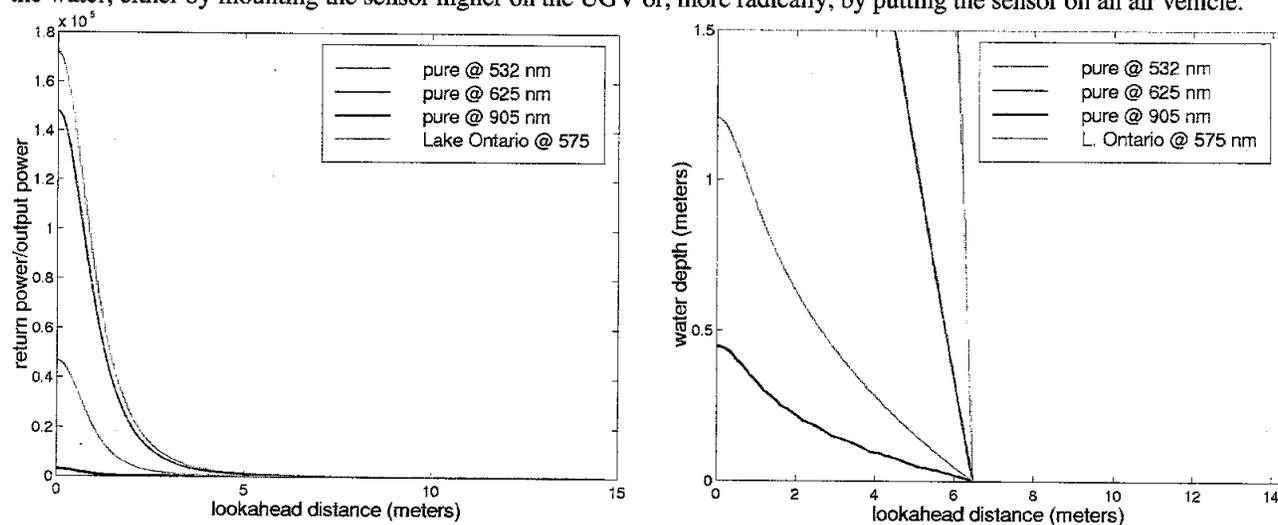


Figure 8: (a) Ratio of return power at the receiver to output power from the transmitter, vs. lookahead distance for a sensor 1 m above the water level. The three curves for pure water model only the absorption coefficients from Figure 7; absorption dominates scattering under these conditions, so this is a reasonable simplification. (b) Maximum depth of water for which the simulated lidar measures a return, vs. lookahead distance for a sensor 1 m above the water level. All curves assume an initially unpolarized beam for simplicity

For the goal of detecting the presence of a water body, we conclude that there are four distinct cases:

1. At the shortest ranges, visible and near infrared ladars may give a return from the water bottom, depending on the water depth.
2. For a range interval beyond that, water will completely attenuate the beam due to surface reflection and attenuation in the water column, but dry ground will give a return in this interval. For the simulations in Figure 8b, this interval starts at 6.5 m ahead (or closer, depending on water depth) and ends at 11 m. In this interval, lack of a return is diagnostic of water.
3. Beyond that, even dry level ground produces no return, so lack of return does not indicate a water body.
4. Others have observed that at high angles of incidence the ladar beam can bounce off water surfaces, hit a background object beyond the water, and bounce back off the water to the receiver, giving a range to the reflected object [11]. We have verified this experimentally ourselves. Determining at what lookahead distance this would be the dominant component of the return power requires more work.

In summary, designing a ladar interpretation algorithm to detect the presence of water bodies based on lack of return would have to take all of these cases into account. Except for the possibility of range to reflections of more distant terrain features, the useful lookahead distance for this technique will be limited to the range at which dry terrain would no longer give a return.

One other conclusion to draw from this analysis is that ladars using the 1550 nm wavelength will be very strongly affected by water on the ground; thus, heavy rain may dramatically reduce the range at which they perceive the ground.

4. SHORT-WAVE INFRARED (SWIR)

In the remote sensing literature, it is well known that water bodies of any appreciable depth appear very dark in near infrared, overhead imagery [12]; the reason for this is clear from the absorption coefficients in Figure 7. Given the rapid increase in the absorption coefficient at even longer wavelengths and the availability of SWIR cameras with sensitivity from about 0.9 to 1.7 μm , the question arises whether passive SWIR imagery might also be useful for water detection. We have observed that under a thick layer of marine fog there is still SWIR illumination, so clouds may not rule this out.

Useful lookahead ranges could be limited by surface reflections, as illustrated by the reflection coefficients in Figure 6b. Snow and ice also have very strong absorption beyond about $1.4 \mu\text{m}$ [13,14]; therefore, the wavelength region around 1.5 to $1.6 \mu\text{m}$ may be useful for recognizing water, snow, and ice. In this section, we demonstrate these possibilities empirically with sample imagery from an unfiltered SWIR camera with response from 0.9 to $1.7 \mu\text{m}$.

Figure 9 shows imagery of water taken with this camera. Figures 9a and 9b show images looking straight down into a plastic pail with 0 cm and 2 cm of water, respectively; the absorption in the latter case is clearly evident. Figure 9c shows an SWIR image of the same reservoir seen in Figure 2, though with a much smaller field of view than the color images in Figure 2. The water is very dark at the bottom of the image where it reflects the sky; the angle of incidence at that point is at least 80 degrees. Beyond that, strong reflections of clouds and trees are evident on the water surface. Note that the vegetation all around the reservoir is highly reflective at these wavelengths. Figure 10 shows a color and an SWIR image of a pile of ice cubes on a lawn; it is clear that the ice is much darker in SWIR than the lawn and the plastic cooler, though not as dark as the car tires. We expect that the water and ice in these images would be much darker if the camera was spectrally filtered to a window around 1.5 to $1.6 \mu\text{m}$. We conclude that these images show good potential for using SWIR to detect water at moderate angles of incidence and some potential to discriminate snow and ice from other terrain material.

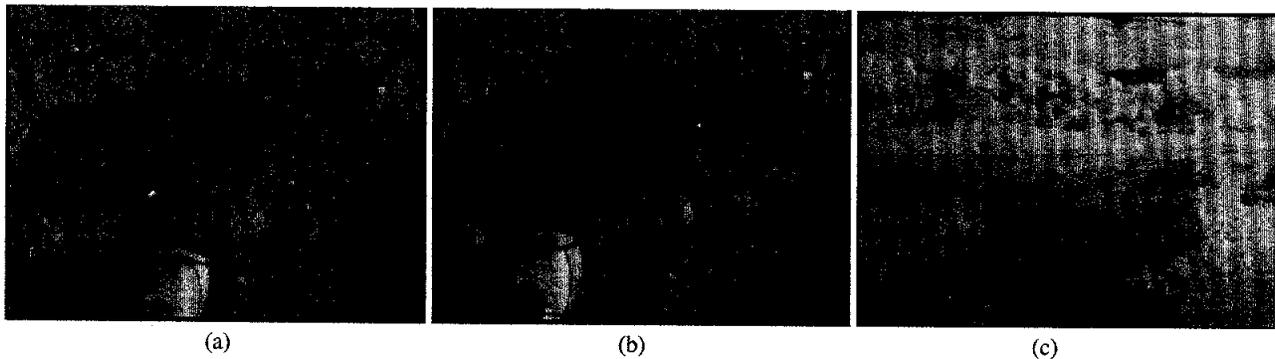


Figure 9: Imagery acquired with an unfiltered SWIR camera with response from 0.9 to $1.7 \mu\text{m}$. (a) Looking straight down into an empty plastic pail. (b) Looking at the same pail containing 2 cm of water. (c) Looking at the same reservoir as the color imagery in Figure 2; note that the field of view here is much smaller.

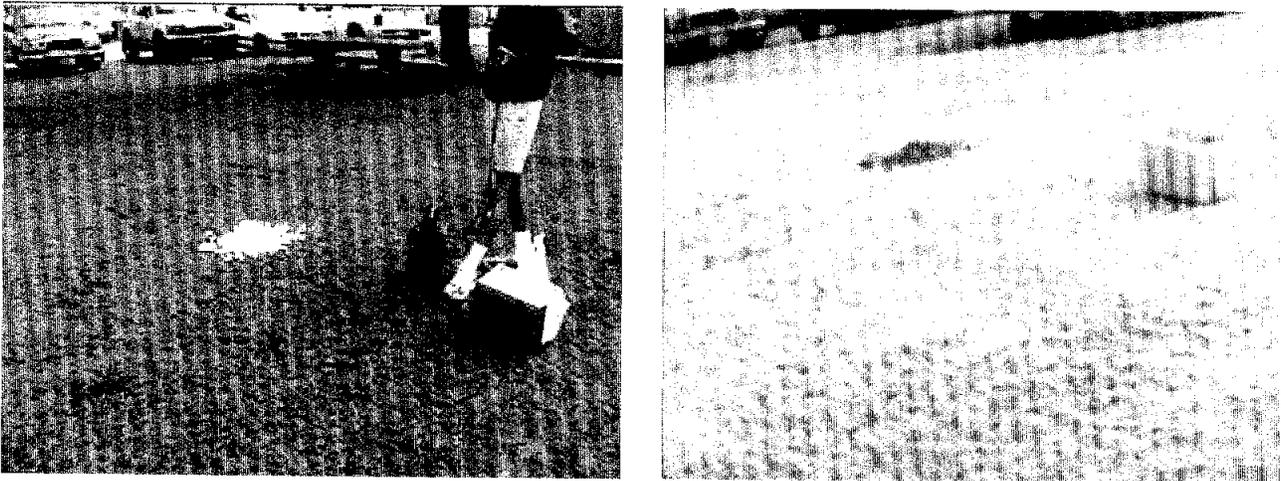


Figure 10: Color image (left) and SWIR image (right) of a pile of ice cubes on a lawn.

5. THERMAL INFRARED

For night operation, color and SWIR imagery is not likely to be useful for water detection. Ladar may be useful, but with limitations on the lookahead range as noted in section 3; moreover, it is desirable to have a non-emissive sensor option for this function. Another well known, qualitative observation in the remote sensing literature is that water bodies tend to be cooler than surrounding terrain during the day and warmer at night [12]. While this will be affected by the size

of the water body, in that very small bodies may equalize temperature with the terrain fairly quickly, it may still be useful for water bodies large enough to be a concern for UGV navigation. Moreover, water has a higher emissivity than other terrain materials, which will contribute to the night-time contrast. Surface reflections again come into play; in this case, their energy will add to that of the thermal emission from the water itself, so at night this may not be a problem. To test the usefulness of thermal infrared, we acquired a 24-hour data set of mid-wave thermal infrared imagery (MWIR) of the Chatfield reservoir at the same time that the color and SWIR imagery was acquired (ie. Figures 2 and 9c). Figure 11 shows sample MWIR images taken at 3 pm and 4 am; the water shows the expected contrast with the rest of the scene. We manually labelled the imagery into regions of water, soil, vegetation, sky, and "other", then plotted the mean brightness in each region over the whole 24 hour period (Figure 12). Except for two small intervals, the average brightness of the water region was higher than the averages of the other regions from about 11:00 pm to 8:00 am, and it was distinctly brighter from 12:30 am to 5:00 am. This supports the qualitative observation from the remote sensing literature. Local brightness variations within each region resulted in cases where some small regions of terrain were brighter than the water; in work in progress we are attempting to overcome this to develop a robust, MWIR-based water classifier for night operation. In the night image in Figure 11, the water also appears to darken near the far shore. We do not fully comprehend this phenomenon; it could be due to a thermal gradient, to non-Lambertian emission at high emission angles, or to some other cause; in any event, this might limit the lookahead distance possible with this technique.

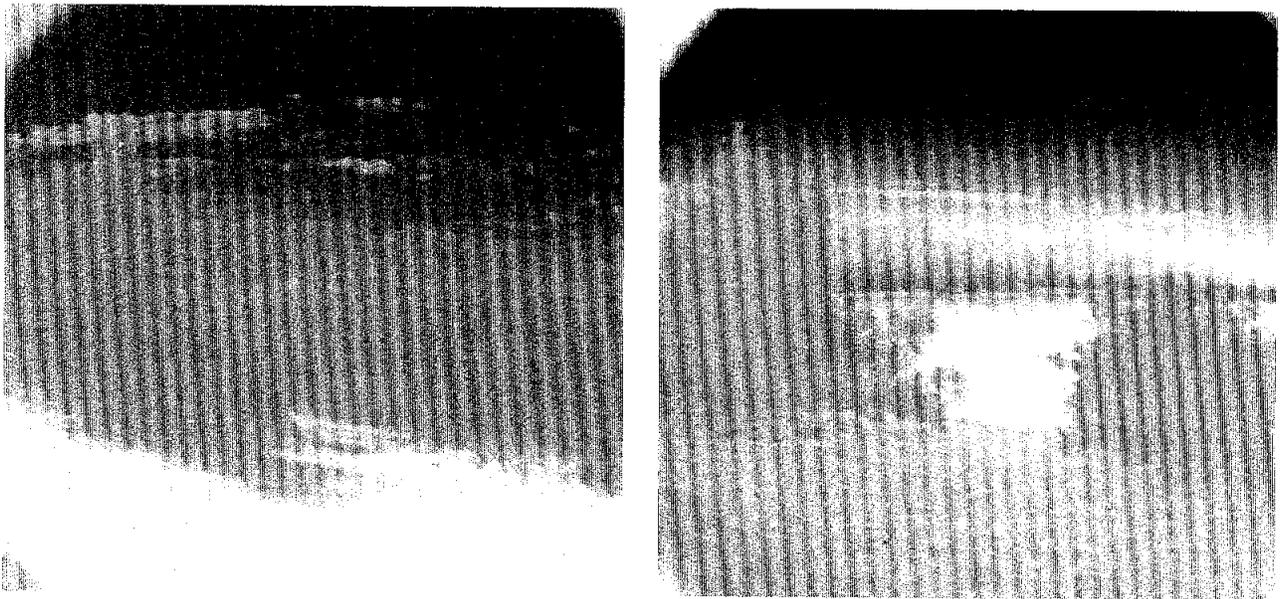


Figure 11: Mid-wave infrared (MWIR) imagery of the Chatfield reservoir taken at 3 pm (left) and 4 am (right).

6. CONCLUSIONS AND FUTURE WORK

For day-time detection of water bodies reflecting sky, we have demonstrated that RGB image classification works quite well to label water in imagery for clear, partly cloudy, and overcast skies and for smooth or rippled water surfaces. When still water reflects other terrain features, like trees, hills, or buildings, RGB classification is not adequate. We have shown that stereo vision often produces range data on such reflections and that this data is likely to enable a solution to this problem.

We presented a propagation model for lidar pulses incident on water bodies that leads to several useful conclusions. Surface reflection and beam attenuation in the water column are both important contributors to reduced return signal from water; moreover, the water attenuation is very highly wavelength dependent. Depending on wavelength, water depth, and water turbidity, at the shortest lookahead ranges (a few meters) lidar may give returns off the water bottom. At somewhat longer ranges, water will give no return but dry land will. Beyond 10 to 20 m, even dry level ground gives no return. At high incidence angles, lidar can reflect off the water to give range measurements to tall terrain features beyond the water. All of these cases must be embodied in any algorithm that attempts to use lidar to infer the presence of water hazards. Lidar can also be useful for sensing water depth, particularly in visible wavelengths. This works best

at normal incidence, so the best sensor placement would actually be on an air vehicle that scans the terrain ahead of ground vehicles. We also observe that water absorption at 1550 nm is so strong that ladars at this wavelength are likely to suffer much reduced maximum range to ground surfaces in heavy rain.

The very strong absorption of near infrared and especially short-wave infrared (SWIR) light by water implies that passive imagery in these wavelengths may also be useful for recognizing water hazards, though surface reflections may limit the useful lookahead range for this approach to incidence angles around 60 degrees. Snow and ice are also highly absorptive beyond 1.4 μm , so SWIR may also be useful for recognizing them as well.

Finally, we used 24-hour observations of a reservoir with mid-wave infrared (MWIR) imagery to confirm that large bodies of water tend to be warmer than surrounding terrain at night. This implies that MWIR imagery has potential for recognizing water hazards at night. Algorithm development to exploit this property is still in process.

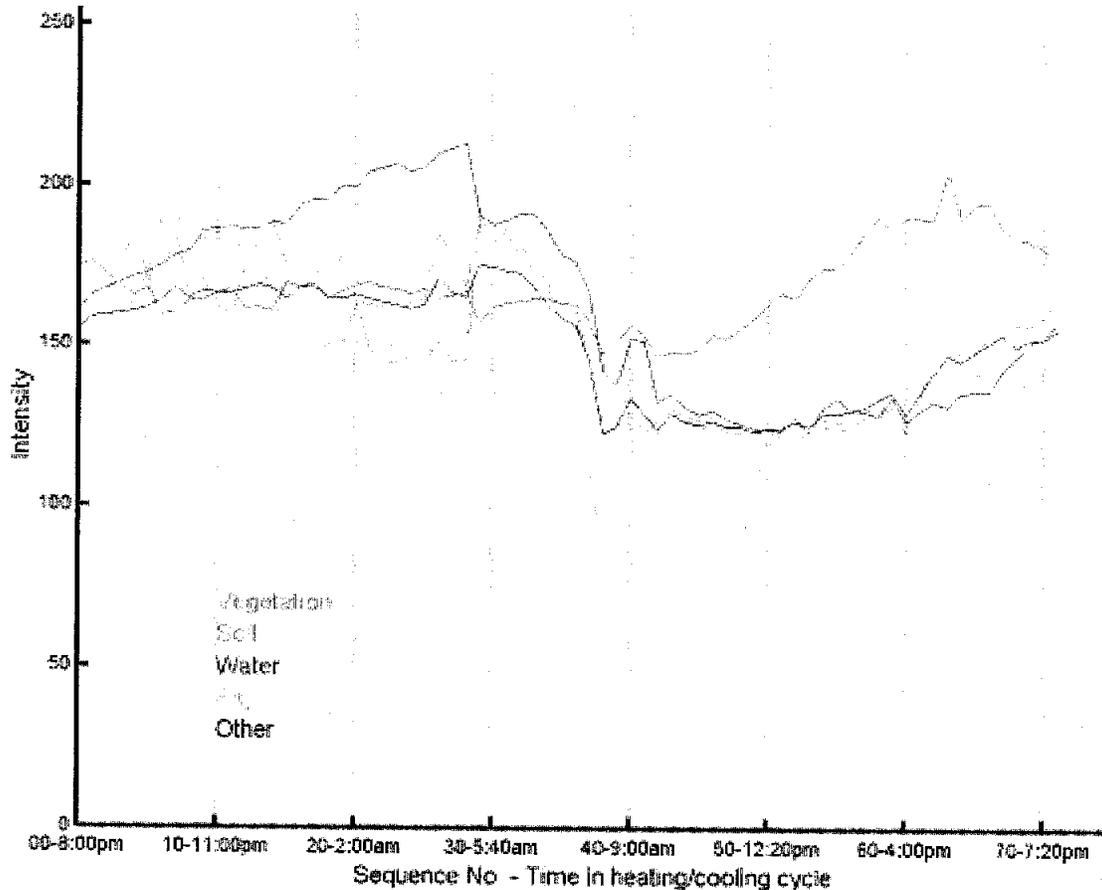


Figure 12: Mean brightness of the vegetation, soil, water, sky, and "other" regions of the Chatfield MWIR imagery over the entire 24 hour data set. On the horizontal axis, the labels give image sequence numbers and the time. Except for brief intervals around 12:30 am and 5:30 am, the water region was brighter than all other regions from about 11:00 pm to 8:00 am.

ACKNOWLEDGEMENTS

This work was sponsored by the Army Research Laboratory, under the Demo III program, and the Defense Advanced Research Projects Agency, under the Perceptor program, through contracts with the National Aeronautics and Space Agency. Staff at SAIC in Denver supported the data collections at the Chatfield reservoir.

REFERENCES

- [1] C.M. Shoemaker and J.A. Bornstein, "Overview and update of the Demo III Experimental Unmanned Vehicle program", *Proceedings of SPIE Conference 4024: Unmanned Ground Vehicle Technology II*, Orlando, FL, April 2000.

- [2] S. Baten, "Autonomous road and contour following with a tracked vehicle", *Proceedings of SPIE Conference 4024: Unmanned Ground Vehicle Technology II*, Orlando, FL, April 2000.
- [3] F. Thomanek and E.D. Dickmanns, "Autonomous road vehicle guidance in normal traffic", *Proceedings of ACCV*, 1995.
- [4] P. Bellutta, R. Mandchi, L. Matthies, K. Owens, and A. Rankin, "Terrain perception for Demo III", *Proceedings of the Intelligent Vehicles Symposium*, Dearborn, MI, October 2000.
- [5] J. Macedo, R. Manduchi, and L. Matthies, "Ladar-based discrimination of grass from obstacles for autonomous navigation", *Seventh International Symposium of Experimental Robotics (ISER)*, Honolulu, HI, December 2000.
- [6] K.L. Coulson, *Polarization and Intensity of Light in the Atmosphere*, A. Deepak Publishing, 1988.
- [7] L.B. Wolff, "Enhanced capabilities for battlefield awareness using polarization sensors", in *Reconnaissance, Surveillance, and Target Acquisition for the Unmanned Ground Vehicle*, O. Firschein and T. Strat (eds.), Morgan Kaufmann Publishers, 1997.
- [8] E. Hecht, *Optics*, 2nd ed., Addison-Wesley Publishing Company, Reading, MA, 1987.
- [9] R. P. Bukata, J. H. Jerome, K. Y. Kondratyev, D. V. Pozdnyakov, *Optical Properties and Remote Sensing of Inland and Coastal Waters*, CRC Press, Boca Raton, FL, 1995.
- [10] C. D. Mobley, *Light and Water: Radiative Transfer in Natural Waters*, Academic Press, San Diego, CA, 1994.
- [11] T. Jochem, private communication.
- [12] R.N. Colwell (ed.), *Manual of Remote Sensing*, 2nd edition, American Society of Photogrammetry, 1983.
- [13] A.W. Nolin and J. Dozier, "Retrieval of snow properties from AVIRIS data", *Proceedings of the 1991 AVIRIS Workshop*.
- [14] R.O. Green and J. Dozier, "Measurement of the spectral absorption of liquid water in melting snow with an imaging spectrometer", *Proceedings of the 1995 AVIRIS Workshop*.