© 2011 California Institute of Technology. Government sponsorship acknowledged.

Daytime Water Detection Based on Sky Reflections

Arturo L. Rankin, Larry H. Matthies, and Paolo Bellutta

Abstract— **Robust water detection is a critical perception requirement for unmanned ground vehicle (UGV) autonomous navigation. This is particularly true in wide-open areas where water can collect in naturally occurring terrain depressions during periods of heavy precipitation and form large water bodies. One of the properties of water useful for detecting it is that its surface acts as a horizontal mirror at large incidence angles. Water bodies can be indirectly detected by detecting reflections of the sky below the horizon in color imagery. The Jet Propulsion Laboratory (JPL) has implemented a water detector based on sky reflections that geometrically locates the pixel in the sky that is reflecting on a candidate water pixel on the ground and predicts if the ground pixel is water based on color similarity and local terrain features. This software detects water bodies in wide-open areas on cross-country terrain at mid- to far-range using imagery acquired from a forward-looking stereo pair of color cameras mounted on a terrestrial UGV. In three test sequences approaching a pond under a clear, overcast, and cloudy sky, the true positive detection rate was 100% when the UGV was beyond 7 meters of the water's leading edge and the largest false positive detection rate was 0.58%. The sky reflection based water detector has been integrated on an experimental unmanned vehicle and field tested at Ft. Indiantown Gap, PA, USA.**

I. Introduction

Detecting water hazards is a significant challenge to unmanned ground vehicle (UGV) autonomous navigation over cross country terrain. This is particularly true for military UGVs navigating over wide-open areas such as clearings. Here, higher vehicle speeds are desired to limit the time a UGV is exposure to ground and aerial detection. But these are areas that typically contain naturally occurring depressions where large water bodies (such as ponds) form during periods of heavy precipitation. The probability of driving into a water hazard increases when a UGV is required to operate at higher than normal speeds, especially since *a priori* water data may not be available as current water hazards may not have existed when the most recent digital terrain map was generated. Thus, the need to traverse wide-open areas at high speed while detecting and avoiding previously unknown water traps necessitates robust water detection.

Water detection for UGV autonomous navigation is still a relatively new research area. A survey of sensors useful for detecting water bodies from a UGV can be found in [1]. Thermal infrared [2], lidar [3], and polarization [4][5][6] sensors have been explored for daytime water detection. But these sensors are all relatively high cost in comparison to commercially available color cameras that have successfully being used for water detection [1][7][8][9][10].

Under the Robotics Collaborative Technology Alliances (RCTA) program, the Jet Propulsion Laboratory (JPL) has developed a water detector for wide-open areas that uses a forward-looking stereo pair of color cameras mounted to the front of a UGV to detect water based on its variation in color. It exploits the knowledge that the change in saturation-to-brightness ratio across a water body from the leading to trailing edge is uniform and distinct from other terrain types [11]. This detector performs well out to 35-50 meters. Beyond this range, however, specular reflections of the sky thoroughly dominate the color of water bodies and water detection based on variation in color is of marginal use. Fig. 1 illustrates sky reflections dominate the color of water at far range for a variety of the sky conditions. As shown in Fig. 2, the reason for this is the fraction of light that is reflected from an air/pure water interface increases with range.

Fig. 1. Sample images of a water body in a wide-open area under a variety of different sky conditions. These images are from a data set collected over a three month span. The three images to the left were acquired from the same side of the water body. The length of this water body exceeds 70 meters.

To detect water bodies in wide-open areas at far range, JPL implemented a water detector based on sky reflections that geometrically locates the exact pixel in the sky that is reflecting on a candidate water pixel on the ground and predicts if the ground pixel is water based on color similarity and local terrain features. When a water body is detected at far range, a UGV's path planner can begin to look for alternate routes to the goal position sooner. Consequently,

Manuscript received February 8, 2011. The research described in this paper was carried out by JPL, California Institute of Technology, and was sponsored by the U. S. Army Research Laboratory (ARL) RCTA program through an agreement with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or JPL, California Institute of Technology.

A. Rankin is with the Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, CA 91109 USA. (phone: 818-354-9269; fax: 818-393-5007; e-mail: Arturo.Rankin@ jpl.nasa.gov).

L. Matthies and P. Bellutta are also with the Jet Propulsion Laboratory. (e-mail: Larry.Matthies@jpl.nasa.gov, Paolo.Bellutta@jpl.nasa.gov).

© 2011 California Institute of Technology. Government sponsorship acknowledged.

detecting water hazards at far range generally reduces the time required to reach a goal position during autonomous navigation.

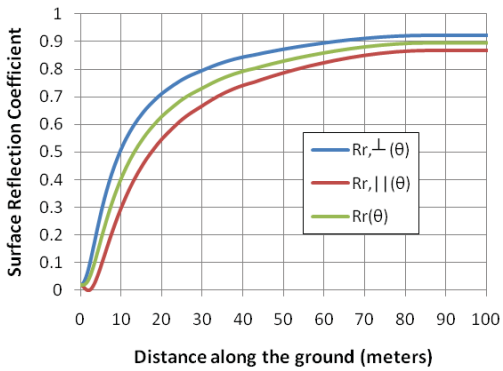

Fig. 2. Theoretical fraction of incident power that is reflected from an air/pure water interface as a function of distance along the ground. A sensor height of 1.5 meters was used. $R_{r,\perp}(\theta)$ and $R_{r,\parallel}(\theta)$ are the Fresnel reflection coefficients for light polarized perpendicular to and parallel to the plane of incidence, respectively. $R_r(\theta)$ is the Fresnel reflection coefficient for unpolarized light.

Since the surface reflection coefficient of water increases with range, the primary factor that limits the range of sky reflection based water detection is water body surface area. Assuming a horizontal ground plane and a water surface modeled as a circle, (1) gives an expression for the minimum detectable water body (diameter D) as a function of the maximum distance to the water's leading edge L, given the sensor height H, the sensor's vertical angular resolution [also called instantaneous vertical field of view ($IVFOV$)], and the minimum number of pixels required for detection N. In the sample plot in Fig. 3, to detect a water body whose leading edge is 100 meters downrange, the water body needs to have a diameter of greater than 138 meters.

$$D = \frac{H\left[L + H\tan\left(2 \cdot IVFOV\sqrt{\frac{N}{\pi}}\right)\right]}{H - L\tan\left(2 \cdot IVFOV\sqrt{\frac{N}{\pi}}\right)} - L \qquad (1)$$

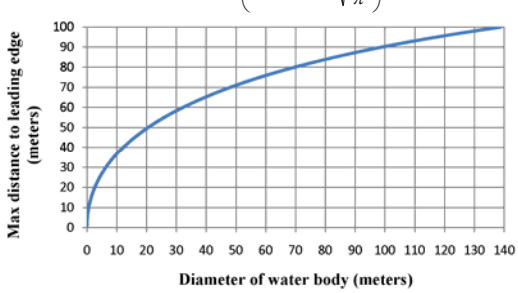

Fig. 3. The maximum range at which a water body can be detected is a function of the surface area of the water body. Here is a plot of (1), using H=1.5m, N=50pixels, and $IVFOV$=1.0925mrad.

The steps in detecting water based on sky reflections are 1) estimate the horizon line, 2) search above the horizon line and detect the sky, and 3) step through the detected sky pixels to determine if the ground pixels they are reflecting on has a similar color and terrain features. In the following sections of this paper, we describe these algorithm steps.

II. Finding the Horizon

The apparent line that separates the earth from the sky is called the horizon line. The horizon line is useful to water detection for two primary reasons. First, if we know where the horizon line is in images of a scene, we can constrain our search for water bodies to the region below the horizon, thereby decreasing the computational cost of water detection as well as reducing the probability of false water detection. Similarly, knowing the location of the horizon line constrains our search for the sky to the region above the horizon line. As we will see in section V, the color of the sky is a strong cue for water bodies in wide-open areas, and knowing the location of the sky in color imagery is a critical component to water detection based on sky reflections.

For a forward-looking imaging sensor mounted to a UGV, the line-of-sight distance to the horizon can be approximated as $[H(2R+H)]^{0.5}$, where H is the height of the sensor above the ground surface and R is the mean radius of the earth (6,371 km) modeled as a sphere. For a sensor height of 1.5 meters, the line-of-sight distance to the horizon is approximately 4,372 meters. The true horizon line may be detectable in color imagery under a limited set of conditions, for example, the earth's surface in the distance is barren terrain or water, there is no sun glare, and changes in a UGV's pitch angle are not causing the horizon line to move outside of the camera's field of view (FOV).

However, if a model of the sensor's optical characteristics is known, including its orientation from an onboard inertial sensor, the horizon line can be directly calculated without having to search for it within an image. Calculating the horizon line from sensor optical characteristics and camera attitude data works even when the earth's surface in the distance is obstructed with natural or man-made structure, there is sun glare, or there are large changes in a UGV's pitch angle.

The horizon line can be approximated as the line of intersection between the camera image plane and the horizontal plane that lies on the optical center of the camera lens. The actual horizon line depends on the effective camera height above ground level, which can vary significantly on open terrain. To approximate the line of intersection, it is easy to calculate the intersection point of the camera image plane with two arbitrary horizontal vectors, $\mathbf{q}_1$ and $\mathbf{q}_2$, originating at the optical center (2). These two image points, (c_1, r_1) and (c_2, r_2) define the horizon line on the camera image plane (3).

The approximation is easily achieved by using the vectors of a **CAHV** camera model [12]. **C** is the coordinate of the camera optical center, **A** is the optical axis extending out of the optical center, and **H** and **V** are the vectors of the horizontal and vertical components of the camera image plane, respectively. The camera model is typically measured relative to flat ground. During UGV operation, the camera model needs to be transformed into a vehicle frame every imaging cycle using inertial sensor attitude data. But since the transformation is already performed every imaging cycle to generate 3D coordinates from stereo matching, there is no additional computational cost to the system.

© 2011 California Institute of Technology. Government sponsorship acknowledged.

$$\mathbf{q}_1 = \{-A[0], -A[0], 0\}$$
$$\mathbf{q}_2 = \{-A[0], \quad A[0], 0\} \tag{2}$$

$$r_1 = (\mathbf{V}\cdot\mathbf{q}_1)/(\mathbf{A}\cdot\mathbf{q}_1)$$
$$c_1 = (\mathbf{H}\cdot\mathbf{q}_1)/(\mathbf{A}\cdot\mathbf{q}_1)$$
$$r_2 = (\mathbf{V}\cdot\mathbf{q}_2)/(\mathbf{A}\cdot\mathbf{q}_2)$$
$$c_2 = (\mathbf{H}\cdot\mathbf{q}_2)/(\mathbf{A}\cdot\mathbf{q}_2) \tag{3}$$

The horizon line may lie outside the camera FOV for extreme positive or negative values of vehicle pitch or pan-tilt unit (PTU) tilt angles. As shown in (4), the expected error E in horizon line estimation is a function of the sensor height H, the radius of the earth R, the camera's vertical angular resolution $IVFOV$, the inertial sensor's pitch accuracy ($\gamma_{inertial}$), and the tilt resolution of the PTU (γ_{PTU}), if one is used to actively point the camera.

For a camera having a resolution of 1024x768 pixels, a pixel pitch of 4.65µm, and a rectilinear lens with a 4mm focal length, the expected error from assuming a horizontal plane in (2) is 0.31 pixels. If, for example, a Smiths Industries Land Navigation System (having a pitch accuracy of 0.873mrad) and a Directed Perception PTU-D48-E (having a tilt resolution of 0.052mrad) are used, the combined expected error is 1.2 pixels. Camera, inertial sensor, and PTU calibration errors can further increase the error in estimating the horizon line. A result from our horizon finding algorithm is shown in Fig. 4. As expected, the horizon line separates the sky from the ground surface regions that may contain water bodies.

$$E = \frac{\sin^{-1}\left[\dfrac{H}{\sqrt{H(2R+H)}}\right] \pm \gamma_{inertial} \pm \gamma_{PTU}}{IVFOV} \tag{4}$$

Fig. 4. A **CAHV** camera model and UGV attitude data from an onboard inertial measurement unit can be used to estimate the horizon line.

III. Detecting the Sky

As seen in Fig. 1, the color of a water body that is out in the open and far away closely matches the color of the sky. Therefore, it is useful to locate the sky in input imagery and determine its color. The color of the sky depends upon atmospheric conditions. The main source of light in the sky is the sun, but light from atmospheric scattering and reflections off clouds is also emitted by the entire sky dome.

There are three primary sky conditions: clear, overcast, and partly cloudy. Under a clear sky, sunlight entering the atmosphere is refracted and reflected, causing sky light to be diffuse. As light passes through a clear sky, longer wavelengths tend to pass straight through but shorter wavelengths are absorbed by gas molecules and radiated in all directions, causing the sky to appear blue. The depth of blue depends upon the UGV height above sea level and the amount of pollution in the air. Nitrogen dioxide makes the color of the sky appear slightly brown or yellow, particularly near the horizon and in the direction of urban areas where concentrations are the highest. Under an overcast sky, the sun and its position are completely obscured by cloud cover. Water particles in the clouds diffusely refract and reflect all wavelengths of sunlight, causing the sky to appear white. Under a partly cloudy sky, the sky contains a combination of blue and white regions.

Cues useful for classifying the sky in color imagery are saturation-to-brightness ratio, monochrome intensity variance, and edge magnitude. As illustrated in Fig. 5, an overcast sky tends to have image saturation, a low saturation-to-brightness ratio, low variance, and a low edge magnitude, a clear sky tends to have a high blue content, low variance, and low edge magnitude, and a cloudy sky tends to have a low saturation-to-brightness ratio. Edge magnitude can be used to find the lower bound of the sky.

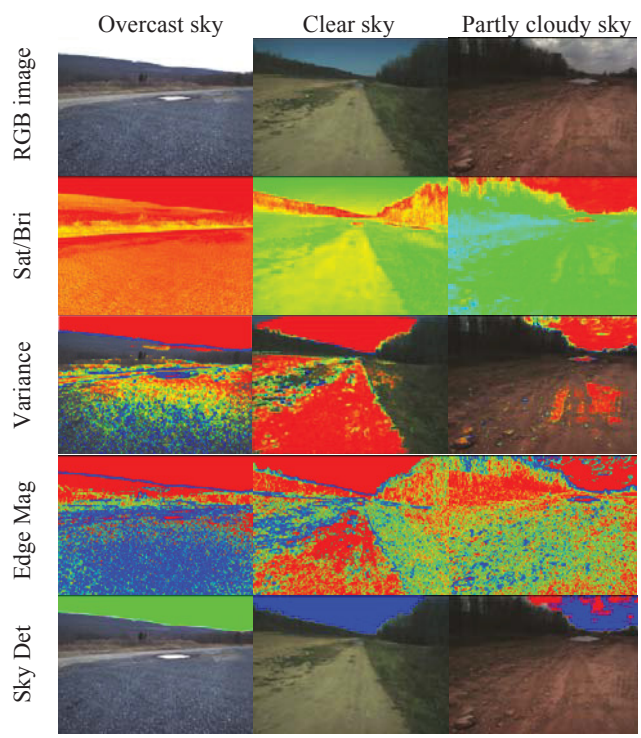

Fig. 5. Cues for sky detection are saturation-to-brightness ratio, image variance, and edge magnitude. In the middle rows, red corresponds to low values, blue corresponds to high values, and the colors in between correspond to an intermediate value. In the bottom row, green indicates overcast sky detection, blue indicates blue sky detection, and a combination of blue and red indicates cloudy sky detection. Edge magnitude is used to find the lower bound of the sky.

© 2011 California Institute of Technology. Government sponsorship acknowledged.

To locate sky pixels, we implemented a sky detection algorithm that performs the following steps on rectified RGB images from the left camera:

1) Convert each native RGB image to hue, saturation, and brightness (HSB) color components.
2) Find the horizon line.
3) Step through the pixels above the horizon line, thresholding brightness to locate saturated pixels, and thresholding hue and brightness to locate pixels with a strong blue content.
4) At each pixel classified as clear or overcast sky, flood fill all connected pixels that have a low monochrome intensity gradient and a low edge magnitude.
5) Threshold the detected sky average saturation-to-brightness ratio to determine if the sky is overcast.
6) If there are bright regions above regions classified as clear sky, classify them as clouds.
7) For pixels classified as overcast, flood fill all connected non-overcast pixels that have a low saturation-to-brightness ratio, low monochrome intensity gradient, and low edge magnitude, and classify them as clouds.

IV. Detecting Water Based on Sky Reflections

The appearance of water bodies in color imagery largely depends on the ratio of light reflected off the water surface to the light coming out of the water body. When a water body is far away, the angle of incidence is large, and the light reflected off the water surface dominates the color coming out of the water body. We implemented two algorithms to detect water bodies in wide-open areas based on sky reflections. In the first method, we ran the sky detection algorithm on left camera rectified color imagery, averaged the color of the lowest N sky pixels in each image column, and then scanned each image column below the horizon for colors that were a close match to the averaged sky color. This technique yielded good results when water was far away, but marginal results at shorter range. In the second method, we segment candidate water pixels in left rectified color imagery based on color and texture [13], geometrically locate the exact pixel in the sky reflecting on each candidate water pixel, and threshold the difference in color.

A water body can be modeled as a horizontal mirror. Fig. 6 illustrates the geometry. A ray of light travels along an incident ray, is reflected off the surface of a water body along a reflected ray, and enters a pixel of a camera's focal plane array (FPA). According to the law of reflection, the angle of incidence is equal to the angle of reflection. A direct ray from the tail of the incident ray (and within the same vertical plane as the incident ray) will enter the camera's FPA at a pixel whose color will indicate the color of the sky being reflected along the reflected ray.

Since the distance between the camera and the light source is much larger than the distance between the camera and candidate water points at normal detection ranges, the direct ray and the incident ray will be nearly parallel and the angle between the direct ray and the reflected ray can be approximated as two times the glancing angle (5). The calculations to locate the pixel a direct ray enters are simple, using a **CAHV** camera model. For any candidate water pixel, convert the 2D image coordinates to a 3D unit vector in a coordinate frame where the z axis is aligned with gravity, negate the z component of the unit vector, and convert the modified unit vector to 2D image coordinates.

$$\begin{gathered} \beta + (\pi - 2\theta) + \alpha = \pi \\ \lim_{|a| \to \infty} \alpha = 0 \\ \beta \approx 2\theta \end{gathered} \tag{5}$$

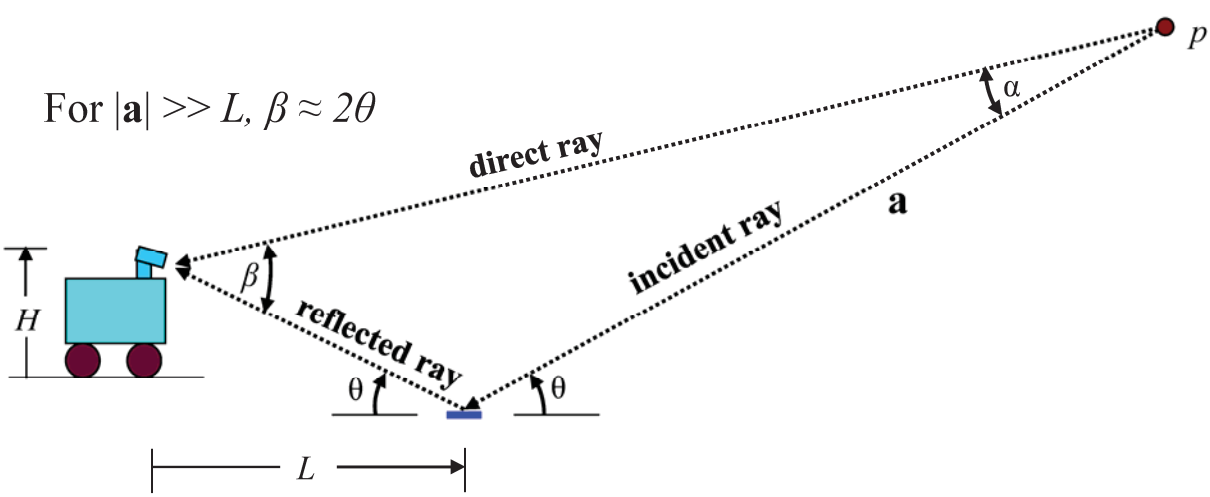

Fig. 6. Since the distance to the light source is much larger than the distance to a candidate water point, the point in the sky reflecting on a candidate water pixel will be found along a direct ray parallel to the incident ray, and in the same vertical plane as the incident ray. That is, the limit of β as point p moves out to infinity along vector a, is 2θ.

© 2011 California Institute of Technology. Government sponsorship acknowledged.

When a candidate water point is far away, the glancing angle θ is small and the sky pixel reflecting on the candidate water pixel is close to the horizon. When a candidate water point is close to the camera, the glancing angle is large and the sky pixel reflecting on the candidate water pixel is closer to the top of the image, or possibly outside the top of the FOV. This is illustrated in the graph in Fig. 7. Assuming the ground is horizontal and using JPL camera parameters for a RCTA experimental unmanned vehicle (XUV), the point in the sky reflecting on a candidate water pixel moves past the top of the camera's FOV once the XUV is within ~6 meters of the candidate water point. By this point, the color coming out of the water body dominates and knowing the color of the sky reflecting on a candidate water pixel is of marginal use.

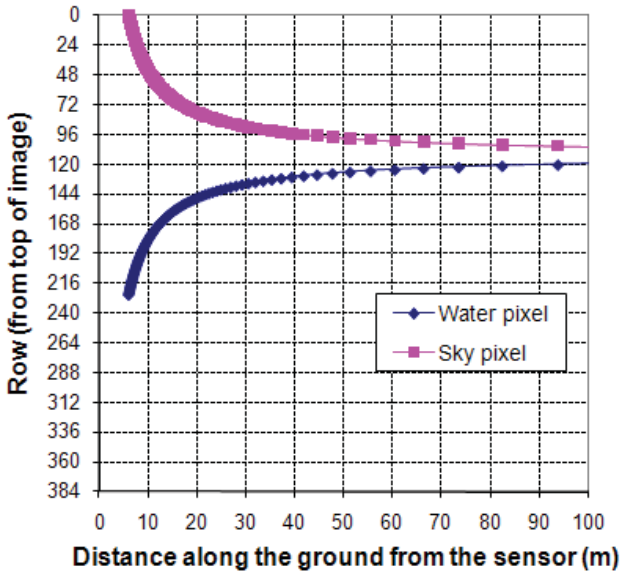

Fig. 7. Using JPL camera parameters for an XUV, the point in the sky reflecting on a candidate water point moves past the top of the camera's FOV once the XUV is within ~6 meters of the candidate water point. (Assumptions: horizontal ground, CameraHeight=1.5m, CameraTilt=10° down, 48° vertical FOV, images processed at a resolution of 512x384 pixels).

Fig. 8 shows a result of locating sky pixels reflecting on a water body at Fort Indiantown Gap (FITG), PA. This algorithm is highly dependent on accurate camera roll and pitch estimates being available. Here, we rely on the vehicle's inertial measurement unit (IMU), assuming that the IMU and cameras are mounted to the same rigid body. Therefore, an accurate IMU calibration is critical for this technique to work well. This algorithm is robust to camera pitch and roll changes.

The vertical line in Fig. 8 illustrates that when there is camera roll, one cannot simply look up a specific image column a certain number of pixels for the sky pixel reflecting on a candidate water pixel. If there is camera roll, then the sky pixel reflecting on the water will be in the antigravity direction. Fig. 9 illustrates the sky pixels reflecting on a water body as an XUV approaches the water body. As expected, the portion of the sky reflecting on the water body moves past the top of the FOV when the water body is close.

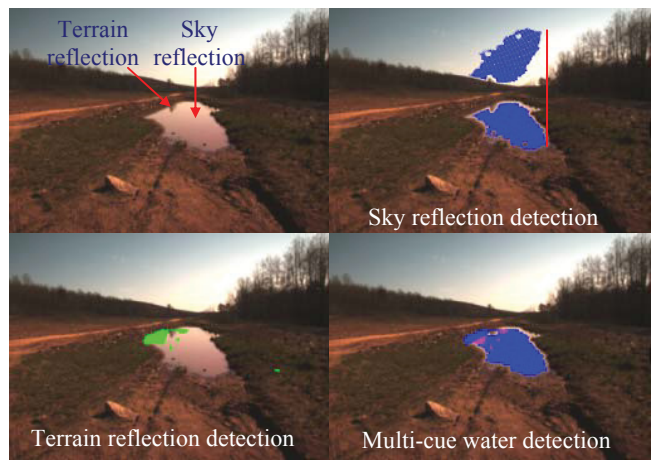

Fig. 8. In this example, water is indirectly detected by detecting sky and terrain reflections and fusing the results. The upper left image indicates the portion of the sky reflecting on candidate water pixels on the ground. When there is camera roll, the sky pixel reflecting on a candidate water pixel may not be in the same image column.

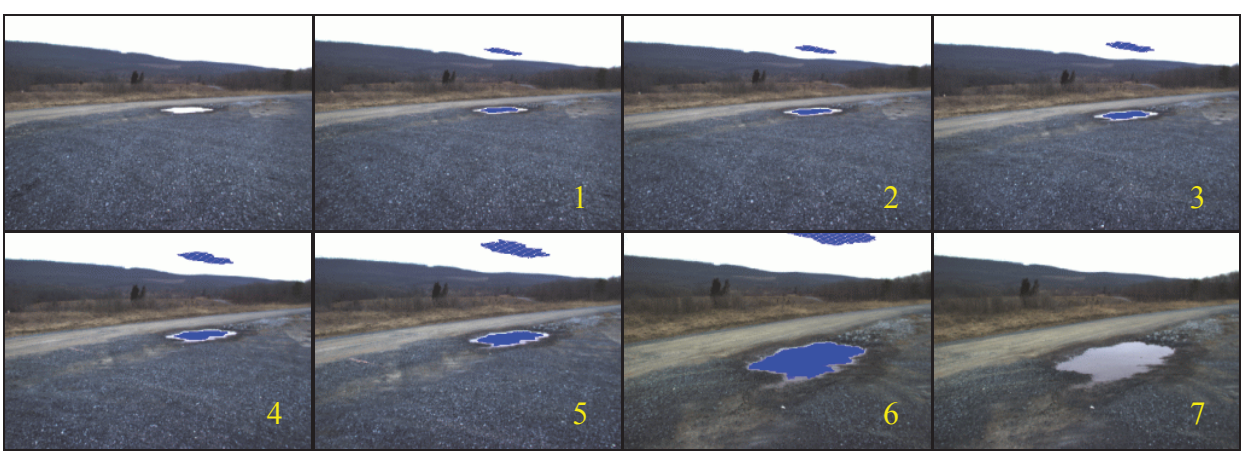

Fig. 9. Sky reflection detection as an XUV approaches a water body at FITG. The upper blue region in each image illustrates localization of sky pixels (in image space) that are reflecting on the water body. As predicted by the graph in Fig. 7, the portion of the sky reflecting on a water body begins to move out of the top of the FOV as water points are within a ~6 meter downrange distance of the stereo cameras (for the JPL camera configuration on a XUV)

© 2011 California Institute of Technology. Government sponsorship acknowledged.

V. False Detection Rejection

False water detection can occur on objects or terrain that has a color similar to the sky. For example, white vehicles, cement pavement, and sun glare all have color components similar to an overcast sky. As illustrated in Fig. 10, bright vehicles and other 3D objects can be rejected by thresholding their slope, estimated using stereo range data. As illustrated in Fig. 11, sun glare can be easily detected since it is characterized by one or more rays of image saturation extending from the sun to the bottom of the image. Water detection is currently skipped for input images where sun glare is detected.

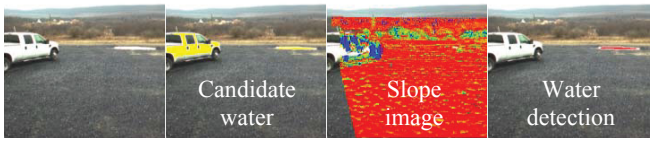

Fig. 10. White vehicles can be disambiguated from water bodies under an overcast sky by using stereo range data to perform slope analysis. The leftmost image contains a white truck and a puddle, the second image labels the pixels similar to the sky (yellow), the third image shows the slope of the scene (where red and blue corresponds to low and high slope, respectively), and the final image shows the truck was correctly rejected.

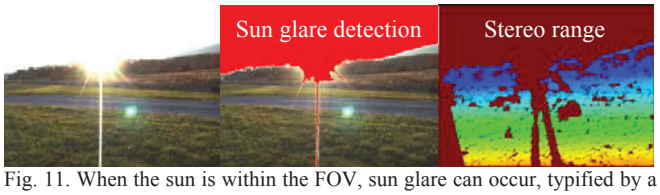

Fig. 11. When the sun is within the FOV, sun glare can occur, typified by a ray of saturated pixels extending from the sun to the bottom of the image. Sun glare causes lines of data loss in stereo range images. Water detection is not performed on images where sun glare is detected.

Disambiguating horizontal cement surfaces and water bodies under an overcast sky is an area we have not yet explored. As illustrated in Fig. 12, patches of white gravel with low slope can be disambiguated from still water based on intensity variance. Still water tends to have an intensity variance similar to the object reflected in it.

When navigating on cross-country trails, it is highly desirable to minimize deviations off the trail due to false positive water detection. To help suppress false positive water detection on trails, we use the algorithm described in [14] to detect the ground surface. If greater than 75% of the pixels detected as a water body were also detected as ground surface, that water detection is not inserted into the vehicle-level map used to plan safe paths.

Fig. 13 illustrates that sky detection is extremely useful for defining the search space for water bodies reflecting the sky. In the right image of Fig. 13, dark blue pixels directly correspond to the ground pixels reflected on by the detected sky. Magenta pixels indirectly correspond to ground pixels reflected on by the detected sky by assuming the sky continues beyond the top, left, and right image borders. Ground pixels with colors similar to the sky can be safely rejected if they fall outside of this space.

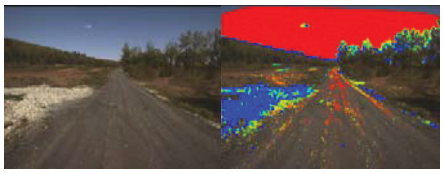

Fig. 12. Grayscale intensity variance may be useful in disambiguating patches of small bright rocks with low slope from water bodies. The right image indicates there is low variance in the sky (red) but high variance on the bright rocks (blue). Variance is not calculated for extremely dark pixels since water bodies tend not to be extremely dark.

Fig. 13. With sky detection results and UGV attitude information, one can locate the pixels on the ground the sky pixels are reflecting on. In the right image, the blue pixels directly correspond to the detected sky pixels, and the magenta pixels indirectly correspond to the sky by assuming the sky continues beyond the top, left, and right borders of the images.

As illustrated in Fig. 14, when driving through a heavy shadow, a camera's automatic gain control (AGC) can cause image saturation above the shadow. The combination of high luminance, low texture, and low terrain slope can cause false positive water detection.

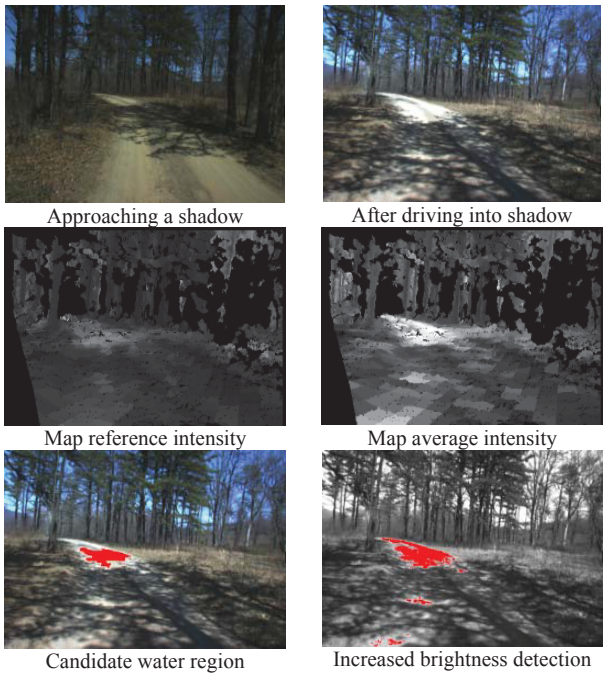

Fig. 14. The combination of high luminance, low texture, and low terrain slope can cause false positive water detection. Trails can be washed out by a camera's AGC when driving through shadows. Software has been implemented to detect terrain that experiences a significant increase in brightness as it is approached. These regions are likely caused by AGC when driving through shadows.

We have three strategies to disambiguate AGC induced image saturation from water:

1) If the mean intensity in the lower half of the image decreases but the blue content increases, we have

© 2011 California Institute of Technology. Government sponsorship acknowledged.

probably driven into a shadow and can ignore any image saturation above the shadow.

2) Since water is typically not saturated under a clear (blue) sky, we can ignore any image saturation when the sky is blue.
3) If the intensity of a region significantly increases over time as that terrain is approached, it is not likely water. A water body typically becomes darker as it is approached because the color coming out of the water body starts to dominate reflections from the sky.

The 3rd strategy was implemented using a 0.4m resolution 100m x 100m terrain map. The first instance of data that occurs in each map cell is averaged and recorded as the reference intensity for that cell. As illustrated in Fig. 14, the current average intensity and the reference intensity is compared for each map cell. Any candidate water regions in cells that experience a significant increase in brightness are ignored. This filter is more useful in cluttered environments (where shadows are common) than in wide-open areas.

VI. RESULTS

Only a portion of a water body needs to have a similar color to the sky for the water body to be detected. We use a flood-fill algorithm to grow water detection regions until the intensity gradient is no longer low. The sky reflection based water detector was run on three image sequences collected under clear, overcast, and cloudy sky conditions (see Fig. 15 and Fig. 16). The clear sky sequence was collected during an approach to the pond in Fig. 1, starting 64 meters from the leading edge, and the overcast and cloudy sky sequences were collected during approaches to the same pond but from the opposite side (from 51 and 47 meters, respectively). Table 1 contains detection results for these three sequences.

The water body was detected in every frame of the sequences collected under a clear and overcast sky. For the sequence collected under a cloudy sky, the UGV's final distance from the water's edge was closer than the other two sequences. For this sequence, the water body was detected in all frames collected beyond 7 meters from the leading edge. It was detected in only 1 of the 24 frames (4.2%) collected closer than 7 meters from the leading edge. This is because at close range the color coming out of water bodies in wide-open areas dominates sky reflections. The only false positive water detection was in one frame of the cloudy sky sequence. To help suppress false positive detection, we required each water detection blob contain at least 50 pixels.

The average time to classify each 512x384 image on a 2.6 GHz Intel Core 2 Duo processor was 105, 111, and 99 milliseconds, respectively, not including stereo processing to localize water or the filter which disambiguates AGC induced image saturation from water. The horizon finding algorithm took less than 0.1ms, preprocessing steps (RGB to HSB color conversion, edge magnitude, intensity variance, and saturation-to-brightness ratio) took an average of 39ms, sky detection took an average of 9.9ms, and sky reflection based water detection took an average of 52.8ms.

Our water detector based on color variation was tested on these three sequences in [11]. There, the true positive detection rates for the clear, overcast, and cloudy sequences started falling of beyond 50, 35, and 35 meters, respectively. For the overcast sequence, the true positive detection rate dropped to 50% between 40-51 meters.

While false positive detections on trails can cause a UGV to deviate off-road into much rougher terrain, or even become stuck, there tends to be milder reactions when avoiding false positive detections in wide-open areas. Therefore, it is more critical to avoid false positive water detection on trails than in wide-open areas. The sky reflection based water detector was tested on the seven trail sequences illustrated in Fig. 14 and Fig. 17. For these tests, the false positive detection rejection algorithms described in section V were enabled. In the 1711 frames processed over a 0.44km distance, false positive water detection only occurred in 3 frames, a rate of 0.18%.

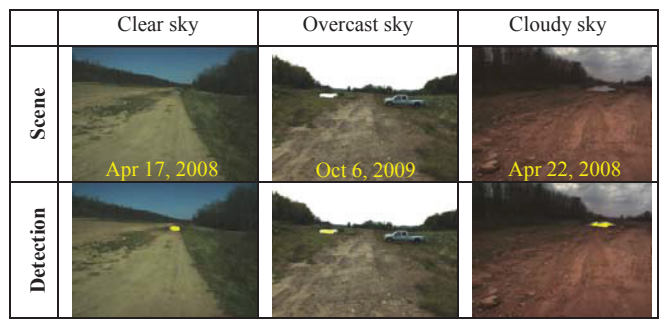

Fig. 15. First-frame water detection results for three sequences collected under clear, overcast, and cloudy sky conditions, approaching the water body in Fig. 1. The approaches on the overcast and cloudy days were from the opposite side.

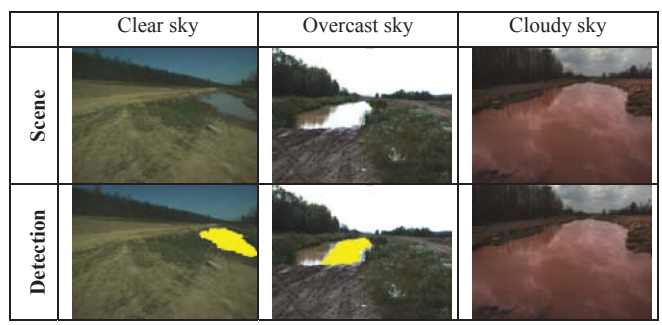

Fig. 16. Final-frame water detection results for the three sequences described in Fig. 15. For the clear sky sequence, the UGV turned away from the water body following this frame. For the overcast sky sequence, only the portion of the water body reflecting the sky is detected at close range. For the cloudy sky sequence, the water body was undetected at the closest range (1.9m).

Table 1. Results for the 3 sequences illustrated in Fig. 15 and Fig. 16.

Scene	1st/last frame distance to water	Num frames	True positive detection rate	False positive detection rate	Avg frame time
Clear	64m 11.4 m	265	100%	0%	105ms
Over-cast	51 m 9.7 m	434	100%	0%	111ms
Cloudy	47 m 1.9 m	173	100%>7m 4.4%<7m	0.58%	99ms

© 2011 California Institute of Technology. Government sponsorship acknowledged.

Fig. 17. Seven sequences of driving on cross-country trails were analyzed for false positive water detection using the sky reflection based detector with the false positive detection rejection algorithms enabled. Six are illustrated here and the seventh is illustrated in Fig. 14. In the 1711 frames processed over a 0.44km distance, false positive detection only occurred in 3 frames, a rate of 0.18%.

VII. Conclusion

A water body's surface can be modeled as a horizontal mirror, still water more so than disturbed water. The appearance of the surface of a water body in color imagery largely depends on the ratio of light reflected off the surface to the light coming out of the water body. When a water body is far away, the angle of incidence is large, and the light reflected off the water's surface dominates. We have exploited this behavior to detect non-traversable water bodies (such as large puddles, ponds, and lakes) that are out in the open on cross-country terrain at mid- to far-range.

This software estimates the horizon line using a **CAHV** camera model and an onboard inertial sensor, searches above the horizon line and detects the sky, segments candidate water pixels below the horizon line based on color and texture, geometrically locates the exact pixels in the sky that are reflecting on candidate water pixels on the ground, predicts if the ground pixels are water based on color similarity to the sky and local terrain features, and rejects detections that are likely false positives (based on slope, ground surface detection, sun glare detection, and the detection of AGC induced image saturation).

The geometry of modeling the surface of water as a horizontal mirror is straightforward. A ray of incident light reflects off the surface and enters a camera's FPA at a candidate water pixel. A direct ray from the tail of the incident ray (and within the same vertical plane as the incident ray) will enter the camera's FPA at a pixel whose color will indicate the color of the sky being reflected along the reflected ray. The angle between the direct ray and the reflected ray can be approximated as two times the glancing angle. The calculations to locate the pixel a direct ray enters involves, for any candidate water pixel, conversion of the 2D image coordinates to a 3D unit vector, negation of the z component of the unit vector, and conversion of the modified unit vector back to 2D image coordinates.

In three color image sequences collected while approaching a pond on different days when the sky was clear, overcast, and cloudy (from starting distances of 64, 51, and 47 meters, respectively), the true positive detection rate was 100% when the UGV was beyond 7 meters of the water's leading edge and the false positive detection rate was 0.58% for the cloudy sky sequence and 0% for the clear and overcast sky sequences. In seven color image sequences collected while driving on a variety of cross-country trails, the false positive detection rate was 0.18%.

Our practice is to run the sky reflection based water detector and a color variation based water detector [11] in parallel to detect water bodies in wide-open areas. These two detectors are complementary. The color variation based water detector performs well in detecting water bodies at close- to mid-range, and the sky reflection based detector performs well in detecting water bodies at mid- to far-range. Long range detection of non-traversable water enables a mid-range planner to make early, mild adjustments to a UGV's route, avoiding more drastic deviations often made by a local obstacle avoidance behavior. Consequently, detecting water hazards at far range generally reduces the time required to reach a goal position during autonomous navigation.

References

[1] M. Iqbal, O. Morel, and F. Meriaudeau, "A survey of outdoor water detection," *The 5th International Conference on Information & Communication Technology and Systems (ICTS)*, Indonesia, 2009.

[2] L. Matthies, P. Bellutta, and M. McHenry, "Detecting water hazards for autonomous off-road navigation," *Proceedings of SPIE*, Vol. 5083, Orlando, FL, pp. 231-242, 2003.

[3] T. Hong, C. Rasmussen, T. Chang, M. Shneier, "Fusing Ladar and Color Image Information for Mobile Robot Feature Detection and Tracking," *7th International Conference on Intelligent Autonomous Systems*, Marina Del Ray, CA, 2001.

[4] A. Pandian, "Robot navigation using stereo vision and polarization imaging," Master's thesis, Institut Universitaire de Technologie IUT Le Creusot, Universit_e de Bourgogne, 2008.

[5] A. Sarwal, J. Nett, and D. Simon, "Detection of small water bodies," *PercepTek Robotics Technical Report*, Littleton, CO, (Available at: http://www.dtic.mil, AD# ADA433004), 2004.

[6] B. Xie, Z. Xiang, H. Pan, and J. Liu, "Polarization-based water hazard detection for autonomous off-road navigation," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, pp. 3186-3190, 2007.

[7] A. Rankin, L. Matthies, and A. Huertas, "Daytime water detection by fusing multiple cues for autonomous off-road navigation," *Proceedings of the 24th Army Science Conference*, Orlando, FL, 2004.

[8] A. Rankin and L. Matthies, "Daytime water detection and localization for unmanned vehicle autonomous navigation," *Proceedings of the 25th Army Science Conference*, Orlando, FL, 2006.

[9] T. Yao, Z. Y. Xiang, J. L. Liu, D. Xu, "Multi-feature fusion based outdoor water hazards detection," *Proceedings of the IEEE Conference on Mechatronics and Automation*, pp. 652-656, 2007.

[10] T. Yao, Z. Xiang, and J. Liu, "Robust water hazard detection for autonomous off-road navigation," *Journal of Zhejiang University*, 10(6), pp. 786-793, 2009.

[11] A. Rankin and L. Matthies, "Daytime water detection based on color variation," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, 2010.

[12] Y. Yakimovsky and R. Cunningham, "A system for extracting three-dimensional measurements from a stereo pair of TV cameras," *Computer Graphics and Image Processing*, 7, 195-210, 1978

[13] A. Rankin, T. Ivanov, and S. Brennan, "Evaluating the performance of unmanned ground vehicle water detection," *Performance Metrics for Intelligent Systems Workshop (PerMIS)*, Baltimore, MD, 2010.

[14] A. Rankin and L. Matthies, "Passive sensor evaluation for unmanned ground vehicle mud detection," Journal of Field Robotics, 27(4), 2010.