

Negative Obstacle Detection by Thermal Signature

L. Matthies and A. Rankin

Jet Propulsion Laboratory, California Institute of Technology
Pasadena, CA. Tel. (818)354-3722 – Fax (818)393-4085
Email: [larry.matthies, arturo.rankin]@jpl.nasa.gov

Abstract

Detecting negative obstacles (ditches, potholes, and other depressions) is one of the most difficult problems in perception for autonomous, off-road navigation. Past work has largely relied on range imagery, because that is based on the geometry of the obstacle, is largely insensitive to illumination variables, and because there have not been other reliable alternatives. However, the visible aspect of negative obstacles shrinks rapidly with range, making them impossible to detect in time to avoid them at high speed. To relieve this problem, we show that the interiors of negative obstacles generally remain warmer than the surrounding terrain throughout the night, making thermal signature a stable property for night-time negative obstacle detection. Experimental results to date have achieved detection distances 45% greater by using thermal signature than by using range data alone. Thermal signature is the first known observable with potential to reveal a deep negative obstacle without actually seeing far into it. Modeling solar illumination has potential to extend the usefulness of thermal signature through daylight hours.

1. Introduction

Autonomous, unmanned ground vehicles (UGVs) are rapidly approaching practicality for a number of applications. UGV perception technology also has application in manned vehicles as a driver aid. A key open problem for UGV perception is detecting negative obstacles. At present, these are essentially impossible for UGV perception systems to detect in time to avoid them at driving speeds above a few kilometers per hour (kph). In this paper, we show that negative obstacles tend to be warmer than the surrounding terrain for most of the night; we also show that this enables detecting negative obstacles in conditions for which other approaches fail. UGVs for military applications are likely to carry thermal infrared cameras for other purposes, such as perception at night and through atmospheric obscurants. Commercial passenger cars are also now available with thermal infrared cameras as driver aids for night operation. Thus, thermal infrared signature has potential to revolutionize negative obstacle detection for both manned and unmanned vehicles, using sensors that are already in the sensor suite of both classes of vehicles for other purposes.

Most prior research on negative obstacle detection has been done by the small community working on autonomous off-road navigation. Past work has relied primarily on geometric analysis of range data [1,2], because these obstacles are largely defined by their geometry and because the most obvious alternative, intensity discontinuities, was considered highly dependent on illumination conditions and subject to false alarms. Informal discussion of other potential sensors for this problem, such as radar or polarization, has not, to our knowledge, led to successful or published results. The difficulty of seeing negative obstacles from ground level has prompted recent work under the DARPA Perceptor program to detect them from the air [3].

In section 2, we review the basic geometry of negative obstacle detection to show analytically what makes this such a hard problem for ground-based sensors. In a nutshell, considering the lookahead distance needed for safe driving, the angular resolution needed to see a negative obstacle grows with the fourth power of velocity; even then, there is no way to know how deep a hole is from the geometry of the visible portion. Section 3 illustrates the results and limitations to date of negative obstacle detection with stereo vision-based range data. We introduce a simple heat transfer model in section 4 that shows that the interior of negative obstacles will tend to be warmer than the surrounding terrain all night long. During the day, the interiors can be warmer or cooler than their surrounds, depending on sky conditions and sun position. Section 5 presents thermal imagery of natural and man-made negative obstacles that confirm the conclusions of the heat transfer model. In section 6, we show initial results of simple algorithms that combine geometric and thermal cues to achieve superior performance than possible with geometric cues alone. Section 7 summarizes our conclusions and near-term plans for additional work in this area.

2. Basic Geometry of Obstacle Detection

Figures 2-1 and 2-2 illustrate the geometry of positive and negative obstacle detection, respectively; as we will see, detecting negative obstacles is much more difficult than positive obstacles. A simple way to characterize obstacle detectability is in terms of the angle θ the obstacle subtends at the sensor. When the range R to the obstacle is large relative to the height H of the sensor above the

ground, so that angles θ and ϕ are fairly small, small angle approximations for trigonometric functions yield $\theta \approx (H - h)/R$, $\theta + \phi \approx H/R$, and therefore

$$\phi \approx \frac{h}{R}$$

where h is the height of the obstacle.

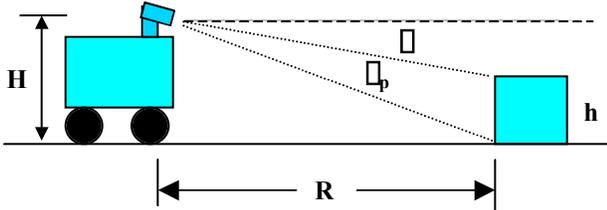


Figure 2-1. Geometry of positive obstacle detection.

For negative obstacles of width w , by the small angle approximation the depression angles are $\theta \approx H/(R + w)$ and $\theta + \phi \approx (H + h)/(R + w)$. By similar triangles, $h = Hw/R$. Combining these equations to solve for ϕ and eliminate h , we get:

$$\phi \approx \frac{Hw}{R(R + w)} \quad (1)$$

Since obstacle widths w can be less than 0.5 m and ranges R of interest are often greater than 5 m, the denominator will be dominated by the R^2 term. Thus, the angle subtended by a positive obstacle decreases by $1/R$ as the range increases, whereas the angle subtended by a negative obstacle decreases by $1/R^2$.

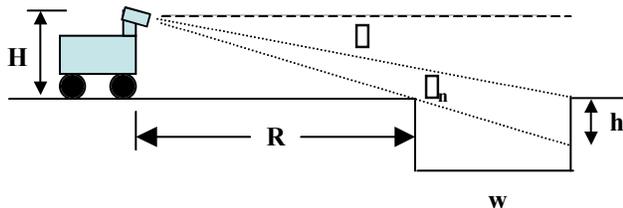


Figure 2-2. Geometry of negative obstacle detection.

The last step we need is to relate the required lookahead distance R to the vehicle velocity v . To have room to stop before hitting an obstacle, it is well known [4] that the lookahead distance needs to be at least:

$$R = \frac{v^2}{2\mu g} + vT_r + B \quad (2)$$

where μ is the coefficient of static friction between the wheels and the ground, g is gravitational acceleration, T_r is the total reaction time, and B is a buffer distance used for safety. For illustration, typical values of these parameters for off-road driving are around $\mu = 0.65$, $T_r = 0.25$, $B = 2$, and of course $g = 9.8$. For these values, the quadratic term will begin to dominate when $v > 3.2$ m/s ($R > 3.6$ m), which is about 11.5 kph (7 mph). Thus,

above this speed, ϕ_p will be dominated by $1/v^2$ and ϕ_n by $1/v^4$. This makes it very clear that negative obstacle detection is harder than positive obstacle detection; moreover, seeing the visible portion of a negative obstacle tells very little about its depth, so there is great potential for false alarms and missed detections.

3. Negative Obstacle Detection with Range Data

We will now show how negative obstacles typically appear in real range data, in this case from stereo vision. We will also describe an algorithm we use to detect them with range data and show the performance that can be expected with real data. Figure 3-1 shows an intensity image of a man-made trench dug to evaluate negative obstacle detection at Aberdeen Proving Grounds (APG) in 1997. This happens to be a thermal infrared image, taken with a cooled camera operating in the 3-5 μ m band. The image was taken in the late afternoon, after there had been sun on the visible side of the trench for several hours; note how bright it is relative to the surrounding terrain. The figure also shows range imagery from a stereo pair of these cameras. Note that there is an irregular band with no range data along the leading edge of the trench; this is typical, though sometimes there will be “mixed pixels” of range data that span the negative obstacle, filling in the range image.

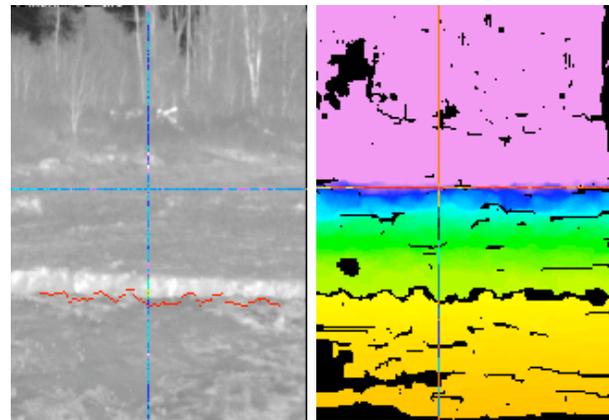


Figure 3-1. Left: thermal image of a trench 0.6 m wide viewed from 5.5 m away at a camera height of 1.0 m. Right: false color range image from stereo vision; yellow is closest, violet furthest, and black represents no data. Cross-hairs in both images are for reference to later figures. The red overlay on the intensity image shows the detection of the leading edge of the trench.

Figure 3-2 shows the range data projected as an elevation plot; each dot shows where one pixel of the range image projects on the ground plane. Height is shown in false color, but since the terrain was quite flat there is little variation in color in this figure. The negative obstacle clearly shows up as a gap in the elevation data. The bright band on the far (right) side of the trench is caused

by multiple range pixels from the vertical, interior wall of the trench projecting to roughly the same place in the elevation plot.

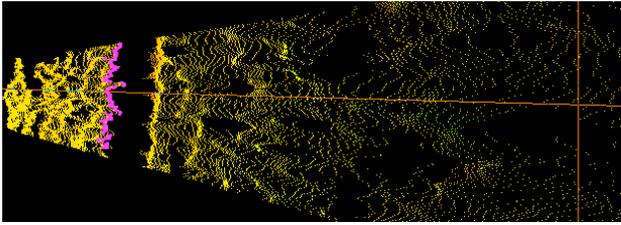


Figure 3-2. Elevation plot of the range data, seen from above. The camera was on the left, looking right. Magenta overlay shows detection of the leading edge of the trench. The cross-hairs mark the row and column marked by cross-hairs in figure 3-1.

Figure 3-3 shows an elevation profile of the column marked by the cross-hair in figure 3-1. The region of near-vertical pixels on the far side of the trench (the “upslope”) corresponds to the visible portion labelled h in figure 2-2. Negative obstacle detection algorithms based on range data generally exploit the gap in the data at the obstacle, the length and slope of the upslope, and the angle between the trend line of the terrain in front of the obstacle and the “downslope” (fictitious line marked red in figure 3-3) from the leading edge of the obstacle to the bottom of the upslope portion. If these measurements are all consistent with negative obstacle geometry as shown in figure 2-2, the pixels immediately bordering the gap are marked as a potential negative obstacle. This has to be “potential”, since all that is known about its depth is the length of the upslope, or h as shown in figure 2-2.

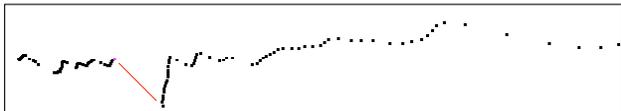


Figure 3-3. Elevation profile for the column marked by the cross-hair in figure 3-1, showing height vs. range, with the height axis exaggerated. The profile crosses the trench at the gap. The near-vertical portion on the far (right) side of the gap is the vertical interior of the trench; this segment is 13 cm tall. This profile has no mixed pixels, but if it did they would fall on the red line, filling in the gap across the obstacle.

In the previous section, we characterized the detectability of an obstacle by the vertical angle it subtends at the sensor. In practice, it is useful to translate that into the number of pixels subtended by the obstacle. In a large set of negative obstacle imagery taken at APG, we found that negative obstacles had to subtend at least 6 pixels for reliable detection. The stereo algorithm we used in those experiments had a correlation window 7 pixels high. This agrees with other qualitative observations we have made in the past that, with stereo, obstacles must be larger than the correlation window to be reliably detected;

otherwise, they get too blurred by the low-pass filtering and other artifacts introduced by the correlation window.

At 0.6 m wide, the trench shown in this section is large enough to be an obstacle for many wheeled vehicles of interest. The cameras used in this data collection had a vertical angular resolution of 2.39 mrad. When the 6-pixel detection rule is applied to a more complicated version of equation (1) that does not use small angle approximations, we get a prediction that this obstacle could not be reliably detected much beyond 6 m. Taking this as the lookahead distance requirement for equation (2) and using the typical parameter values we gave earlier, this implies a maximum safe driving speed of about 5.7 m/s (20.6 kph or 12.9 mph). Even this is optimistic, since the data set used to obtain our 6-pixel rule was based on ideal conditions involving vertical-walled trenches cut in level ground with negligible vegetation.

4. Simple Heat Transfer Analysis for Negative Obstacles

Given the limitations of negative obstacle detection with range data alone, having another observable feature that would either confirm detection or enable detection at greater range would very valuable. The bright interior of the trench in the thermal image in figure 3-1 gives a hint that temperature might be such a feature. Intuitively, negative obstacles are cavities that we might expect to retain heat. Hence, when the sun is not present, intuition suggests that open terrain might cool more quickly than the interiors of negative obstacles. This leads to the hypothesis that negative obstacles might be warmer than surrounding terrain most of the night. This by itself would be of significant value; by modelling the solar illumination during the day, it might be possible to push the usefulness of the temperature signature through much of the day. In this paper, we focus on confirming the intuition for night observations. The balance of this section gives a simple heat transfer analysis that supports the intuition above; section 5 confirms the analysis with real imagery.

The mechanisms of heat transfer are radiation, convection, and conduction [5]. In addition, evapotranspiration can significantly affect heating and cooling rates of vegetation and damp ground [6]. To date, we have only modelled the first three factors; it seems likely that they will dominate for holes that have existed for some time, hence dried out. This is consistent with our experimental results to date.

We will present the intuition behind our analysis before delving into the equations. Radiative heat transfer occurs when a surface exchanges energy by thermal emission and absorption with other surfaces “visible” to it in the hemisphere around it. Horizontal ground surfaces that are not overhung by other structures exchange heat with the sky over the full 2π steradians above them. On the other hand, the interior of a negative obstacle exchanges heat with the sky over only about half of the hemisphere and

with the other sides of the hole over the other half of the hemisphere. Since the sky will generally be colder than the ground, this means that open terrain will tend to cool faster than the interior of a negative obstacle. In still air, radiative transfer is the dominant mode of heat transfer [6].

Convective heat transfer occurs when air motion carries away heat that is conducted from the surface into the immediately adjacent air. The rate of heat transfer will depend on the rate of air motion. In general, we expect that there will be more air motion over the level ground than inside negative obstacles; heuristically, then, it is likely that convection will tend to cool open terrain faster than the interior of negative obstacles.

Conduction cools the ground surface by transferring heat into the subsurface of the terrain. Given the geometry of a negative obstacle, it is not immediately obvious how the effects of conduction will differ between open terrain and the interior of a negative obstacle. Basic considerations of heat flow rate for soils suggest that conduction has the smallest effect on the rate of surface cooling, so that whatever the relative effects on open terrain versus the interior of a negative obstacle, they are likely to be dominated by radiation and convection. Therefore, we expect that radiation and convection will work together to cool the open terrain faster than the interior of a negative obstacle, and that the effect of conduction will be small and unimportant in comparison.

For this initial study, we need enough of a physical model to provide a qualitative comparison of heat flux for these three heat transfer mechanisms, but we do not need accurate, quantitative predictions of actual temperature profiles. Therefore, very simple mathematical modelling will suffice to validate the intuition outlined above; experimental confirmation with real imagery taken over a 24 hour period will then confirm both the intuition and the theory.

4.1. Radiative Transfer

We will approximate the model for heat exchange between open terrain and the sky as two infinite, parallel planes, in which case the net heat flow is

$$q_r = \epsilon \sigma (T_s^4 - T_{sky}^4)$$

where q_r is in W/m^2 , ϵ is the emissivity of the soil, T_s is the soil temperature (in °K), T_{sky} is the effective sky temperature, and $\sigma = 5.67 \times 10^{-8} W/m^2K^4$ is the Stefan-Boltzmann constant. A nominal value of soil emissivity is $\epsilon = 0.9$. The effective sky temperature ranges from a low of 230°K under cold, clear sky conditions to a high of approximately 285°K under warm, cloudy conditions [5]. If we take the warm end of this spectrum ($T_{sky} = 285°K$) and let the soil temperature be $T_s = 300°K$, the heat flux is $77 W/m^2$.

For the interior of a negative obstacle, consider an idealized obstacle represented by a trench with vertical walls. A small surface element near the top of the wall will “see” the sky in half of its hemisphere and the inside of the

trench with the other half. It is easy to show that the net heat flux out of the surface element will be

$$q_n = 0.5 \epsilon \sigma (T_{s1}^4 - T_{sky}^4) + 0.5 \epsilon \sigma (T_{s1}^4 - T_{s2}^4)$$

where T_{s1} and T_{s2} represent the temperatures of the near and far sides of the trench. From our own measurements, the temperature difference between opposite sides of a half meter wide trench are only a few degrees; therefore, q_n will be close to $q_r/2$. That is, the heat flux from the walls of the negative obstacle is about half of that from the surrounding terrain. When the walls of the obstacle are not vertical, the contrast will be less.

4.2. Convective Transfer

In general, modelling convective cooling can be quite complex; however, to get an initial, qualitative assessment we can use the steady state form of Newton’s law of cooling with crude estimates of the convection coefficient [7]:

$$q = h (T_s - T_a)$$

where h is the convection coefficient in W/m^2K , T_s is the soil temperature, and T_a is the air temperature (in °K). h depends weakly on the temperature difference and the characteristic length of the surface the air is moving over, as well as on the nature of the air flow over the surface. From [7], for natural convection over a 0.3 m vertical wall in air (eg. inside of a negative obstacle), h is in the range 2 to 4, whereas for forced convection over a 1 m flat plate (eg. strong wind over open terrain) h is at least an order of magnitude larger. Using the value $h = 2$ and assuming a temperature difference of 10°C, we get $q = 20 W/m^2$ as an order of magnitude estimate for the heat flux due to convection inside a negative obstacle. For open terrain, the heat flux due to convection will be entirely dependent on wind speed, but we can safely expect it to be at least as large as inside the negative obstacle. In any case, the values are the same order of magnitude as our earlier estimates for radiative heat transfer.

4.3. Conductive Transfer

The basic heat flow relation for one-dimensional, steady state conduction is

$$q = -k dT/dx$$

where k is the thermal conductivity ($W/m°K$) and dT/dx is the temperature gradient with soil depth ($°K/m$). For soil, $k \approx 0.5$; for sand, $k \approx 0.25$ [5]. dT/dx will of course vary over the diurnal cycle due to the periodic solar heating. The remote sensing community has created models of temperature versus depth under these conditions [8], which could be used to find limits for dT/dx . We have not yet undertaken that derivation, but an adequate bound on dT/dx for our purposes can be estimated from modelling results presented in [8]. These show surface diurnal temperature variations ranging between 20°C and 50°C and show that effects of the diurnal variation only propagate about one meter into the ground. Using 35°C for the

daily surface temperature variation and approximating the daily maximum of dT/dx as half that value, then the maximum value for q is $O(10) \text{ W/m}^2$. It will be less than that for much of the night, because dT/dx will be less. Moreover, we are interested in the difference in conductive cooling of the open terrain versus the interior of a negative obstacle. Our measurements show temperature differences between the inside of a trench and surrounding terrain of less than 5°C ; this suggests a small difference in dT/dx , hence also a small difference in q , between the two locations. Thus, it appears that conductive effects on the temperature difference will be dominated by radiative and convective effects.

We conclude that radiative and convective cooling will both cause the interior of negative obstacles to cool more slowly than open terrain. The effect of conduction on the relative rates of cooling will be minor in comparison. Thus, following some transitional period after sunset, the interior of negative obstacles should be warmer than the surrounding terrain all night long, hence amenable to detection with thermal imagery.

5. Experimental Validation of Heat Transfer Model

We have used several data sets to confirm the model of section 4:

- Images taken through about half the night of a construction trench and man-made pothole in soil around JPL;
- Images taken over a 24-hour period from north, south, east, and west viewing directions of an L-shaped, man-made trench cut in soil by collaborators at SAIC in Denver;
- Images taken at various times of natural and man-made negative obstacles at Fort Polk, Louisiana.

All images were taken with cooled, 3-5 μm thermal cameras. Figure 5-1 illustrates results with the first data set taken at JPL. The trench was 0.5 to 1 m wide and over 1 m deep. Both sides and the bottom are clearly warmer than the surrounding ground surface; this suggests that thermal signature may be useful for both ground-level and airborne imagery. The pothole was about 0.4 m wide and deep. Note that open spaces under the construction debris are also warm in these images. The sky was clear.

To confirm these observations more systematically, we dug an L-shaped trench in soil, so as to capture the widest range of solar illumination conditions, and observed it from all four cardinal directions every hour for a 24-hour period. Figure 5-2 shows the trench. Sky conditions were clear for this period; sunrise occurred between 6 and 7 am and sunset between 7 and 8 pm. The full data set included observing the trench from six different distances, ranging from 4.6 to 18.3 m (15 to 60 feet). For all four viewing directions, the interior of the trench was warmer than the surround from 9 pm to 6 am. As expected, conditions varied during the day based on sun position and time of

day. Figure 5-3 shows sample images taken at 9 pm and 7 am for what should be the worst-case viewing condition for night-time thermal contrast, looking west. Since all other imagery taken between these times showed similar contrast, we do not show the rest here.



Figure 5-1. Top: color and 3-5 μm thermal infrared imagery of a construction trench at JPL, taken at 10 pm. Bottom: color and 3-5 μm thermal infrared imagery of a pothole dug in soil at the same construction site, taken at midnight.

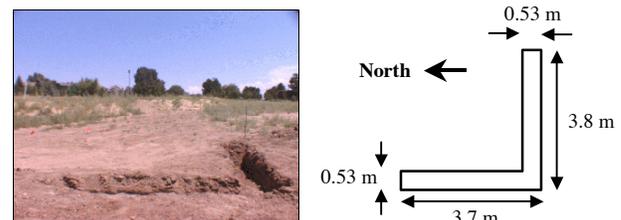


Figure 5-2. Man-made trench, 0.75 m deep, dug for 24-hour observations.

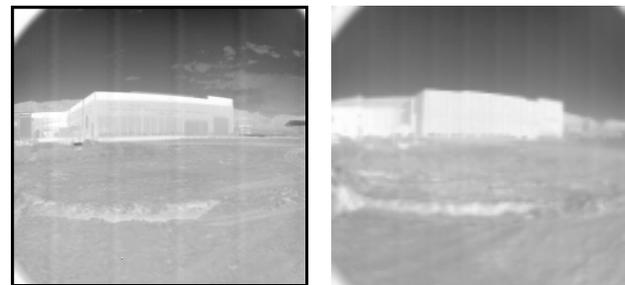


Figure 5-3. Worst-case viewing condition for night-time thermal contrast, looking west at north-south arm of trench. Left: 9 pm. Right: 7 am.

Finally, we acquired imagery of a number of natural and man-made negative obstacles during UGV field testing at Fort Polk, Louisiana, in November of 2003. The schedule did not permit systematic variation of viewing times and directions, as in the previous data set; however, this set includes more variety of terrain and weather conditions. Figures 5-4 and 5-5 show a sampling. Figure 5-4 is a natural erosion gully, or wadi, which was warmer than the rest of the terrain at 9 pm. Figure 5-5 is a man-made foxhole, which was very wet inside due to recent rain. It was not warmer at 10 pm, but was distinctly warmer when it was next observed, at 7 am.

In summary, all of these data support the theoretical model of section 4.

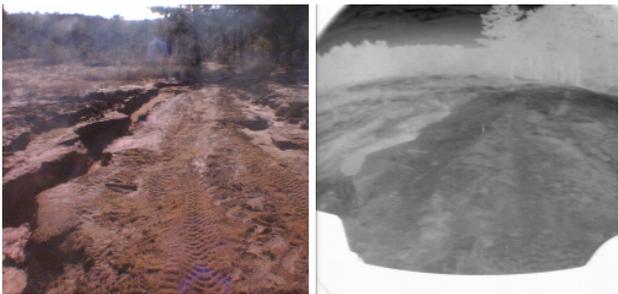


Figure 5-4. Wadi seen in (a) color at 9 am and (b) thermal IR at 9 pm. The bright area at the bottom of the IR image is structure on the vehicle.



Figure 5-5. Imagery of foxhole. (a) color, 5 pm. (b) Thermal IR, 5 pm. (c) IR, 10 pm. Structure in the bottom of the image is part of the vehicle. (d) IR, 7 am.

6. Obstacle Detection Results

Our work to show that the confirmed thermal signature actually enables improved obstacle detection is still at a fairly early stage. So far, we have developed a simple algorithm that looks for bright spots in thermal imagery that could be negative obstacles and applies simple geometric checks to rule out gross false alarms. The brightness test scans up each column from the bottom of the image, computing the running average brightness. When a pixel is found that exceeds the average ground plane intensity by a threshold, the pixel is flagged as the start of a candidate negative obstacle. The algorithm then continues to scan up the column until it finds the next pixel within the threshold value of the average ground plane intensity; this marks the end of the candidate obstacle. Tests on the relative geometry of the start pixel, end pixel, and any range pixels in between are then used to support or reject the negative obstacle hypothesis. Once all columns are processed in this manner, we group connected components of column-wise obstacle blobs into obstacle regions. One difficulty is that the obstacle interiors can be sufficiently warmer that they saturate the image, which prevents stereo from obtaining range data on their interior. In this case, we base geometric checks just on the relative height of the start and end pixels in each column. This rules out warm building structures as potential negative obstacles, but retains the saturated true obstacles.

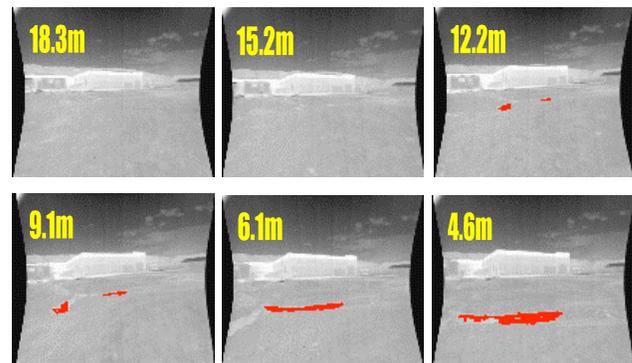


Figure 6-1. Trench detection results, 9 pm facing west.

We have applied this algorithm to the 24-hour trench data set and the foxhole imagery. Both data sets were taken with cameras 1.2 m above the ground and a vertical image resolution of 4.95 mrad/pixel. Note that this resolution is over twice as coarse as the camera used for the data set described in section 3, which was taken with different equipment several years earlier. Using the 6-pixel detection rule from section 3 and the geometric model from section 2, we predict that our standard geometry-based algorithm could detect the trench in figure 5-2 (0.53 m wide) to a distance of about 4.2 m. The trench data set includes imagery taken at 4.6, 6.1, 9.1, 12.2, 15.2, and 18.3 m from the trench. We processed the imagery for all distances, four times of night (9 pm, 12, 3, and 6 am), and all four approach directions. As illustrated in figure 6-1, the trench was detected reliably to 6.1 m in all but one

frame of the data set. Over all 96 frames in the data set, there were three frames with false alarms. These occurred on wheel tracks that were warmer than the rest of the terrain. Since the algorithm used here is extremely simple, we anticipate that such false alarms can be mitigated by more elaborate geometric analysis. Even with potential for such false alarms, the combination of thermal imagery and range data is superior to range data alone at detecting true obstacles. Figure 6-2 shows results with the foxhole imagery from Fort Polk. The foxhole was not detected as an obstacle with range data alone acquired by color or thermal stereo vision or by ladar. Figure 6-2c shows why stereo fails. Although there is a gap in the range data at the obstacle, there is no upslope, as there is for the example in figures 3-1 to 3-3; hence, the detector failed. Note that other areas that are not obstacles also produce gaps in the range data and elevation map. Figure 6-2d shows that using both thermal and geometric cues did enable detection.

7. Summary, Conclusions and Future Work

Negative obstacle detection is one of the key tall poles in perception for autonomous UGV navigation. By analyzing the basic geometry of both positive and negative obstacle detection, we showed why this is so hard: whereas the angle subtended by a positive obstacle decreases as $1/R$ with increasing range, for negative obstacles is decreases as $1/R^2$. Therefore, detecting negative obstacles with range data alone has a very short effective lookahead distance, which translates into low safe driving speeds.

We then presented a new approach to negative obstacle detection, particularly at night, based the observation that the interior of negative obstacles cools more slowly than the surrounding terrain. We included a simple mathematical model for heat transfer in and around negative obstacles that reveals this property and experimental results that confirm it with imagery acquired over a 24-hour period for a man-made trench. Less exhaustive observation of other natural and man-made negative obstacles also confirmed the property. The duration of the night for which negative obstacles are warmer does appear to be affected by weather and by the width of the obstacle; for example, rain during the day reduces temperature differences and very large negative obstacles appear to have smaller differences, as one would expect from our mathematical model. We developed a simple algorithm that combines thermal and geometric cues to detect negative obstacles and showed that it achieved superior detection performance to the use of range data alone.

We conclude that these results have fairly well established that negative obstacles do remain warmer than their surrounds at night under many conditions and that this does lead to better obstacle detection. This has important potential applications for UGV navigation and for driver aid for manned vehicles.

Future work needs to address several issues, including evaluating the heat transfer properties of negative obstacles under more weather conditions, developing more sophisticated algorithms to combine thermal and geometric analysis for negative obstacle detection, and modelling solar illumination to attempt apply thermal signature to day-time negative obstacle detection.

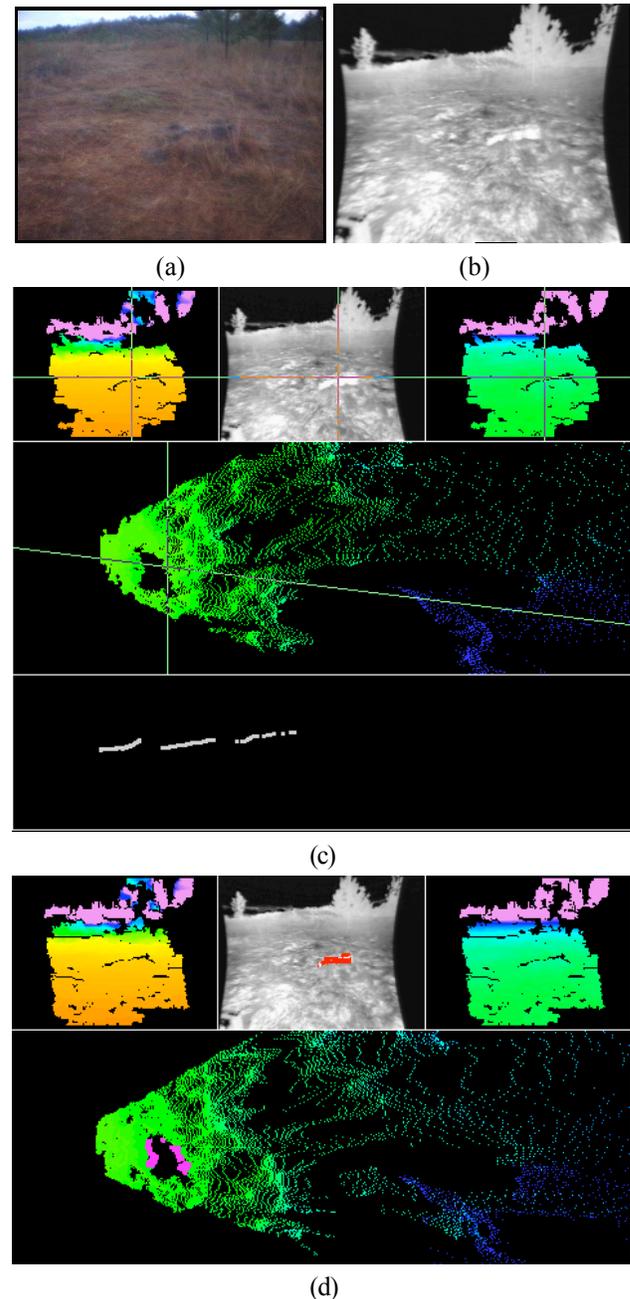


Figure 6-2. Foxhole detection. (a), (b): color and thermal imagery taken at 7 am at a distance of 2.8 m. (c) Results using range data alone (no detection). (d) Results with thermal and geometric cues (detection).

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the DARPA/TTO Perceptor Program through an agreement with the National Aeronautics and Space Administration. We thank members of the prime contractor team at SAIC in Denver for support with the data collections in Denver and Fort Polk. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

8. References

- [1] L. Matthies, A. Kelly, T. Litwin, and G. Tharp, "Obstacle Detection for Unmanned Ground Vehicles: A Progress Report", *7th International Symposium of Robotics Research*, Munich, Germany, 1995.
- [2] L. Matthies, T. Litwin, K. Owens, A. Rankin, K. Murphy, et al., "Performance Evaluation of UGV Obstacle Detection with CCD/FLIR Stereo Vision and LADAR", *IEEE Workshop on Perception for Mobile Agents*, Santa Barbara, CA, June 1998.
- [3] A. Stentz, private communication.
- [4] L. Matthies and P. Grandjean, "Stochastic Performance Modeling and Evaluation of Obstacle Detectability with Imaging Range Sensors", *IEEE Transactions on Robotics and Automation*, 16(12), December 1994.
- [5] F.P. Incropera and D.P. DeWitt, *Fundamentals of Heat and Mass Transfer*, John Wiley and Sons, New York, NY, 1985.
- [6] R.G. Reeves, A. Anson, and D. Landen (eds), *Manual of Remote Sensing*, 1st edition, American Society of Photogrammetry, Falls Church, VA 1975.
- [7] J. Lienhard and J. Lienhard, *A Heat Transfer Textbook*, Phlogiston Press, Cambridge, MA, 2002.
- [8] C. Elachi, *Introduction to the Physics and Techniques of Remote Sensing*, John Wiley and Sons, New York, NY, 1987.