

Performance of a scanning laser line striper in outdoor lighting

Christoph Mertz¹

Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, USA 15213;

ABSTRACT

For search and rescue robots and reconnaissance robots it is important to detect objects in their vicinity. We have developed a scanning laser line striper that can produce dense 3D images using active illumination. The scanner consists of a camera and a MEMS-micro mirror based projector. It can also detect the presence of optically difficult material like glass and metal. The sensor can be used for autonomous operation or it can help a human operator to better remotely control the robot. In this paper we will evaluate the performance of the scanner under outdoor illumination, i.e. from operating in the shade to operating in full sunlight. We report the range, resolution and accuracy of the sensor and its ability to reconstruct objects like grass, wooden blocks, wires, metal objects, electronic devices like cell phones, blank RPG, and other inert explosive devices. Furthermore we evaluate its ability to detect the presence of glass and polished metal objects. Lastly we report on a user study that shows a significant improvement in a grasping task. The user is tasked with grasping a wire with the remotely controlled hand of a robot. We compare the time it takes to complete the task using the 3D scanner with using a traditional video camera.

Keywords: structure light, 3D sensor, eye safety

1. INTRODUCTION

Search and rescue robots and reconnaissance robots need to be able to detect a whole range of natural (e.g. rocks, plants, dirt) and man-made (e.g. wires, glass, metal, electronics) objects. These vary in their size, shape, albedo and optical properties. We have developed a scanning laser line striper that is able to make dense 3D maps of objects and is able to classify their optical properties. In the paper [1] we gave a detailed description of the sensor. In this report we want to evaluate this sensor in detail, with an emphasis on the scanning of difficult materials that can be found in a rubble pile (glass, metal) and materials found in an IED or other explosive devices (wires, phones, pipes, ammunition)². We will also evaluate how the sensor can be used to improve remote controlled grasping tasks. We will first give a description of the sensor itself, explain the evaluation method and finally give the evaluation results.

2. DESCRIPTION OF SENSORS



Figure 1 Left: Camera and PicoP® projector. Middle: Scene in direct sunlight with one laser line. The ambient sunlight is suppressed by the fast shutter and a bandpass filter. Right: Reconstructed 3D scene, color indicates depth

The scanning laser line striper uses the principle of structured light [2]. It uses a PicoP® projector to illuminate the scene with laser lines, a camera to observe the line and a computer to analyze the image (Figure 1 left). A key to this system is the working principle of the projector. It has a laser beam that is steered by a micro-mirror steers to draw a video stream

¹ cmertz@andrew.cmu.edu, phone: 412-2687955

² No real ammunition or explosives were used in any of the experiments. We used blank ammunition, mockup IEDs or inert explosives.

at 60Hz per image and 30kHz per line. The camera takes images with a very short shutter, about 30 μ s. It therefore sees only one laser line (Figure 1 middle). Crucially, during the short shutter time only a small amount of the ambient light is integrated. The camera has a bandpass filter that suppresses the ambient light even further (in Figure 1 (middle) the laser line is clearly seen in direct sunlight). Lastly we employ background subtraction to remove the rest of the ambient light. The laser line is scanned by changing the trigger delay between the camera and the projector. Figure 1 (right) shows a complete 3D scan.

2.1 Resolution of structured light sensors

Structured light is a well-established method and their basic metrics are known. The resolution is:

$$\Delta r = (r^2/bf) \Delta d$$

With r the radius (distance sensor to object), b is the baseline (distance camera to projector or to second camera), f is the focal length of the camera and d the disparity. We will use $\Delta d=0.3$ to calculate the nominal resolution, it is the standard deviation if the error is a constant distribution of ± 0.5 pixel.

2.2 Camera and projector properties

The basic properties of the scanning line striper used in this evaluation are listed in Table 1.

Table 1 Basic properties of the scanning line striper used in our experiments.

	camera	projector
horizontal resolution	640	848
vertical resolution	480	480
update rate [Hz]	120	60
horizontal FOV [deg]	32	43
vertical FOV [deg]	25	25
focal length [pixels]	1100	1090
baseline [m]	0.09	
resolution at 1 m [mm]	3.0	
resolution at 0.2 m [mm]	0.1	

3. TECHNICAL ISSUES

3.1 Trigger issues for scanning line striper

There are two trigger issues, one is the stability of the projector sync out and the second is the programmable trigger delay of the camera.

The syncing of the projector and the camera needs to be accurate to about 0.1 μ s to get a stable single horizontal line. The current trigger we have is not stable enough which in the end causes the resolution to be 2 pixels instead of the expected 1 pixel. We are currently investigating if we can find a better sync out from the projector.

We change the trigger delay in software to scan the laser line. Sometimes the change in trigger delay is not fully in effect by the time the next image is taken and we get a different laser line than expected. This will cause the calculated z-position of the line to be incorrect. This problem can be circumvented by waiting long enough to make sure the trigger delay is fully implemented. But this will reduce the update rate by a factor of 2 or 3. Another solution would be to use an additional microcontroller to change the trigger timing instead of the trigger delay in the camera. This would involve additional hardware.

3.2 Update rate, CPU usage and power consumption

For about half the rows (1-245) the update rate is 60 Hz. For the other rows it is 30 Hz. The reason for this is that at one point the trigger delay plus the analysis time is more than one cycle time (17 ms). This could be improved in the future by running the trigger delay and the image capture in a different thread from the analysis.

We used a MacBook Pro with an Intel Core 2 Duo CPU at 2.26 GHz running Ubuntu 10.04 to test the CPU usage. The process takes 15% and 40% CPU power respectively of the two cores when running the scanner at 60 Hz (rows 1-245). At 30 Hz (rows 246 – 480) it is 6% and 23%.

The scanning line stripper uses the power from one USB for the projector (<5W) and one Firewire port for the camera (<2.5 W).

4. DATA ANALYSIS

4.1 Range, resolution, and relative error

Method

We placed a calibrated target at various distances in front of the scanning line stripper as shown in Figure 2.



Figure 2 Measurement setup. The scanning line stripper was sitting on one table. The calibrated target was placed on a cart.

The target consists of five squares with respective sizes of 5 cm, 10 cm, 20 cm, 30 cm, and 40 cm. The distance between the squares is 10 cm. The target is measured by the sensor and the measured widths, heights, and locations are compared to the ground truth. A typical raw point cloud from the scanning line stripper is shown in Figure 3. One can clearly see three full squares (two of them partially occluded) and one part of a square. Some spurious points are also present. This data was taken in daylight with few clouds and the background subtraction was not turned on.

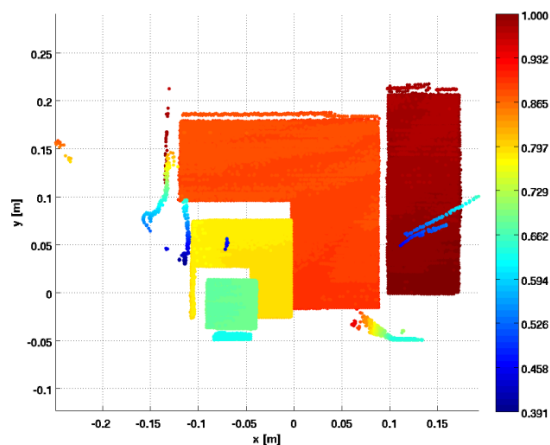


Figure 3 Raw point cloud of calibrated target at a distance of about 0.9m.

We apply cuts to the data to only retain the points that belong to the squares and we rotate the points so that the squares are parallel to the x- and y-axes. The clean and aligned data are shown in Figure 4.

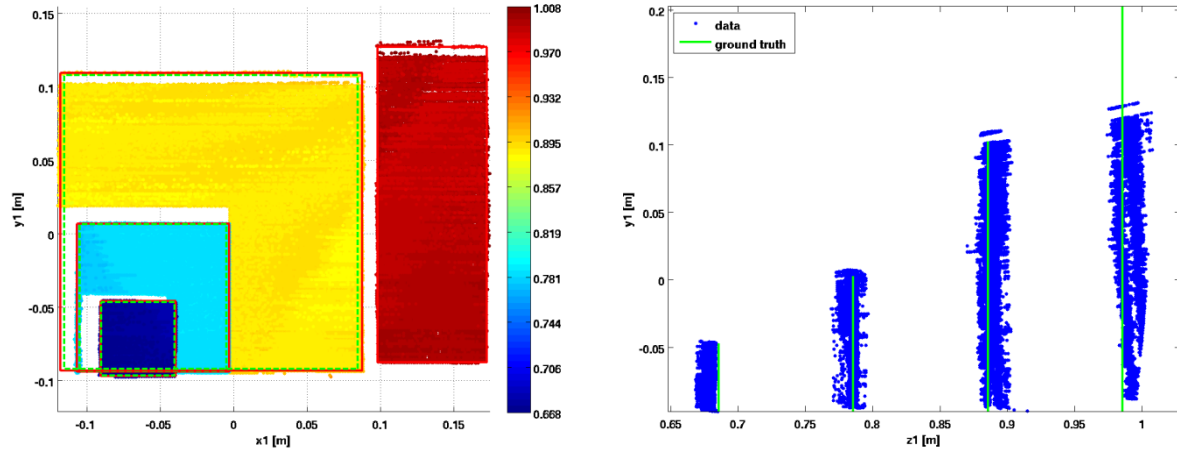


Figure 4 Front and side view of the clean data. The red boxes on the left indicate the measured edges of the squares. The dashed green boxes indicate the true size of the squares. The green lines on the right are 10 cm apart, indicating the ground truth.

Next we put additional cuts on the data to get the points for each square. A plane is fit to each square as can be seen in Figure 5.

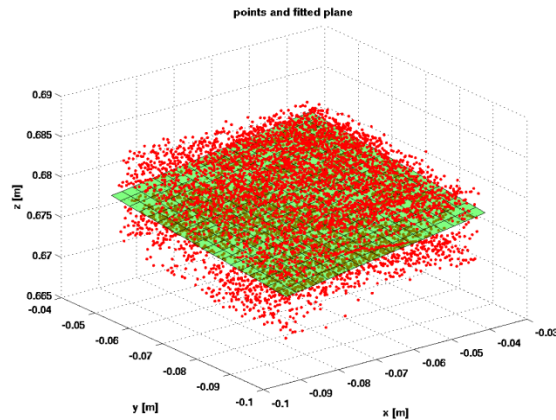


Figure 5 A plane is fit to the points of one square.

The standard deviation of the point-plane distances is a measurement of the resolution in the z-direction and the z-location of the planes is a measurement of the z-location of the square. These measurements can be done even if only parts of the squares are seen. The differences between the z locations are compared to the ground truth and are a measure of the z error.

Finally, lines are matched to the edges of the squares with following method:

1. Count the number of points N in the square.
2. The approximate number of points on each edge is $n = \sqrt{N}$.
3. The n^{th} leftmost (rightmost, highest, lowest) point is the measured location of the left (right, up, down) edge of the square.

The difference of the left-right (up-down) edges are compared to the ground truths and are a measure of the x (y) error. Figure 4 left shows the measured (solid red line) edges of the squares and the ground truth (dashed green line) size of the square. The ground truth of the absolute location of the squares is not known, the center of the green squares we placed at center of the measured (red) squares.

All the data taken with the calibrated target are shown in Table 2. We took data at distances of 0.5 m, 0.9 m, 1.7 m, and 2.6 m. All were taken outside in daylight. We took data during cloud cover and in full sunlight at a distance of 1.7 m. The data at a distance of 2.6 m was also taken in full sunlight.

Results: Resolution and relative error

Table 2 shows the setups and the raw data.

Table 2 Calibrated target at various distances and the resulting data.

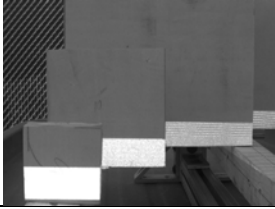
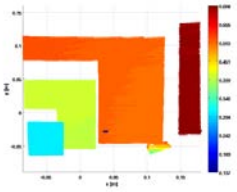
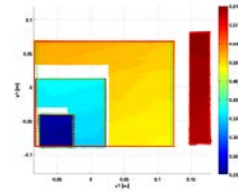
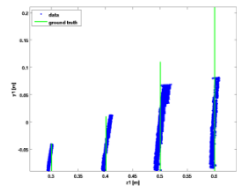
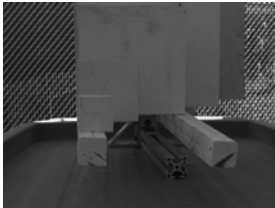
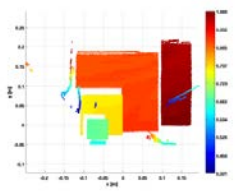
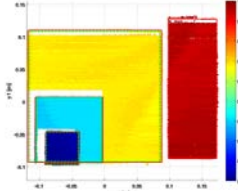
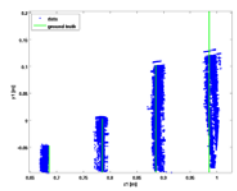

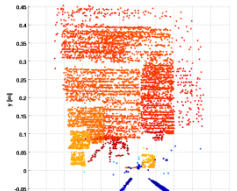
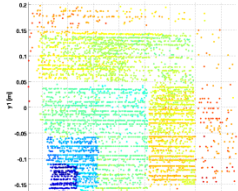

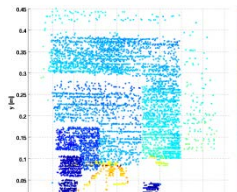
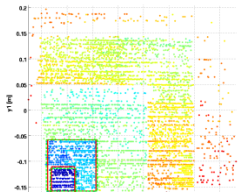
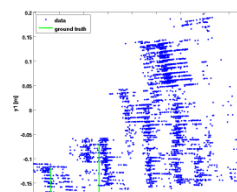

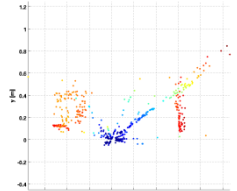
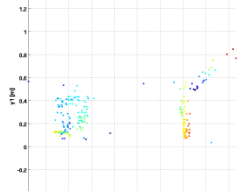
setup	image	All 3D points Striper: x-y, color = z	Cut 3D points, x-y, color = z, red = measured green = ground truth edges	Side view of cut points, green line = ground truth
Near: 0.5 m, under tent				
Medium: 0.9 m, outside, cloudy				
Far: 1.7 m, outside sunny				
Far: 1.7 m, outside cloudy				
Very far: 2.6 m, sunny				

Figure 6 (left) shows the resolution of the scanning line striper for different distances. As expected, the resolution increases with the square of the distance. The data is compared with the 1-pixel and 2-pixel resolution. The data points are between the 1.5- and 2-pixel resolution points. We discussed above the trigger issues of the scanning line striper as the reason that the resolution is not 1-pixel.

The relative errors are shown in Figure 6 (right). No dependence on the distance is apparent. The errors are around $\pm 5\%$. We believe that the error can be improved by improving the calibration. The side view of the data taken at 0.5 m (Table 2 upper right) shows a slight slant of the vertical planes, another indication that the calibration can be improved.

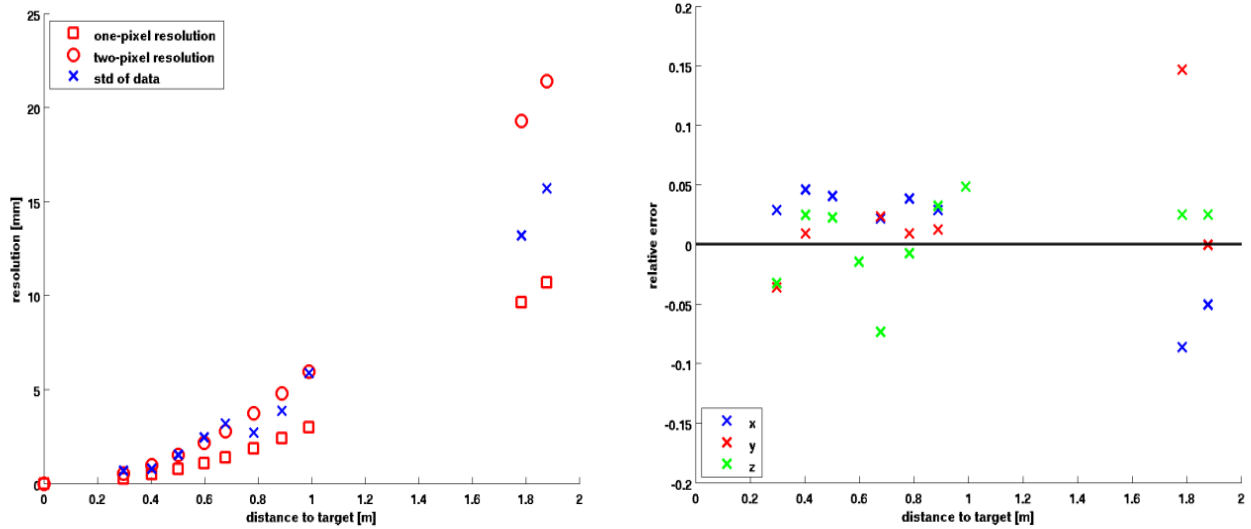


Figure 6 Left: Resolution vs. distance. The data is compared to the 1-pixel and 2-pixel resolution. Right: Relative error of x, y, and z for different distances.

Results: Range

The scanning line striper has no problems at 1 m distance. At the 1.7 m distance some deterioration of the quality is apparent; there are some missed points at the furthest plane. At 2.6 m the quality is very poor. Notice that this data was taken in bright sunlight. With more favorable ambient light the quality would probably be better.

4.2 Missed points and spurious points

At short and medium (around 1 m) distances the scanning line striper has basically no missed points. At far distances (1.7 m) one notices some missed points, especially for the farthest plane. At very large distances (2.6 m) there are hardly any points on the target.

There are some spurious points evident in the medium distance case. For that run we forgot to switch on the background subtraction and therefore the ambient light was not fully removed. Otherwise there are only spurious points in the sunny far and very far cases. The spurious points appear as points at close distances. We do not employ any algorithm that removes spurious points (e.g. de-noising), but will so in the future.

4.3 Various objects

In the next sections we investigate how good the sensor is in seeing various objects. We put an emphasis on difficult materials that can be found in a rubble pile (glass, metal) and materials found in an IED or other explosive devices (wires, phones, pipes, ammunition)². The setups, data, and comments are listed in Table 3 (glass and metal), Table 4 (wires), Table 5 (pipes and phones), and Table 6 (mockup IEDs, blank RPG, and inert explosives).

Table 3 Glass and metal.


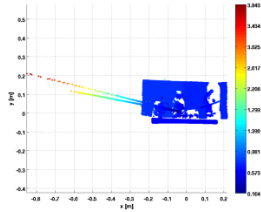
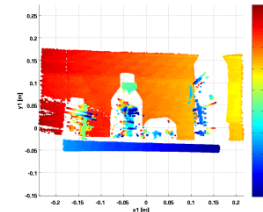
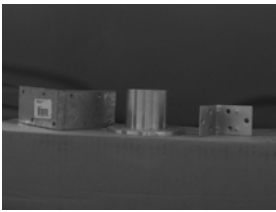
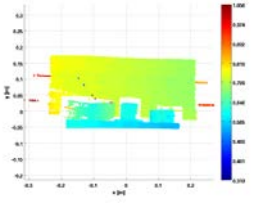
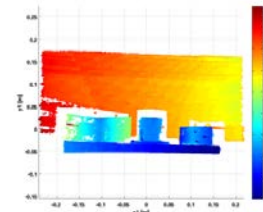
material	image	All 3D points x-y, color = z	Cut 3D points x-y, color = z	comments
glass				Lots of spurious returns, basically no correct points on glass.
metal				Few spurious returns. Small details like holes are apparent.

Table 4 Wires.

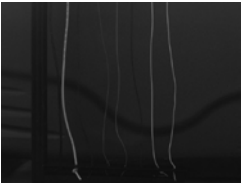
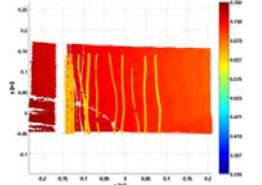
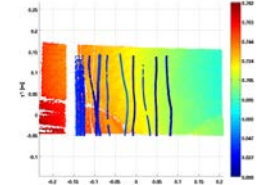

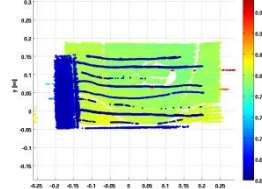
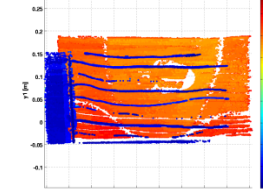

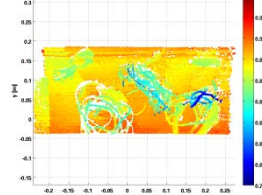
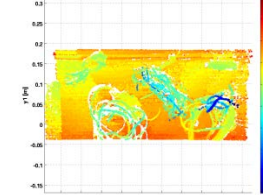
material	image	All 3D points Striper: x-y, color = z	Cut 3D points x-y, color = z	comments
Wires: vertical				All wires can be seen. Some missed points at the third (thin green) and fifth (thin black)
Wires: horizontal				All wires can be seen. Thick black: some misses. Thin green and black: many misses
Wire Bundles				Details of the wire bundles can be seen.

Table 5 Pipes and phones.

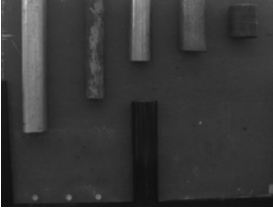
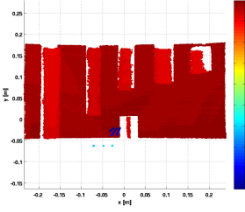
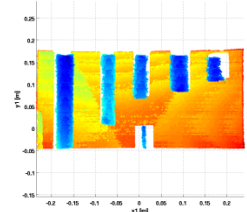

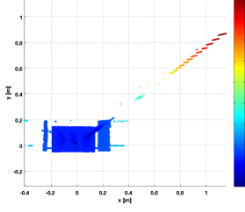
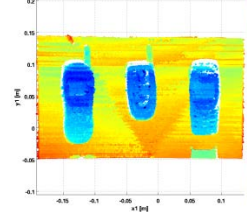

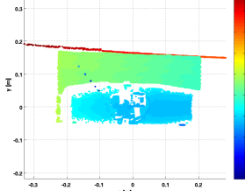
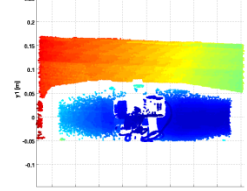

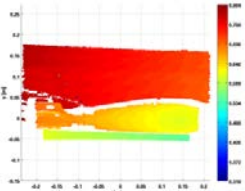
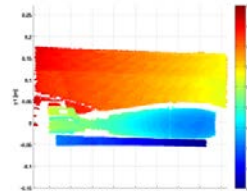

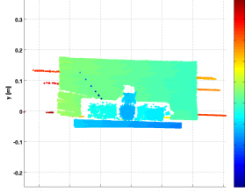
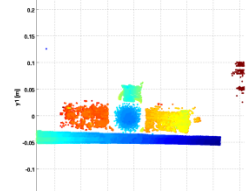
material	image	All 3D points Striper: x-y, color = z	Cut 3D points x-y, color = z	comments
Pipes				Details of pipes seen. Some spurious points. Black reflective pipe: Only points on top surface.
Cell phones				Details of phones can be seen. Some spurious points.

Table 6 Mockup IED, blank RPG, and inert explosives.

material	image	All 3D points Striper: x-y, color = z	Cut 3D points x-y, color = z	comments
Mockup IED				Details can be seen. One stray line (trigger delay issue)
Blank RPG				Details of blank RPG can be seen.
Inert Explosives				Details of inert explosives can be seen. Three stray lines (trigger delay issues)

Mockup IEDs in grass

In order to have a scenario that closely resembles a real-world event we placed a mockup IED in grass and observed it with the scanning line striper mounted on a the robot arm (Figure 7). The sensor was fully integrated into the robot. It received its power from the robot battery and all the computing was onboard. Wireless connections allowed us to control the data collection on the robot as well as a display of the data on a user interface.

Table 7 Mockup IEDs in grass.


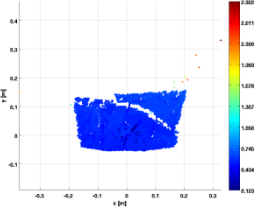
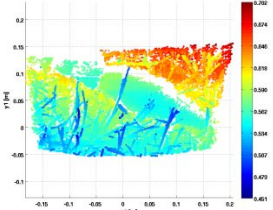

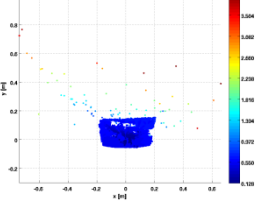
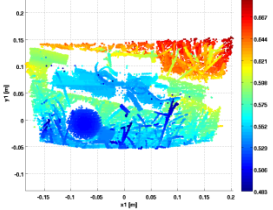
material	image	All 3D points Striper: x-y, color = z	Cut 3D points x-y, color = z	comments
Mockup IED in grass, in shadow				Details of grass and mockup IED can be seen. Some spurious points.
Mockup IED, blank Hand-grenade in grass, in shadow				Details of mockup IED, grass, and balnk grenade can be seen. Some spurious points.



Figure 7 Scanning line striper mounted on a robot arm and observing a mockup IED.

5. PRACTICAL USEFULNESS FOR GUI

We investigated how the 3D data from the striper can be used to improve the usefulness of the GUI compared to a common video stream.

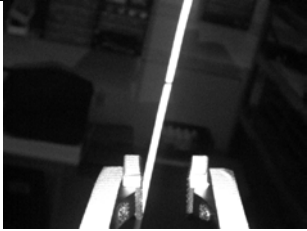
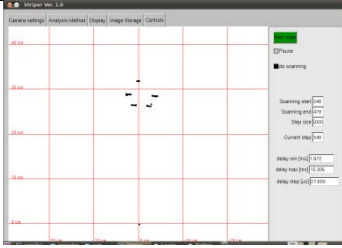
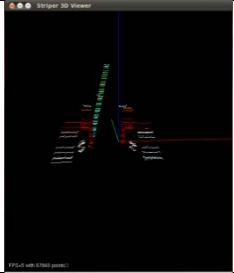
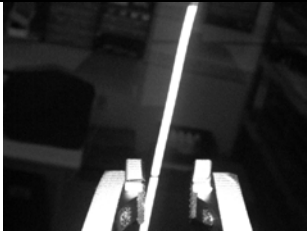
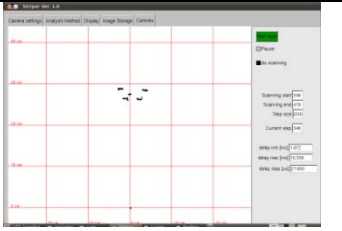
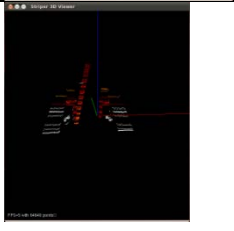
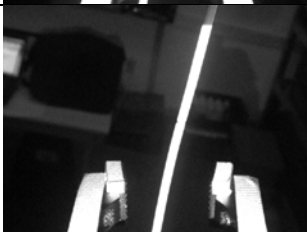
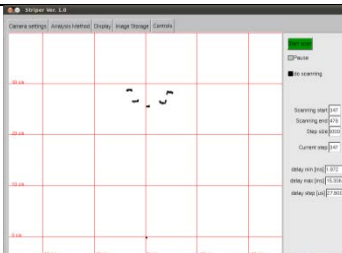
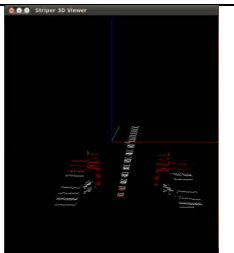
A white wire was hung in front of the robot (Figure 8). The user console was placed on a table so that the back of the operator was facing the robot. During the experiment the operator was only allowed to watch the video streams on the console. However, he was able to hear the robot move. The operator was using only two controls, one moved the robot arm back and forth and the other closed or opened the gripper.



Figure 8 Setup for GUI testing: A wire was hung in front of the robot while the operator (left person) was guiding the arm and the gripper.

In the experiment we tested how long it takes the operator to grip the wire when using three different visualizations (Table 8). The first was a live 30 Hz video stream from the striper camera. The second was a bird's-eye view of the wire and the gripper using one slice of the striper, also at 30 Hz. The third was a color coded 3D view with an update rate of about 1 Hz. The color code was as follows: White for distances outside the gripper and jet-color (i.e. blue to red) for distances inside the gripper. In Table 8 we show the different visualizations for three situations: The wire is too far, inside the gripper, and too close.

Table 8 Three different visualizations of the wire and gripper.

Wire position	video	Bird's-eye view	3D view
Too far			
In gripper			
Too close			

The operator received instructions on how to use the console and was allowed to practice for a short while. During the practice he was able to turn around and watch the robot. During the experiment itself he was allowed to watch the robot activity only through the console.

Table 9 Experimental results of GUI test. Time it took the test subjects to grasp the wire and their comments.

person	Video [s]	bird's eye [s]	3D [s]	comments		
				video	bird's eye	3D
1	57	10	44	noise was important, many tries, closing, went passed, back up etc.	much easier	time delay slowed things down
2	45	17	24	almost impossible to tell, need thickness of wire, need to probe	much easier	harder than bird's eye, with training this might be better
3	63	7	28	moderately difficult, especially when it is swinging	very easy	too much lag
4	65	9	11	moderately difficult	easy	easy, color coded is useful
5	72	15	-	very difficult to judge relative position		
mean	60.4	11.6	26.8			

The experimental results are shown in Table 9. On average it took the test subjects 60.4 seconds to grasp the wire using only the video. The fastest on average was the bird's eye view with 11.6 seconds and with the 3D view it took on average 26.8 seconds. With the bird's eye view the grasping was significantly faster (more than 5 times) which was also reflected in the comments of the test subjects. It turned out that while using the video only the users employ a trial and error approach and often used the grasper as a probe. E.g. the grasper was closed and moved towards the wire. It was noticed when the grasper touched and moved the wire. Then the user knew that the grasper was close enough. The 3D view was slower than the bird's eye view. From the users comments it appears that the main reason for that was the lag of it, i.e. it was updating only about once a second.

Overall it is clear that the depth information from the striper improves the grasping significantly.

6. CONCLUSION AND OUTLOOK

The evaluations in the previous sections showed that the scanning laser line striper performs well outside while remaining eye-safe, even in direct sunlight. It is well suited to make 3D maps of natural objects like grass and man-made objects like IEDs and their components (explosives, wires, cell phones)². Its performance can still be improved for optically challenging materials like glass. This 3D sensor can also significantly improve the performance of remote controlled grasping tasks compared with video.

We are still doing active research on the scanning laser line striper to improve its performance. Using the next generation PicoP® projector and USB3 cameras will significantly increase the update rate and resolution of the sensor. Further software development will also enable us to get 3D maps of more optically challenging objects.

ACKNOWLEDGEMENT

This work was conducted through collaborative participation in the Robotics Consortium sponsored by the US Army Research Laboratory (ARL) under the Collaborative Technology Alliance Program, Cooperative Agreement W911NF-10-2-0016.

REFERENCES

- [1] C. Mertz, S. Koppal, S. Sia, and S. Narasimhan, "A low-power structured light sensor for outdoor scene reconstruction and dominant material identification," 9th IEEE International Workshop on Projector-Camera Systems, June, 2012.
- [2] P. M. Will and K. S. Pennington. "Grid coding: A preprocessing technique for robot and machine vision." AI, 2, 1971