

A low-power structured light sensor for outdoor scene reconstruction and dominant material identification

Christoph Mertz*

Sanjeev J. Koppal[†]

Solomon Sia*

Srinivasa Narasimhan*

Abstract

We introduce a compact structured light device that utilizes a commercially available MEMS mirror-enabled hand-held laser projector. Without complex re-engineering, we show how to exploit the projector’s high-speed MEMS mirror motion and laser light-sources to suppress ambient illumination, enabling low-cost and low-power reconstruction of outdoor scenes in sunlight. We discuss how the line-striping acts as a kind of “light-probe”, creating distinctive patterns of light scattered by different types of materials. We investigate visual features that can be computed from these patterns and can reliably identify the dominant material characteristic of a scene, i.e. where most of the objects consist of either diffuse (wood), translucent (wax), reflective (metal) or transparent (glass) materials.

1. Introduction

Structured light has been a popular research topic for many decades [28]. However, some difficult research challenges still remain. Ambient illumination can be a problem; for example, the Microsoft KinectTM works quite well under cloudy illumination and at night, but struggles under direct sunlight [2]. Another issue is the presence of materials that exhibit complex visual effects. While the separation of global illumination [20] can address some of these issues, highly reflective or refractive scenes with many caustics can still create errors. In this paper, we show that an off-the-shelf, commercial, visible light projector can be used to create a structured light sensor that allows 3D line-striping even in the presence of direct sunlight. When the scenes have complex material properties, our sensor obtains visual features that help identify scenes that have a dominant material property. By “dominant” we mean that a user-defined fraction of the visible scene points belong to either diffuse, translucent, transparent and highly reflective surfaces.

Our key idea is to exploit certain characteristics of Microvision ShowwxTM hand-held projectors. Recently these devices have enabled new possibilities [3] since they have the advantages of low-power and very large depth of field.

*Carnegie Mellon University

[†]Harvard University

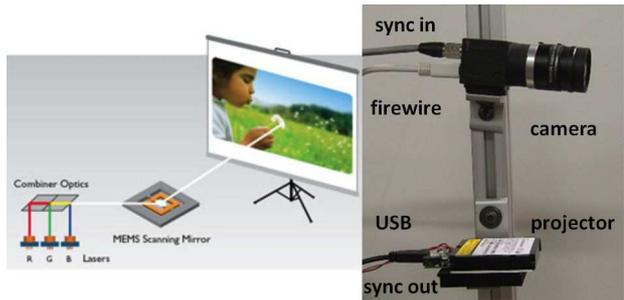


Figure 1. Left: Workings of the ShowwxTM projector (from Microvision website [1]). Right: Our structured light sensor with a Flea3TM camera and the ShowwxTM projector.

While our approach also has these features, we are, instead, interested in the ShowwxTM’s steering micro mirror that controls red, green and blue lasers [1] at high frame rates (18kHz horizontally and 60Hz vertically). These features offer three advantages for removing ambient light. First, since the lasers have a narrow-bandwidth, an appropriate filter placed on the camera can block some of the unwanted ambient light. Second, if we block much of the remaining ambient light with a low image exposure, the very high horizontal frequency of the MEMS steering device still allows the detection of a horizontal line stripe. Finally, the vertical frequency of the MEMS mirror allows background subtraction at near real-time rates, further removing all visual input except the line stripe. This final sensor line-striping data appears as if taken in the dark.

Such a sensor that enables line-striping in ambient illumination would be useful to a variety of platforms. These include search-and-rescue robots that find and dispose of dangerous materials. These robots work outside, perhaps in daylight, and need to detect various materials. In addition, they are limited in payload and power. Our sensor enables two applications for these types of platforms.

The first is low-power and low-cost reconstruction of diffuse scenes under strong ambient lighting (e.g. direct sunlight). Most active 3D devices that work in the outdoors use time-of-flight sensors (an exception is [17]), which are expensive in terms of cost and power consumption. Our sensor uses a hand-held projector, and therefore consumes less power. Furthermore, the fast micro mirror can potentially allow for near real-time reconstruction.

The other application of our sensor relates to the scene’s material properties. Since the measured data appears without ambient light, we can detect subtle visual effects, such as caustics, that would otherwise be washed out. In this sense, our sensor is a kind of “light-probe”. For future work, we would like to exploit this type of data to extend previous research on reconstruction of glass [18], environment matting [4] and relighting [25] to scenes with ambient lighting. In this paper, we explore the visual features that show preliminary promise for disambiguating scenes into groups that have a dominant material property. We will consider four types of effects: diffuse reflectance, sub-surface scattering/medium scattering and reflective or refractive caustics. We investigate a variety of materials such as paper, wood, metals, glass and milk and demonstrate the potential of our sensor with two heuristically based algorithms, one fast and one off-line, for dominant material classification.

1.1. Previous work

Projectors are commonly used in reconstruction [28], relighting [10], light transport analysis [20, 25] and defocus [9]. Our paper is related to a recent trend towards building projector-camera systems that utilize reflective micro-mirrors for high-speed control of light. For example, [19] build a programmable camera by exploiting a projector’s DMD (digital micro-mirror device) and show applications such as adaptive dynamic range and optical filtering. A single-pixel camera has been demonstrated by applying random projections onto incoming radiance with DMDs [26]. Interactive applications [5, 23] are possible when the illumination changes faster than humanly perceivable. DMD kits are used in stereoscopic displays [11] and high-speed range finding [15]. Off-the-shelf projectors can be used for reconstruction of dynamic scenes [14].

Unlike these previous efforts, which exploit the optical features of larger projectors, we are interested in the new wave of smaller, more compact designs. The optics of these projectors are fundamentally different from larger projectors: many use laser light sources and do not suffer from focusing issues. Recent efforts in the fields of augmented reality, human-computer interaction and ubiquitous computing exploit the properties of these new projectors [8, 24, 6]. However, almost all this previous human-centered work is targeted at indoor scenes (such as offices). Instead, we want to enable vision applications, outdoors and under sunlight.

Recently, there has been work on classification of glass and other materials based on applying learning techniques to visual cues obtained by passive sensors [16, 7, 12]. Unlike our active sensor, however, once these methods detect the lack of complex material (diffuse scene), they cannot obtain depths. [27] obtains BRDF slices for classification, but requires an illumination dome, while our setup is low-power and mobile.

2. Sensor description

We use the ShowwxTM evaluation kit (Fig. 1), which contains a PicoP display engine and has a sync connection that is used to trigger a 480x640 120Hz Flea3TM camera. The projector consists of three lasers, red, green, and blue, that are optically combined into one beam. The color of the beam is controlled by changing the intensity of each laser. A MEMS scanning mirror steers the beam horizontally and vertically and thereby produces the projected image. We investigated the distortion of this projected image and found that it is not well described by the standard radial and tangential lens distortion. It appears that the manufacturer adjusted the timing for drawing the horizontal line pixels so that there are no vertical distortions. Therefore, vertical lines are straight and horizontal lines have a “bow-tie” distortion. To correct for the distortions we created a lookup table that corrects each individual pixel. In all the experiments in this paper, the projector displays only a single image of uniform intensity.

An advantage of this projector is that its displayed images are in focus everywhere. Another feature, one we explicitly take advantage of, is that the high-speed MEMS scanning mirror illuminates each scene point only for a very short time. The update rate per image is 60Hz in scan line order and, since there are 480 rows, it projects 28.8k lines/s. Each line takes 28 μ s to draw and each pixel is illuminated for only 33ns. After each image nothing is projected for about 3.4ms in order for the mirror to return to its origin. This is a 15 lumens projector with nominal maximum output power and laser wavelengths as R: 90mW/635nm, G: 60mW/526nm, B: 50mW/440nm. We measured the wavelength of the red laser to be 637.9nm at room temperature (22.6°C) with a FWHM of 0.8nm. The projector eventually heats up and the wavelength drifts to 644.2nm (FWHM 1.0nm), showing a temperature drift of 0.1nm/°C.

3. Ambient light suppression

One challenge with structured light setups in the outdoors is the very strong ambient lighting, e.g. sunlight. There are 5 common ways to increase the signal-to-ambient light ratio: (1) Increase the power of the projector’s light source (2) Use a wavelength where the ambient light is low (3) Narrow band filter (4) Pulsed light source and fast shutter (5) Background subtraction.

For completeness we want to mention that there are also the methods of spatial differencing, using polarization and spectral estimation (see [21] for a discussion). The power of a laser is restricted by eye-safety concerns. For our particular projector we choose the red laser which at its maximum power of 90mW is brighter than the green and blue lasers. The ambient light from the sun is usually similar across red, green and blue wavelengths. Compared to 635nm (red)

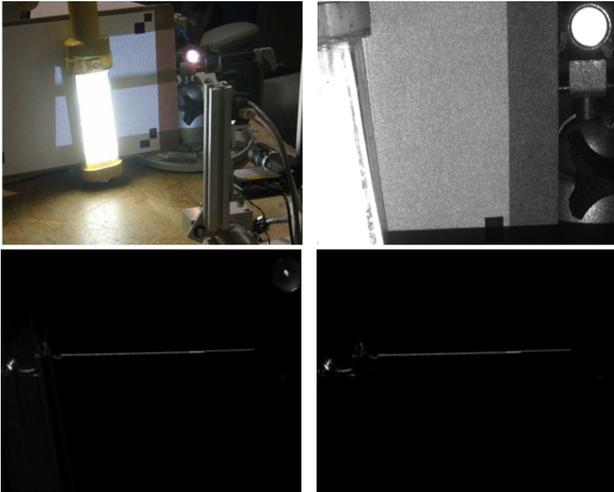


Figure 2. Ambient light suppression: Upper left: Picture taken of the scenario. A fluorescent lamp and a flashlight are in the field of the projection. Upper right: Scene taken with 18ms shutter. The direct light saturates the image. Lower left: Scene taken with $56\mu\text{s}$ shutter. Lower right: Scene taken with $56\mu\text{s}$ shutter and background subtraction.

there is 7% more ambient light at 526nm (green) and 7% less at 440nm (blue). To have significantly less ambient light one needs to go to the UV or IR. A particular interesting wavelength is 938nm where there is 78% less ambient light from the sun. An ambient reduced projector could be specially constructed for that wavelength.

We are using a bandpass filter with a central wavelength of 640nm and a FWHM of 20nm. This bandwidth takes into account the uncertainty of the central wavelength (20% of bandwidth), the drift of the laser with temperature and that the central wavelength of the filter changes with incident angle. The fact that the projector draws the lines sequentially (see previous section) makes it effectively a pulsed light source. For one line the system is on for $28\mu\text{s}$ and off for the remaining $16640\mu\text{s}$, a ratio of about 1/600. We use a camera that has a shutter down to $30\mu\text{s}$ and is triggered by the projector. We can choose different trigger delays in software to select the line we want to observe. It turns out that the trigger is not stable enough at these short times and we need to extend the shutter time to $56\mu\text{s}$. That gives us a consistent line that is 2 pixels thick.

Lastly the background can be suppressed by taking an image with the projector on and one with the projector off. The second image is subtracted from the first to get the response without background. It is not actually necessary to shut the projector off; instead, we choose a different trigger delay which effectively moves the location of the projected line. In this way, one gets two images with the same background but with different projected lines. Subtracting one from the other and keeping only the positive values gives us a single line-stripe.

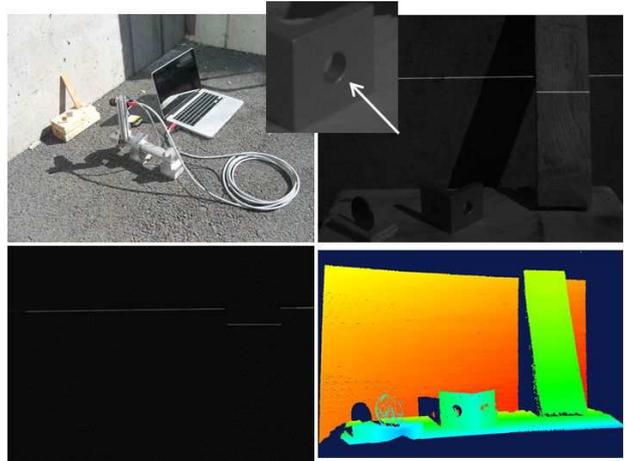


Figure 3. We show a setup (top left) to reconstruct a scene in direct sunlight. At a shutter speed of $56\mu\text{s}$, the ambient light is greatly reduced and one line stripe remains visible. There is some specular reflection off the metal bracket from sunlight (see insert). Background subtraction further reduces the ambient illumination and specular reflection (lower left). The quality of the remaining line allows for the reconstruction of the scene (lower right).

Fig. 2 demonstrates how effectively the ambient light can be suppressed. A fluorescent lamp and a flashlight are positioned in the field-of-view of the sensor (Fig. 2 upper left). Even with the bandpass filter in the camera the direct light saturates the image when the shutter is 17ms (Fig. 2 upper right). Notice that at this shutter speed we can still see most of the scene. When the shutter time is reduced to $56\mu\text{s}$ only a line of the projection is seen (Fig. 2 lower left). The ambient light is also greatly reduced. The fluorescent light is barely visible and from the flashlight only the light bulb itself is clearly seen. After image subtraction, there is no more background (Fig. 2 lower right). Notice that the projected line is visible even directly on the fluorescent lamp; the brightness of the lamp does not destroy any information.

The ambient light suppression we have described enables our sensor to reconstruct scenes in daylight. In Fig. 3, we show our sensor reconstructing a scene in bright sunlight. Note that the system is powered simply by a laptop: it is a mobile scanner. Also, when the projector is switched on, the projected uniform image cannot be seen. However, the ambient light suppression is effective (Fig. 3 upper right and lower left) allowing the successful detection of the line and the reconstruction of the scene (Fig. 3 lower right).

At a shutter of $56\mu\text{s}$ the projection is a 2 pixel-wide line and is clearly noticeable. Notice that the metal bracket has a small bright spot due to the specular reflection of the sun (Fig. 3 insert). This spot causes spurious data points if it is not eliminated through background subtraction. We have calculated the sum-of-squared differences for a patch on the wall in the figure (by fitting a plane to it) which is 0.8mm at a distance of about 0.5m.

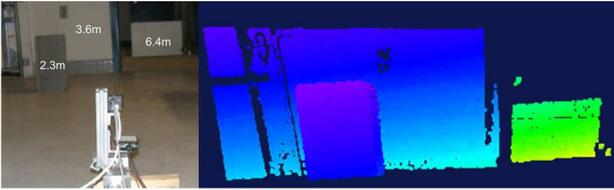


Figure 4. Indoors, our sensor can reconstruct scenes within a 6 meter range. We have noticed that this range is approximately halved in sunlight.

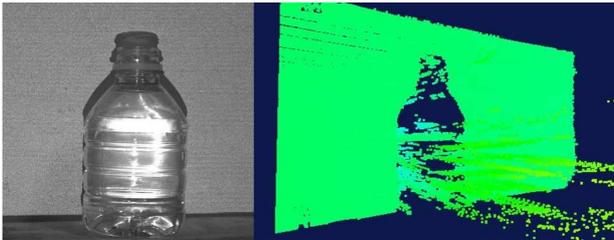


Figure 5. We show a transparent scene, whose complex refractive properties cause our sensor to give incorrect depths.

Our sensor has a range of around 6 meters for diffuse scenes, as we demonstrate with a white poster board as target and moderate ambient light (see Fig. 4). In bright sunlight, this maximum range decreases to about 3 meters. Note that we obtain this with a 15 lumens projector in broad daylight, and that the Kinect does not work for even short ranges in such bright sunlight [2]. We mounted the sensor on a 100 lb reconnaissance robot and showed that its low power and weight makes it a viable sensor for the target platforms.

In Fig. 5 we show the reconstruction results for a plastic bottle with water. This scene creates caustics when a line plane intersects it, creating many errors in the reconstruction. While we cannot reconstruct these materials, we are able to identify that these regions are erroneous. Furthermore, we will show how analysis of the caustics and reflected patterns can be used to classify scene materials.

4. Identifying a scene’s dominant material character

In this section, we demonstrate how our system acts as a “light-probe”, allowing optical measurements of a scene’s illumination patterns, that are usually washed out under ambient light. Distinctive visual features are computed from the observed light patterns and could be used to identify the material properties of a scene.

We will outline the potential power of these visual features by focusing on a simpler sub-problem: finding a scene’s dominant material characteristics. We will assume that the scene is placed in front of a diffuse background and that most of the objects display a common optical characteristic. Our goal will be to identify scenes with a “dominant” property, i.e., where a user-defined fraction (or higher) of

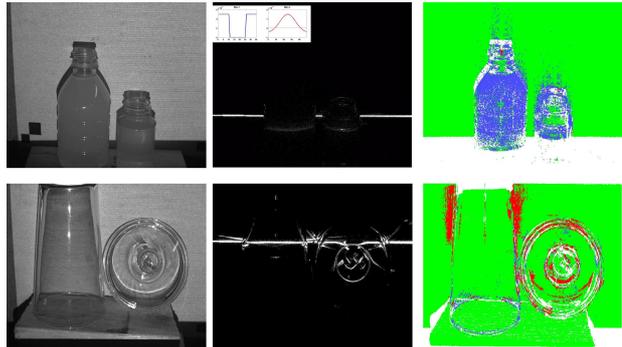


Figure 6. We show fast material classification for two scenes, a pair of bottles of milky water and a scene with glass objects. The middle frames are a single image from the line-stripe scan. The inset in the first row shows the two filters we use. Green=diffuse, blue=milky, red=glass.

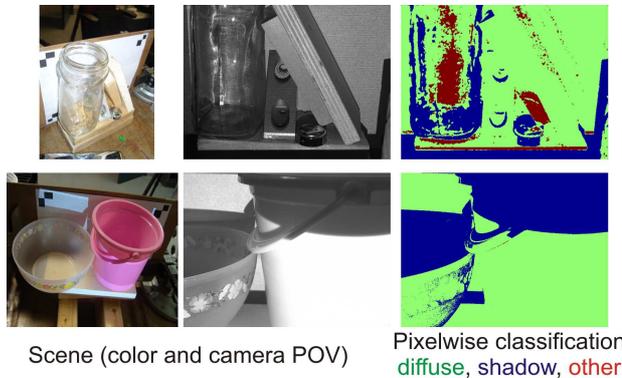
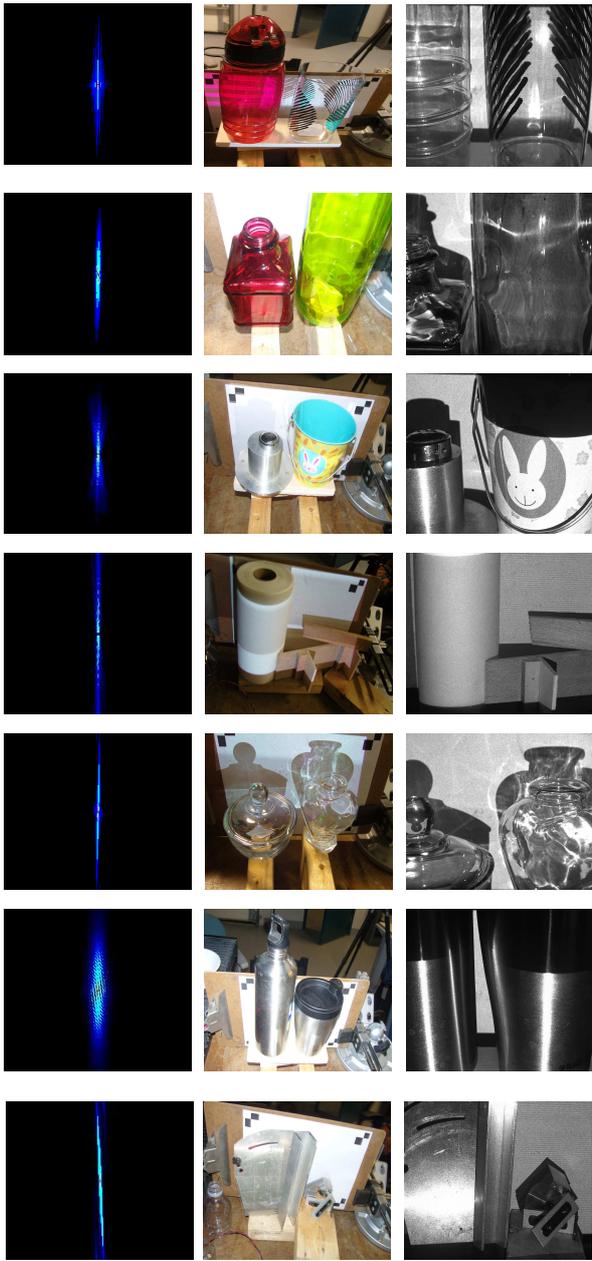


Figure 7. The top row has a scene with glass and metal. The bottom row has diffuse bowls. We classify each pixel by looking at the number of intensity maxima. Zero maxima imply a shadowed region. A single maxima implies diffuse/lambertian and many imply a more complex material.

the observed scene points belong to either (1) lambertian materials, (2) dispersive media and subsurface scattering materials, (3) reflective surfaces or (4) refractive surfaces. We set the default fraction to be $\frac{1}{2}$ (majority), and this can be varied for other applications.

In our experiments, the scenes are static. However, since our sensor scans the scene at a 60Hz, the features we investigate could also be used for dynamic scenes. We present a two-tiered, heuristic-based approach to classifying complex materials. The first is a per-frame fast method, while the second takes in as input a whole scan of the scene. The first method makes no assumption as to the scene geometry. The second method applies fourier analysis, and makes the assumption that the scene geometry has low spatial frequencies. This is a weak, smoothness assumption on scene geometry and, from our empirical results, only prevents us from recognizing the most difficult scenes (such as those dominated by broken glass).

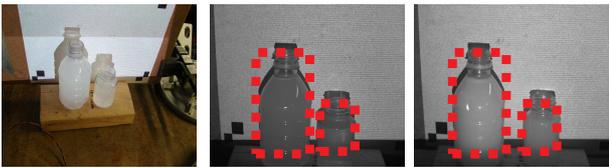
The advantage of the first method is its speed, while the



Glass, plastic and metals (power spectrum has large support)

Figure 8. We take the power spectrum of the three dimensional Fourier transform of each scan video, and integrate the time frequency dimension. The resulting 2D matrix is mostly sparse. Low non-zero support gives an indication of scattering and subsurface scattering.

second approach could be used for in-depth analysis of the scene materials. In particular, we point out that our ability to remove the ambient light is critical, and allows us to measure high-frequency reflection and refraction effects, such as caustics, that might be too faint to notice for a 15 lumens projector under ambient illumination. It is these types of effects that allow us to classify and detect challenging ma-



Scene with two concentrations, showing analyzed patches



Scattering media show small support

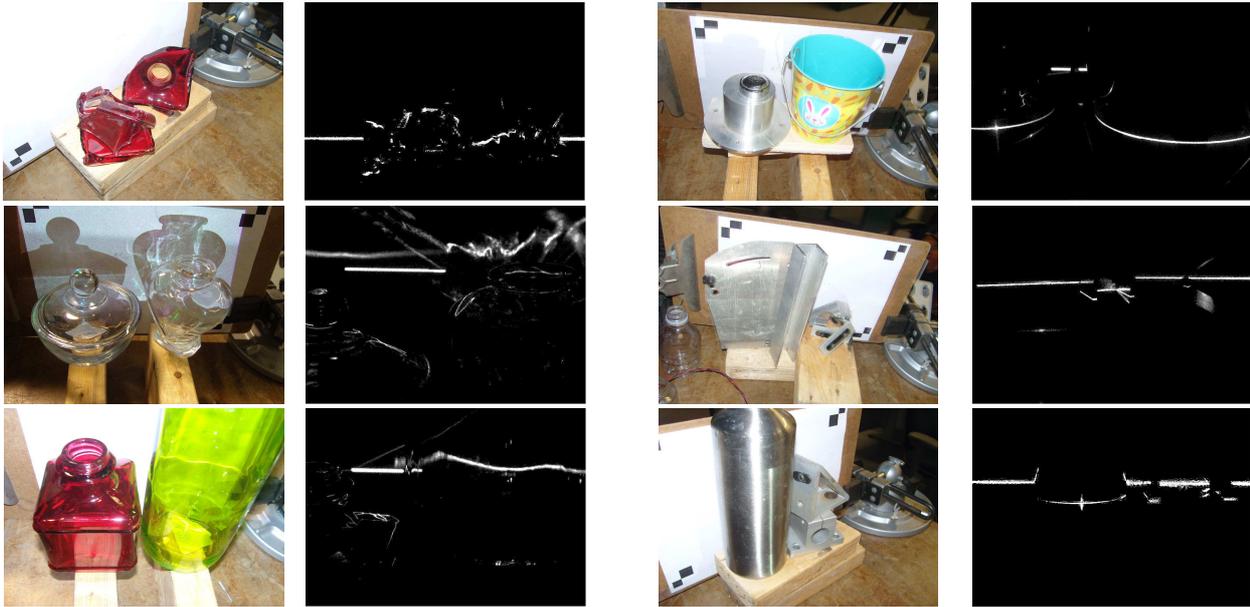


Wax materials show small support



Broken wax has anomalously large support

terials. Furthermore, while there has been previous work on using structured and controlled illumination to obtain the full light-transport matrix [25] (and therefore identify the scene material properties), we attempt to do the same with a single slice of the light transport function. In this sense, our sensor enables a low-power, low-cost version of dome-based material identification work [27].



(a) Glass (left) and metal (right) scenes with one frame of line-stripe video showing caustics

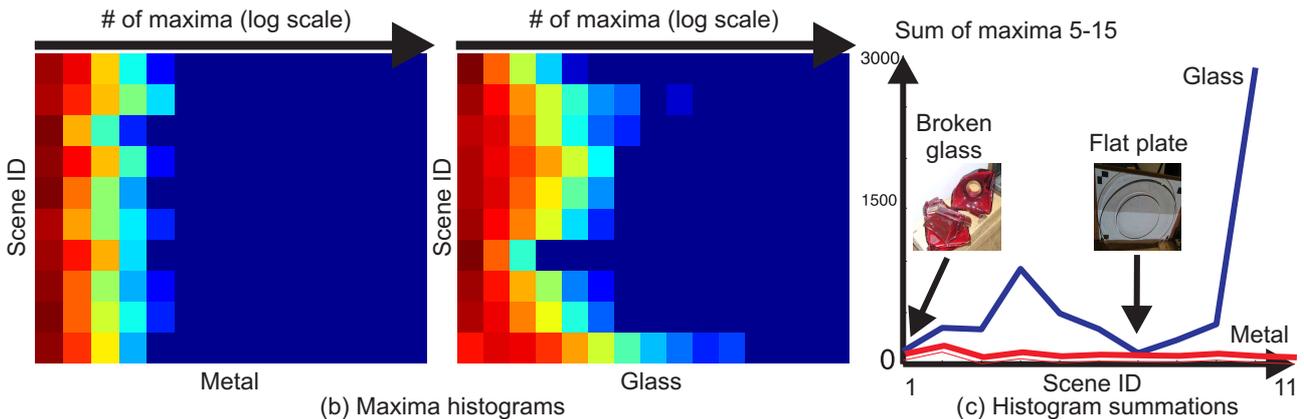


Figure 9. Histograms of the number of maxima at each scene point can distinguish between metal and glass scenes. At the top we show a few of the scenes that we investigated. At the bottom left and center we show maxima histograms for each scene, from 1 to 15 (log scale). Most of the metal maxima are concentrated to the left (most scene points have few maxima). Integrating the histograms from 5-15 gives a feature that clearly separated metals from glass (bottom right). The two failure cases are due to broken glass and the glass plate.

4.1. Fast per-frame analysis

Fig. 6 shows a per-frame analysis of a scene with milky water bottles and another with glass objects. Our method has five steps: (1) For each column, find the maximum intensity pixel (2) At this pixel, apply two filters (see figure inset), (3) If filter 1's response is greater than a threshold, it is glass (4) Otherwise, if the response to the second filter is greater than a second threshold, label as milk. If there is no labeling, then it is a diffuse material. In the figure, we have marked glass as red, milk as blue and diffuse as green. The biggest errors are for clear glass when the camera sees mostly the background. This is a fast classification, since for each column the filters are only applied once.

4.2. Full scan analysis

For a diffuse object, with no inter-reflections, each scene point is illuminated exactly once by the line scan. However, for complex scenes, the light stripe can be scattered, reflected and refracted such that scene points might be illuminated many times. We present heuristic features that exploit this temporal variation in a scene point's incident illumination. First we point out that simple detection of a single, sharp peak in a scene point's appearance profile [13] strongly suggests lambertian/diffuse reflectance. If the profile has no peak, then the projector is not illuminating this pixel and therefore it is in shadow. Fig. 7 shows two scenes where pixels are classified into shadow (blue), dif-

fuse (green) or other (red). If most of the scene is classified as “other”, we continue with further analysis.

Once we eliminate the possibility of diffuse material dominance, we can use a simple, heuristical feature which can potentially detect scenes with scattering media or sub-surface scattering surfaces. We have compared 7 scenes containing wax candles, diffuse styrofoam and milky water (two concentrations) with seven other scenes containing glass, metal, paper, and wood (Fig. 8). For each scan video, we compute the three dimensional Fourier transform, and integrate the power spectrum over the time frequencies. The resultant features are shown to the left of the scenes. This 2D feature’s support is small for the wax and milk scenes, but large for other scenes. In particular the y-axis has many non-zero components in the non-scattering scenes. We believe that thresholding the spectral support has the potential to indicate the dominance of scattering materials. We show a failure case of a broken candle, where the presence of non-smooth edges may add higher terms to the power spectrum.

With the possibility of diffuse or scattering material dominance removed from consideration, we now focus on distinguishing between reflective or refractive surfaces. We conducted line scan experiments on 11 metallic scenes and 10 glass scenes including colored metals and glass, broken glass, brushed and polished metal pieces. Fig. 9(a) shows three scenes from each class, with frames showing the projected line stripe. Note the different types of reflective and refractive caustics that are produced. We have empirically found that the number of intensity maxima in the appearance profile at each pixel can be very discriminative. An intuitive explanation is that since reflective caustics are caused by opaque objects, the number of observed caustics at each scene point is less than in a refractive material, where the camera can view through the material onto the diffuse background, allowing the observance of many more caustics.

In Fig. 9(b), we show two visualizations of the heuristical feature, one each for glass and metal. In the first image, each row corresponds to one of the 11 metal scenes we investigated. The rows are a histogram of the number of observed maxima in that scene’s scan video, ranging from 1 to 15 (in log-scale). The second image shows similar data for 10 glass scenes. Note that, the histograms for the metallic scene is more heavily concentrated in the first few entries. This means that metallic scenes have fewer maxima in their appearance profiles than glass scenes. We can exploit this by counting the number of scene points that exhibit maxima greater than a threshold (four in our case). Fig. 9(c) shows the plot of these scene points for the glass (blue) scenes and metallic scenes (red). Note that, except for two scenes, there is a clear difference between the metals and glass. These two anomalous scenes are easily explained; one is a glass plate with almost no caustics, and the other is the broken glass scene which breaks our smoothness assumption.

In addition to the intensity maxima at each pixel, we also empirically investigated four other features for disambiguating metals from glass in Fig. 10(a)-(e). In (a) and (b), we show two non-discriminative features; the histogram of edges and histogram of fourier frequencies. As before, rows correspond to the scene IDs and columns to the entries of the histogram. In (c) we show the raw features obtained from a low-res histogram of gradients (HOG). The top three discriminative features (d) for metal and glass show promise, but we believe more data is needed before a discriminative hyperplane can be learned. In Fig. 10 (e), we create an approximate “indirect” image, as in [19], by integrating each appearance profile after removing the intensity maxima. The images have reduced direct component, and visually appear discriminative. While simple edge features do not capture this difference, better features could be tried.

5. Conclusion: Towards mixed scenes

A combination of the online and offline features introduced above, should enable a fast “light-probe” for material detection. As a proof-of-concept (Fig. 11), we implemented a cascaded approach to classifying a scene with mixed materials. We first apply the fast per-frame analysis and then the diffuse/shadow/other classification to each pixel. Then, we apply the fourier and maxima histogram analysis to a 6x6 tiling of the scene. Note that without ambient light submission, there are many errors in the scene.

In conclusion, we have designed and built a structured light sensor that can work in ambient light, enabling reconstruction under direct sunlight. We also show that the removal of ambient light allows the measurement of visual effects, such as caustics, and can lead to classification of a scene’s dominant material property. In the future, we would like to exploit multiple viewpoints (similar to [22]) to show reconstruction of objects with complex materials.

Acknowledgements: 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. Narasimhan was partially supported by an ONR Grant N00014-11-1-0295, an NSF Award IIS-0964562 and a SAIT GRO grant. Koppal was supported by NSF award IIS-0926148, ONR award N000140911022, and the US ARL and the US Army Research Office under contract/grant number 54262-CI. The authors thank Todd Zickler for early discussions.

References

- [1] Microvision product page. In <http://www.microvision.com/technology/picop.html>.
- [2] Videos of kinect in sunlit outdoors. In <http://www.iheartrobotics.com/2011/02/playback-kinect-outdoors.html>.

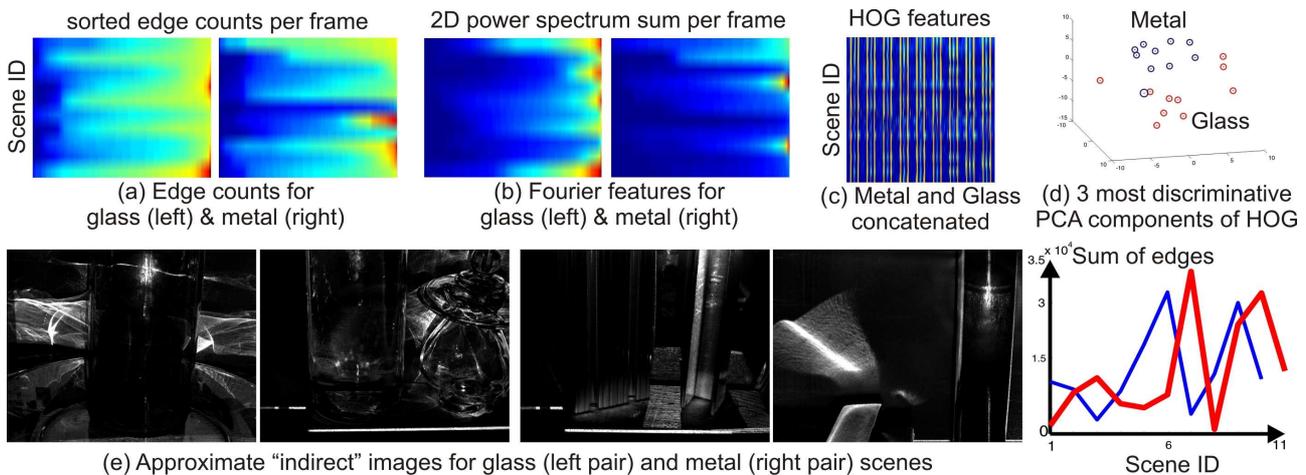


Figure 10. We show other features that were empirically investigated for glass/metal discrimination. (a) and (b) show poor features such as histograms for edges and power spectrum sums. (c) and (d) show better low-res HOG features on which PCA is applied and which has some discriminative power. We also created approximate indirect images [13] which visually appear to show potential for discrimination.

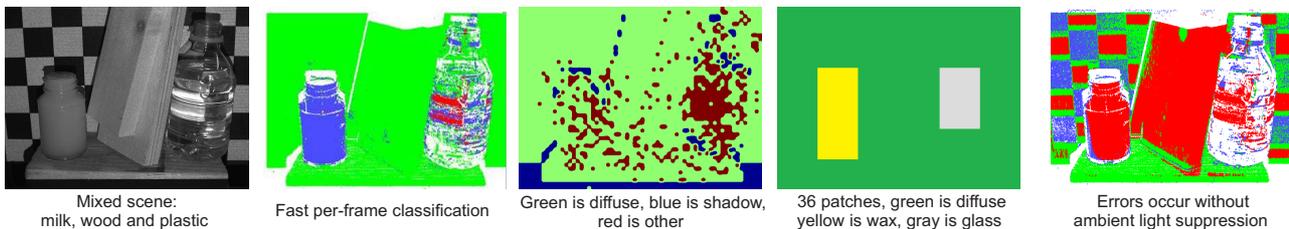


Figure 11. Analysis of a mixed scene

- [3] Wikipedia and article. In http://en.wikipedia.org/wiki/Handheld_projector.
- [4] Y. Chuang, D. Zongker, J. Hindorff, B. Curless, and D. Salesin. Environment matting extensions: Towards higher accuracy and real-time capture. *SIGGRAPH*, 2000.
- [5] D. Cotting, M. Naef, M. Gross, and H. Fuchs. Embedding imperceptible patterns into projected images for simultaneous acquisition and display. In *ISMAR*, 2004.
- [6] L. Cowan, N. Weibel, W. Griswold, L. Pina, and J. Hollan. Projector phone use: practices and social implications. In *Personal and ubiquitous computing*, 2012.
- [7] M. Fritz, M. Black, G. Bradski, and T. Darrell. An additive latent feature model for transparent object recognition. In *NIPS*, 2009.
- [8] A. Greaves, A. Hang, and E. Rukzio. Picture browsing and map interaction using a projector phone. In *MobileHCI*, 2008.
- [9] M. Gupta, Y. Tian, S. Narasimhan, and L. Zhang. A combined theory of defocused illumination and global light transport. *IJCV*, 2011.
- [10] A. Jones, A. Gardner, M. Bolas, I. McDowall, and P. Debevec. Performance geometry capture for spatially varying relighting. *ECVMP*.
- [11] A. Jones, I. McDowall, H. Yamada, M. Bolas, and P. Debevec. Rendering for an interactive 360 degree light field display. In *ACM SIGGRAPH*, 2007.
- [12] V. R. Kompella and P. Sturm. Detection and avoidance of semi-transparent obstacles using a collective-reward based approach. In *ICRA*, 2011.
- [13] S. J. Koppal and S. G. Narasimhan. Clustering appearance for scene analysis. In *CVPR*, 2006.
- [14] S. J. Koppal, S. Yamazaki, and S. Narasimhan. Exploiting dlp illumination dithering for reconstruction and photography of high-speed scenes. In *IJCV*, 2011.
- [15] I. McDowall and M. Bolas. Fast light for display, sensing and control applications. In *IEEE VR Workshop on Emerging Display Technologies*, 2005.
- [16] K. McHenry, J. Ponce, and D. Forsyth. Finding glass. In *CVPR*, 2005.
- [17] C. Mertz, J. Kozar, J. R. Miller, and C. Thorpe. Eye-safe laser line stripper for outside use. In *IVS*, June 2002.
- [18] N. J. Morris and K. N. Kutulakos. Reconstructing the surface of inhomogeneous transparent scenes by scatter trace photography. *ICCV*, 2007.
- [19] S. K. Nayar, V. Branzoi, and T. Boult. Programmable imaging using a digital micromirror array. *IEEE CVPR*, 2004.
- [20] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM SIGGRAPH*, 2006.
- [21] D. Padilla, P. Davidson, J. Carlson, and D. Novick. Advancements in sensing and perception using structured lighting techniques : An ldrd final report. In *Sandia National Lab report*, 2005.
- [22] J. Park and A. Kak. 3d modeling of optically challenging objects. In *IEEE TVCG*, 2008.
- [23] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future : A unified approach to image-based modeling and spatially immersive displays. In *ACM SIGGRAPH*, 1998.
- [24] J. Schoning, M. Lochtefeld, M. Rohs, and A. Kruger. Projector phones: A new class of interfaces for augmented reality. In *IJMHCI*, 2010.
- [25] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual photography. *ACM SIGGRAPH*, 2005.
- [26] D. Takhar, J. Laska, M. Wakin, M. Duarte, D. Baron, S. Sarvotham, K. Kelly, and R. Baraniuk. A new compressive imaging camera architecture using optical-domain compression. In *SPIE*, 2006.
- [27] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using brdf slices. In *CVPR*, 2009.
- [28] P. M. Will and K. S. Pennington. Grid coding: A preprocessing technique for robot and machine vision. *AI*, 2, 1971.