Towards Active Reflectometry for Segmentation

Alejandro Perez
MIT CSAIL
aperez@csail.mit.edu

Abstract

In this report we discuss the progress towards the development of a system that measures reflectance using an actuated light source, a light meter, and a camera. Existing reflectometry methods require complicated hardware or estimate reflectance using standard images. The process of empirically measuring reflectance is complicated and timely as it requires controlled illumination and the direct measurement of all light reflected around the hemisphere of the object in question. Similarly, estimating the reflectivity from a standard image is difficult as all light sources are often not completely known and reflective behavior, i.e., specularity and gloss, can cause saturation and loss of information. Moreover, strong assumptions regarding the material, shape, and surface are often required. We present an approach that is simple and can efficiently detect changes in reflectance across objects, even when color, depth, and all other information is similar. An actuated light source with a known pose is used to illuminate the different parts of a scene. Simultaneously, a light meter is used to measure the change in reflected light incurred by the new source. The resulting information can then be used to augment the segmentation process as different materials vary greatly in reflectance regardless of color and shape.

1. Introduction

The reflective properties of different surfaces and materials have been studied for centuries [11]. Indeed, the understanding of the behavior of visible light is used for photorealistic rendering [9], computer animation [17], material classification [5], object recognition [16], and shape estimation [6]. Many formalizations of the reflectance properties of different materials have been developed. In particular, the Bidirectional reflectance distribution function (BRDF) [12] has been widely used for decades [15], [13], [18]. BRDFs are four dimensional functions that determine how light is reflected at a given point of a surface. Because the function considers four or more variables, its measurement and specification is a complicated process. However, reflectance information has proven to be a powerful way of understanding and even simulating different materials and environments. For example, the BRDF is commonly empirically computed to validate and research different materials [2]. One device that is widely used for this process is the gonioreflectometer [1]. The instrument measures the BRDF directly by illuminating a target and measuring the light reflected around its hemisphere with a sensor or camera.

Figure 1. Inferring material properties from an image is a difficult task. In the image above, even when all surface and shape information is known, classifying the material as plastic, paper, water, or steel, is not a straightforward task. The reflectance properties however, vary regardless of color, shape, and surface similarities.

Unfortunately, despite the recent advancements in photographic hardware, there is no simple mechanism or sensor that can be used to quickly extract reflective information from a scene. More recently there has been an increased interest in estimating and learning these properties by considering different features such as those found in 'micro-textures' [16] and the use of HDR images with controlled illumination [14]. In this report, we explore different methods to measure reflectance. Our motivation is the possible application of a reflectometry technique to the segmentation of deformable, unknown objects, in scenes where color and depth information do not provide sufficient information.

2. On Perceiving Reflectivity

A human tasked with analyzing and classifying the object shown in Figure 1 will possibly move and change its viewpoint slightly until either gloss, specularity, or transparency is apparent. Intuitively speaking, the perception of the reaction a surface has to light plays a key role in day-

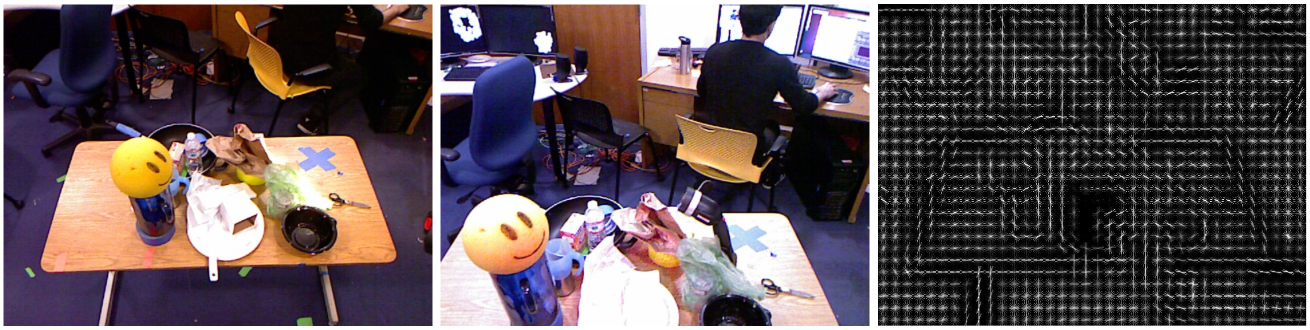

Figure 2. (a) A light source is added to the scene. Note the bright region above the dark plastic plate. (b) Red structured light is projected as the camera is moved. (c) HOG features were computed in real-time as the light source was moved.

to-day segmentation. This approach is adopted by humans at very early states. Indeed, infants as young as 20-weeks of age have been observed to use reflectance for color constancy [4]. In very general terms, the variations in surface reflectance among different materials is so sharp and apparent, that it can be useful even when information (and previous experience) is minimal. Consider the following tasks: investigating if a plastic bag left outside is wet; verifying if an image displayed on a building is a projection, monitor, graffiti, or a reflection; searching for a phone that fell into a pile of clothes. For all these tasks, colors, depth, or even knowledge of all shapes involved might not be enough. However, the slightest hint of gloss, specularity, or just abrupt change in color could possibly help a person in these situations. More importantly, this behavior can be leveraged without empirically computing the path of the light, the material properties of the surface, or even the shape of the object. In many cases, simply detecting the difference in perceived reflectivity is enough.

In an attempt to develop a technique that provides similar information, we implemented and tested several approaches. At the highest level, we considered three main paths: (i) movement, (ii) controlled illumination, and (iii) controlled reflectometry. As seen in Figure 2, our first arrangement included several objects with gloss, color changes, transparency, and specularity that were apparent or could be perceived by moving the camera slightly. Even though in some cases these provided salient information about the objects in the scene, we found this is only the case for some materials, and noticed that these were often present solely in small parts of the surface of a particular object. The next idea was to attempt to induce these type of 'events' by adding an actuated source of light into the scene. In our case, we tried both structured light (see Figure 2 b) and a robotic arm with an LED light source (see Figure 2 a). These produced seemingly noticeable changes in the environment. However, the red structured light, which was designed to extract shape information [10], usually affected all objects equally. In the case of the LED light, we found it to be too disruptive, as it would frequently saturate the image resulting in loss of information. One promising approach we explored was tracking changes in HOG (Histogram of Oriented Gradients) [3] descriptors as a light source was projected at different parts of the image (see Figure 2 c). We found that objects with similar material properties would have the HOG descriptors corresponding to its location change in similar ways. However, this observed behavior was useful but coarse. More specifically, we found that it could be used to differentiate paper and plastic on a white 'blob'. However, for more challenging scenarios, e.g., cardboard and a similarly colored paper bag, the approach would require exhaustively searching for the right parameters. For example, selecting the appropriate cell size for an object with a texture that is not known in advance can be very challenging. Because the goal of our project is generality and salience, we moved on to measuring what causes these changes instead of the changes themselves.

Measuring the amount of visible light reflected from the environment is the standard approach to photography and vision. Yet, any information obtained from a single viewpoint will usually contain light from various sources. It appears that infants [4] quickly develop the ability to extract reflectance information in order to achieve color constancy. Similarly, the measurement of reflected light seems to be a powerful tool in material analysis [2]. Moreover, our experience using actuated light was positive; the information is useful because it does not require a complete understanding of all the illumination in the scene and simply focuses on the changes caused by actions that are well understood. The next step was to develop a way to have this information available whenever needed and not only when reflectance became perceivable through controlled illumination. This would ensure maximum generality and reliability. Following a 'certum quod factum' mentality, we moved on to exploring the idea of directly measuring the reflected light.

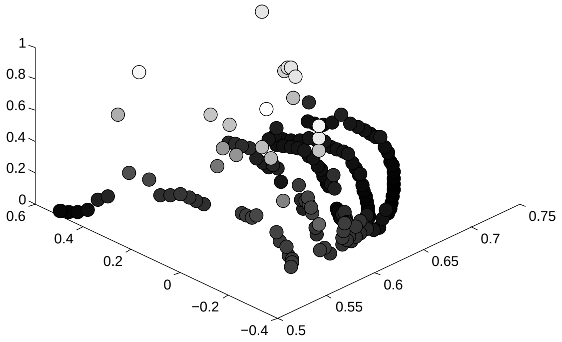

Figure 3. Illuminance readings obtained from the scene in Figure 6. The robot moves the light source and meter in a rectangular pattern. The Z axis represents the lux value. The X, Y axis the estimated pose of the sensor.

3. Proposed System

The goal of the proposed technique is to measure reflective information that can be used in conjunction with an image or depth-image to better segment the objects in the scene. We tested a hardware configuration that includes an LED light source that is controlled by a robotic arm. Therefore, the pose of the light source is known. Additionally, a light meter [8] is positioned immediately next to the light source. Like the light, the pose of the light meter is given by the robotic arm. The light meter returns a continuous reading of Lux, i.e., a measurement of the intensity of 'visible light' at its location. It has a maximum range of 400,000 Lux and returns 1.5 measurements per second. The detector itself is 115mm by 60mm. Therefore, when the robotic arm is orthogonal to a surface, a small beam of light is projected onto it and the sensor measures the amount of light reflected back. In very general terms, in order to measure all light reflected by a surface a sensor would have to move through the entire hemisphere around it. However, with this configuration, a salient measurement of some of the light reflected is obtained. Information regarding ambient light is unnecessary, as the measurement can be the difference in intensity caused by the controlled light source. Moreover, because the meter is in close proximity to both the light source and the surface, this resulting change serves a substantial sample of reflected visible light. Finally, a standard camera observes the target environment from a fixed or known location. This configuration was selected to circumvent saturation, self-occlusion, and the need of extensive measurement trajectories. The controlled light source will not result in saturation of the lit area as the light meter has a very wide range. The self-occlusion caused by the robotic arm is manageable as the pose of the arm is known and the fixed camera is able to see the scene when it moves out of view. Finally, because the arm is able to move close to the objects in the environment, a considerable amount of the reflected light is measured by the sensor without having to scan the hemisphere around the target or to calculate the path of the light.

4. Experiments

We tested our configuration on several configurations of deformable, cluttered, objects where segmentation is challenging even with strong color, normal, and depth information. The robot moved the controlled light source and the light meter in a 'lawn mower' pattern above the objects on the table. The end-effector remained orthogonal to the table at all times and at a constant height. The entire pose of the robot was logged continuously along with Lux readings, images from a standard camera, depth-image from a Kinect sensor, and time.

Our three main scenarios are depicted in (i) Figures 3 (b) and 7 (b), (ii) Figures 3 (c) and 6 (b), and (iii) Figure 7. The data obtained in these environments is presented below.

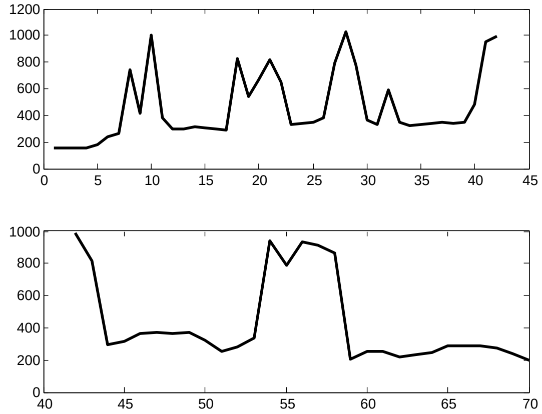

Figure 5. The lux values obtained during the first two 'swipes' shown in Figure 6. Peaks correspond to the two pieces of plastic wrapping, the CD jewel case, foam, and the plastic spatula. The second swipe displays the two plastic cups.

5. Discussion

The measurement of the change in luminance caused by a controlled light source proved to be salient and allowed for the detection of several objects. We were able to correspond the luminance readings to points in depth and pixel images. This is due largely to the fact that the camera and light meter pose information was readily available. Additionally,

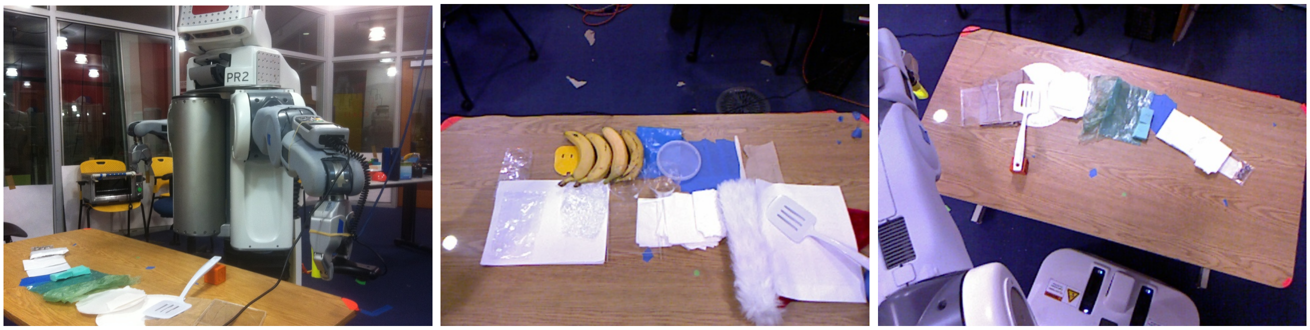

(a) (b) (c)

Figure 4. (a) The PR2 robot scanning a scene with a light source and light meter. (b) an arrangement of objects on a table (from left to right, bottom-up): different plastics on paper, a transparent CD jewel case, napkins, foam, a hat, a plastic spatula, empty plastic cup, plastic up with water, yellow paper, bananas, blue paper wrapping, a plastic lid, blue tape, a small knife, brown paper. (c) Another arrangement (from left to right): Transparent CD case, brown paper bag, plastic spatula, paper plate, paper, green plastic wrapping, green foam, blue tape, different kinds of paper, foam, a plastic bag with metallic nuts.

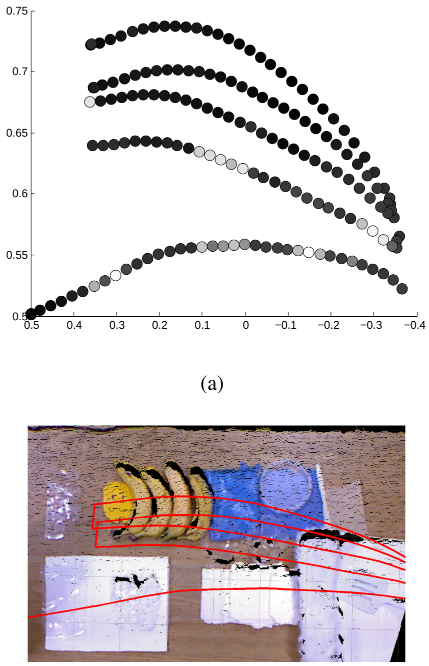

(a)

(b)

Figure 6. (Below) the trajectory traveled by the sensor while scanning the scene. (Above) Illuminance values (in Lux) obtained at their respective X,Y locations; darker colors indicate lower intensity.

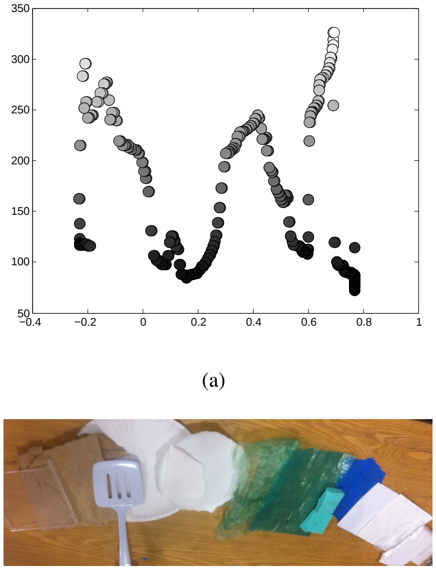

(a)

(b)

Figure 7. Resulting illuminance values for the scene depicted above. Values of high intensity correspond to the CD jewel case, plastic spatula, green plastic bags, and foam.

we were able to obtain useful readings from objects that are hard to perceive by cameras and depth-cameras (e.g., transparent plastic). However, the process of moving the robotic arm is time consuming and prone to error. Even though we are able to track the actual trajectory followed by the arm, specifying a 'lawn mower' pattern proved to be a difficult task. This is because the robot was programmed to maintain a constant orientation and height while scanning the table with the light and sensor. It is possible to write a joint-by-joint routine that allows the robot to move quickly and efficiently, but this would limit the types of environments that can be evaluated. It is apparent that this type of approach can greatly help with the segmentation of unknown objects in unstructured environments, i.e., scanning the sensor over a stack of papers to find a plastic binder. Yet, it is not clear if would be practical to use this approach to assign a value to all parts of an image as is done by a LIDAR or Kinect sensor. Perhaps measurements outside the visi-

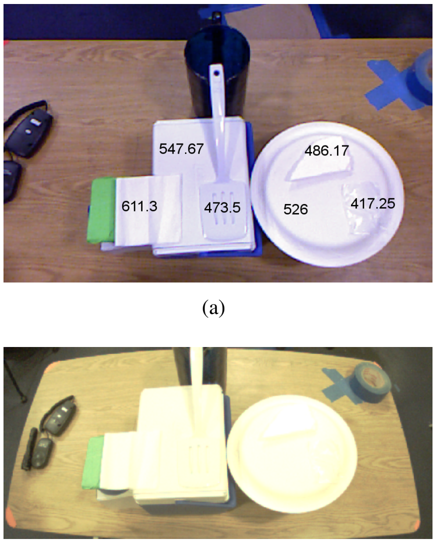

(a)

(b)

Figure 8. (Above) Several white objects were placed at equal height and measured with the device. The resulting difference in Lux obtained from adding and removing a controlled source of illumination were, Napkin: 611.3, Fridge: 547.67, Plastic spatula: 473.5, Paper plate: 526, Foam: 486.17, Plastic wrap: 417.25. (Below) An image taken from another camera on the robot.

ble light spectrum can be used to augment the segmentation process without the need of an actuated light and sensing end-effector. Cameras such as the FinePix S3 Pro [7] are able to provide IR and UV information that might directly related to texture, material properties, and reflectance.

References

[1] Frank Benford and WA Ruggles. Some characteristics of metal mirrors and a new gonioreflectometer. *JOSA*, 32(3):174–174, 1942.

[2] DB Betts, FJJ Clarke, LJ Cox, and JA Larkin. Infrared reflection properties of five types of black coating for radiometric detectors. *Journal of Physics E: Scientific Instruments*, 18(8):689, 1985.

[3] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005.

[4] James L Dannemiller. A test of color constancy in 9-and 20-week-old human infants following simulated illuminant changes. *Developmental Psychology*, 25(2):171, 1989.

[5] W Erb. Computer-controlled gonioreflectometer for the measurement of spectral reflection characteristics. *Applied optics*, 19(22):3789–3794, 1980.

[6] Roland W Fleming, Antonio Torralba, and Edward H Adelson. Specular reflections and the perception of shape. *Journal of Vision*, 4(9), 2004.

[7] E Raymond Hunt, W Dean Hively, Greg W McCarty, Craig ST Daughtry, Patrick J Forrestal, Robert J Kratochvil, James L Carr, Nathaniel F Allen, Joseph R Fox-Rabinovitz, and Christopher D Miller. Nir-green-blue high-resolution digital images for assessment of winter cover crop biomass. *GIScience & Remote Sensing*, 48(1):86–98, 2011.

[8] Loyd A Jones and HR Condit. The brightness scale of exterior scenes and the computation of correct photographic exposure. *JOSA*, 31(11):651–678, 1941.

[9] Konrad F Karner, Heinz Mayer, and Michael Gervautz. An image based measurement system for anisotropic reflection. In *Computer Graphics Forum*, volume 15, pages 119–128. Wiley Online Library, 1996.

[10] Kurt Konolige. Projected texture stereo. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 148–155. IEEE, 2010.

[11] J. H. Lambert. *Photometria sive de mensure de gratibus luminis, colorum umbrae*. Eberhard Klett, 1760.

[12] Fred E Nicodemus. Directional reflectance and emissivity of an opaque surface. *Applied Optics*, 4(7):767–773, 1965.

[13] Stephen M Pompea, Donald F Shepard, and Scot Anderson. Brdf measurements at 6328 angstroms and 10.6 micrometers of optical black surfaces for space telescopes. In *32nd Annual Technical Symposium*, pages 236–247. International Society for Optics and Photonics, 1989.

[14] Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler. Passive reflectometry. In *Proceedings of the 10th European Conference on Computer Vision: Part IV*, ECCV '08, pages 859–872, Berlin, Heidelberg, 2008. Springer-Verlag.

[15] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994.

[16] Lavanya Sharan, Ce Liu, Ruth Rosenholtz, and Edward H Adelson. Recognizing materials using perceptually inspired features. *International Journal of Computer Vision*, pages 1–24, 2013.

[17] Eric Tabellion and Arnauld Lamorlette. An approximate global illumination system for computer generated films. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 469–476. ACM, 2004.

[18] K Tomiyasu. Relationship between and measurement of differential scattering coefficient and bidirectional reflectance distribution function (brdf). *Geoscience and Remote Sensing, IEEE Transactions on*, 26(5):660–665, 1988.