(CD-ROM is included for the Movie)

A Non-Photorealistic Camera: Detecting Silhouettes with Multi-flash

Ramesh Raskar[*]
Mitsubishi Electric Research Labs

Jingyi Yu[†]
MIT

Adrian Ilie[‡]
U of North Carolina at Chapel Hill

Abstract

We describe a non-photorealistic camera ('NPR camera') to robustly capture edge information in real scenes. New stylized images are synthesized by highlighting or de-emphasizing the detected edges.

1 Introduction

Stylization of photographs is traditionally considered a post-process [DeCarlo et al 2002] but the task can be greatly simplified by modifying the capture setup. We show how a multi-flash camera (Figure 1, top row) can robustly classify the intensity edges into depth edges and texture (or material) edges. Depth edges correspond to the pixels with depth discontinuity and are loosely referred to as silhouettes (internal or external) in NPR literature.

Conventional passive or active computer vision methods are not robust enough to capture or classify scene edges. For example, intensity edge detection cannot distinguish between edges due to shape and due to material changes [Forsyth and Ponce 2002]. Active methods including photometric stereo (PS) and shape from shadow [Savarese et al 2001] work well for smooth surfaces but fail at depth discontinuities. More robust Helmholtz stereo [Zickler et al 2002] requires solving a correspondence problem. All such active vision methods, although they produce more information like surface orientation, require elaborate setup. So far, they have not been shown to work on complex scenes with high depth complexity and low contrast like the flower plant. Our approach uses a surprisingly simple setup but achieves high-quality edge-based stylization effects. Further, the NPR camera can be readily built as a self-contained device. It does not require external components.

2 Exposition

We take successive pictures of a scene, each with a different light source placed *very close* to the camera's center of projection (COP). Due to the small baseline between the camera COP and the light source, a narrow sliver of shadow appears attached to each depth edge. Since the sliver of shadow is several pixel wide, we need to determine which segment of the shadow region boundary corresponds to a depth edge. Note that the shadows appear on the 'opposite' side of the light i.e. if the light is to the left of the camera; the shadows appear to the right of the depth edge. Hence, in this image, the left part of shadow boundary is a depth edge. The shadow position changes in each image, and by combining information from multiple images, we detect all the depth edge pixels. (The complete algorithm is explained in the video on the CD-ROM and following pages).

In our NPR camera prototype, the baseline is about 50 mm and we can detect all pixels with depth discontinuity of 5mm or more at 2000mm. The technique is robust as shown in Fig 1 but it does fail for highly specular objects and requires a background within a finite distance to be able to detect the depth edge. We assume a moderately controlled illumination environment.

After detecting depth edges, other edges in the intensity image are classified by a process of elimination. In the Fig 1 second row, the shape of the 'bone' is easier to understand when the image is superimposed with depth edges. In the fourth row, depth edges are highlighted with bold black lines while texture edges (leaf interiors) are de-emphasized by diffusion blurring.

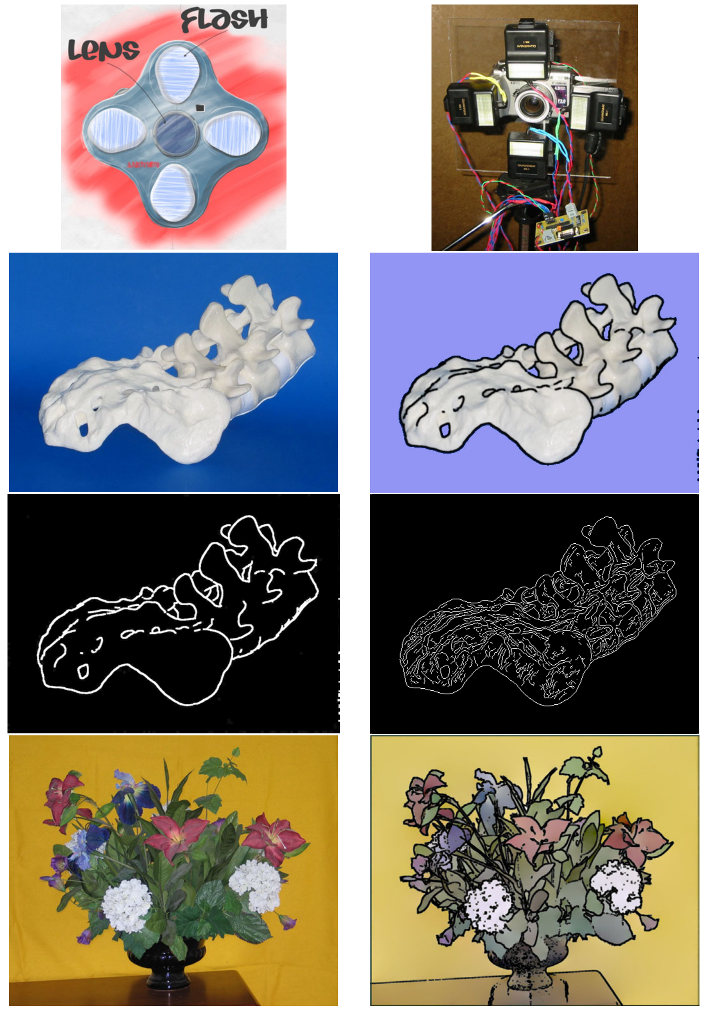

Figure 1. (Top row) The four-flash NPR camera prototype (Second) Source image, enhanced with depth edges (Third) Depth edges detected by NPR camera, compares with Canny edge detector (Fourth) Flower image, highlighted depth edges with de-emphasized texture.

References

DECARLO, D., AND SANTELLA, A. 2002. Stylization and Abstraction of Photographs. In *Proc. Siggraph 02, ACM Press*.

FORSYTH, AND PONCE. 2002. *Computer Vision, A Modern Approach*.

SAVARESE, S., RUSHMEIER, H., BERNARDINI, F., AND PERONA, P. 2001. Shadow Carving. In *Proc. of the Int. Conf. on Computer Vision*.

ZICKLER, T., BELHUMEUR, P. N., AND J., K. D. 2002. Helmholtz Stereopsis: Exploiting Reciprocity for Surface Reconstruction . In *ECCV*.

---

[*] raskar@merl.com, [†] jingyi@lcs.mit.edu, [‡] adyilie@cs.unc.edu

(CD-ROM is included for the Movie)

Additional Supporting Images

(Not part of the first page of the sketch)

Source

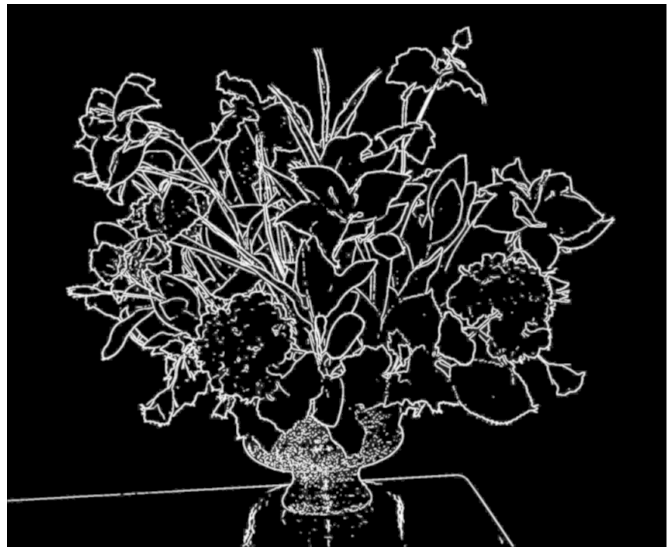

Depth discontinuity Edges using NPR Camera

Intensity edges using Canny Edge detector. Notice how a pure intensity image operation leads to spurious edges.

Example of a NPR rendering, Depth edges are superimposed to highlight the shape and texture edges de-emphasized to reduce the detail in the interior.

The basic algorithm to detect silhouette is as follows.

Capture n pictures I_k, $k = 1..n$ with a point light source at P_k
For all pixels x, $I_{max}(x) = \max_k(I_k(x))$, k = 1..n
For each image k,
    For all pixels x, $R_k(x) = I_k(x)/I_{max}(x)$
    Traverse in the image R_k radially outwards from epipole e_k
    Find pixels y with step edges with negative transition
    Mark the pixel y as silhouette

The image I_{max} describes a no-shadow image.

Pixel e_k in camera image is the image of the (idealized point) light source at P_k in the camera image. The notion is similar to the epipole of one camera in another camera's image.

(CD-ROM is included for the Movie)

How our method works

Image due to top flash

Image due to left flash

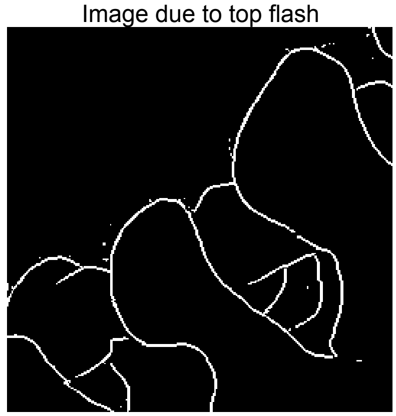

Resultant depth edges
detected by NPR camera

Image due to right flash

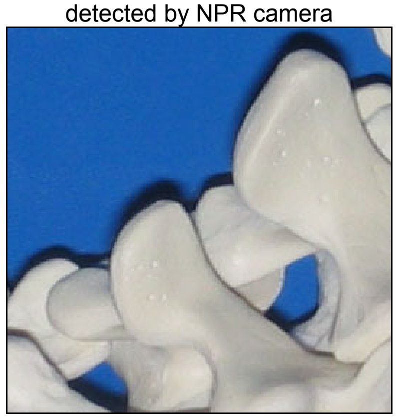

Image due to bottom flash

Note how the shadow attached to the depth edges change with the flash position.

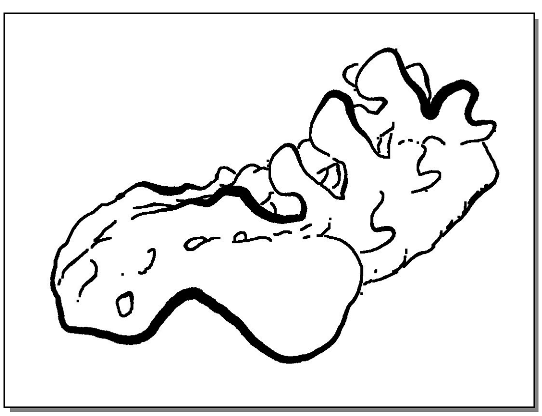

A stylized output with variable silhouette width