Casual 3D Photography

Askold V. Strat[1,3] Manuel M. Oliveira[2]
strat@symbol.com oliveira@inf.ufrgs.br
[1]Symbol Technologies, Inc. [2]Instituto de Informática – UFRGS [3]SUNY at Stony Brook

Figure 1. Picture of a dummy head showing projected laser stripes (left). A photograph of the dummy's face used for texturing (center). Reconstructed texture-mapped model (right).

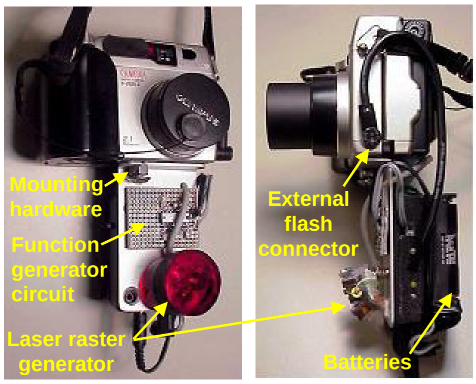

Figure 2. The 3D camera prototype. A laser pattern generator is attached to a consumer-grade camera and electrically connected to it through the external flash connector.

Abstract
We describe the construction of a lightweight, point-and-shoot 3D color camera, intended for reconstruction of smooth surfaces. Our approach extracts 3D shape from a single image and is completely automatic. The 3D camera consists of an eye-safe laser pattern generator attached to a consumer-grade digital camera and some image processing software.

Categories and Subject Descriptors: I.4.8 [Image Processing and Computer Vision]: Digitization and Image Capture – Imaging Geometry; B.4.2 [Input/Output and Data Communications]: Input/Output Devices.

Additional Keywords: 3D camera, 3D photography, 3D shape acquisition.

Description
We have designed and built a truly portable 3D color camera suitable for shape acquisition of smooth surfaces. This is our second-generation 3D camera prototype, and consists of a consumer grade digital camera, an eye-safe laser pattern generator and some image processing software.

The 3D camera acquires a pair of pictures with a single press of the camera's trigger. One of the pictures contains a set of projected laser lines (created by the pattern generator) and is used to reconstruct the object's imaged geometry using triangulation. The other picture is used for texture mapping. Model reconstruction is carried out in three steps: first, the laser lines are extracted and each such line is assigned an elevation angle. This information is then used to compute depth for each of the lines' pixels. Next, the 3D coordinates of the recovered points are processed to remove outliers, minimize the effects of possibly incorrectly assigned angles, and fill in holes in the model. After 3D geometry has been recovered, the second picture is used for texture mapping.

Figures 1 (left) and (center) show a pair of images acquired by our 3D camera. Figure 1 (left) is a photograph of a dummy head (about the size of a human head) exhibiting the laser lines. Figure 1 (center) shows the image acquired for texturing. Figure 1 (right) shows the reconstructed model seen from a novel viewpoint. The acquisition of the pair of pictures takes less than one second. Figure 2 shows a prototype of our 3D camera based on an Olympus C-2020 zoom (top portion of Figure 2). In order to keep the design as modular as possible, the only electrical connection between the camera and the raster generator is through the external flash connector, whose signal is used for synchronization. Such a design provides great flexibility, allowing any consumer-grade digital camera with support for external flash and exposure time control to be quickly turned into a 3D camera.

The laser raster generator produces an eye-safe laser (Class IIa at 650nm), requiring about 2.7V at 80mA. It projects a set of quasi-horizontal scanning lines while a narrow laser beam bounces onto two mirrors, which oscillate around perpendicular axes. The laser beam bounces off a mirror with a vertical axis (x-mirror), oscillating at 2KHz, and then encounters the second mirror (y-mirror), which oscillates at 6Hz. This motion composition results in a zigzagged line. A raster-control circuit is used to turn off the laser every other line, eliminating the retrace.

Like all other laser-based techniques, our approach is subject to several physical limitations due to the optical material properties, ambient illumination, occlusions, and speckle noise, which make the task of shape extraction significantly harder. A limiting factor in the amount of geometric detail that can be captured by our 3D camera is the density of the raster lines used to sample the surface. Currently, the raster generator projects sixty lines. While a larger number could potentially allow the recovery of more surface detail, reducing the spacing between two consecutive lines makes the extraction and disambiguation task more difficulty. The existence of depth discontinuities can make the proper segmentation of lines considerably difficult. An important step towards making our approach more robust is to use a global analysis for correctly handling the projection of raster lines at surface discontinuities.