Document photography in vitro

G. Nagy[1], B. Clifford[1], A. Berg[1], G. Saunders[1], E. Barney Smith[2], D. Lopresti[3]
[1]Rensselaer Polytechnic Institute, [2]Boise State University, [3]Lehigh University,
nagy@ecse.rpi.edu

Abstract

Experiments on evaluating and improving the mechanical configuration and target illumination in a prototype camera-based ballot counter are presented. The constraints on the mechanical design are gravity paper feed, portability, ease of use and low cost. The constraints on illumination are dictated by image processing requirements. Initial results are reported on the effects of transparent ballot cover plates, geometry, and light conditioning (color, diffusion, polarization).

1. Introduction

Photographing a document behind glass is surprisingly troublesome. Why then photograph a ballot behind glass? As explained in a companion ICDAR paper [1], a camera based ballot reader offers potential advantages over the customary scanner based systems. It is desirable to complement such a system with a simple gravity-feed paper transport that does not require a motor and is less prone to jamming. Vacuum platforms to keep the ballot flat in the near vertical paper path are too expensive, complex, and power hungry for our application. Therefore we use a transparent cover plate to flatten the ballot in case it has been previously folded, crimpled or wetted.

Sheet-fed optical scan systems are regarded as an appealing option by voting advocates who tout their support for manual recounts. While several commercial vendors sell such systems for use in elections, recent experience has raised intrinsic reliability concerns. During the 2008 U.S. Presidential election, wet ballots were found to jam scanners in Virginia (it had been raining that day and voters' hands were wet) [2]. Dust build-up on poorly maintained machines was blamed for uncounted ballots in Michigan [3]. A survey conducted by the U.S. Election Assistance Commission after the November 2004 general election found 541 instances of scanner failures in the 210 jurisdictions that used optical scan systems, where "failure" in this case was defined as "A malfunction or interruption of a paper ballot reading device that either renders the device incapable of counting votes or renders the tabulated results inaccurate" [4]. Indeed, scanner failures were the most prevalent type of reported error across all voting systems and technologies used in that election. Although such instances are still relatively rare, they are non-negligible and could trigger a transition from scanner based to camera based ballot counters.

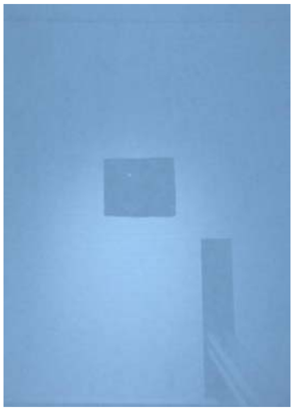

Figure. 1. Ghost image (reflection) of the camera and camera support in the transparent cover plate. The target behind the cover plate is matt white paper.

Our design comprises a thin vertical chute with an opaque plate behind the ballot and a transparent cover plate in front of it. Before the ballot is photographed, it must enter the chute far enough to make it difficult or impossible for the voter to retrieve it. (Some electoral jurisdictions allow two-sided ballots, which would require two transparent cover plates and two cameras.) It is the reflections from the cover plate (Fig. 1) that hamper uniform illumination and motivated this study.

The above-cited ICDAR paper references objections to direct recording electronic (DRE) ballot counters, lists the requirements of viable paper-based election technology, proposes a camera-based Portable Ballot Counter (PBC), and compares characteristics of scanners and cameras that are relevant to image digitization and vote mark extraction. The major difference between contact digitization (scanner) and remote digitization (camera) is the variability in the mapping of the reflectance of the target to gray values. The range of gray levels observed on a contact-scanned uniformly colored document is only about 10, and there is no problem with glare. Furthermore, at ~260 dpi the diameter of the point spread function of the scanner is just over half of that of the camera.

In this paper, we present the current status of the design and construction of a prototype ballot counter and our on-going efforts to achieve near-uniform illumination of the ballot. In the following sections we describe the evolution of the prototype, discuss imaging problems that arise in this application, and present our observations of the relative merits of various configurations of illumination. We seek the help of the camera-based document recognition community towards making further progress. For researchers interested in the technical issues behind the controversies surrounding recent elections in the United States, we reference two lively monographs on the subject [5, 6].

Prototype design and construction

The project's goal of increased reliability and accuracy requires a rigid mechanical foundation for the ballot holder, camera and lights (Fig 2). For the frame and the ballot feed, sturdy materials were chosen for accurate positioning of the ballot with respect to the camera. Machined aluminum parts contribute to ease of modification and assembly.

The design for the ballot feed mechanism must accept a wide range of paper weights and of ballot heights and widths. The camera is mounted such that the largest allowable ballot occupies the full field of view of the camera. Both of these parameters must be easily adjustable. In the prototype, the entire feed tray can be tilted and moved forward and backward along rails. The paper guide rails can be slid in and out to accommodate the ballot size used in a particular jurisdiction (Fig. 3).

The design for the ballot feed mechanism must accept a wide range of paper weights and of ballot heights and widths. The camera is mounted such that the largest allowable ballot occupies the full field of view of the camera. Both of these parameters must be easily adjustable. In the prototype the entire feed tray can be tilted and moved forward and backward along rails. The paper guide rails can be slid in and out to accommodate the ballot size used in a particular jurisdiction (Fig. 3).

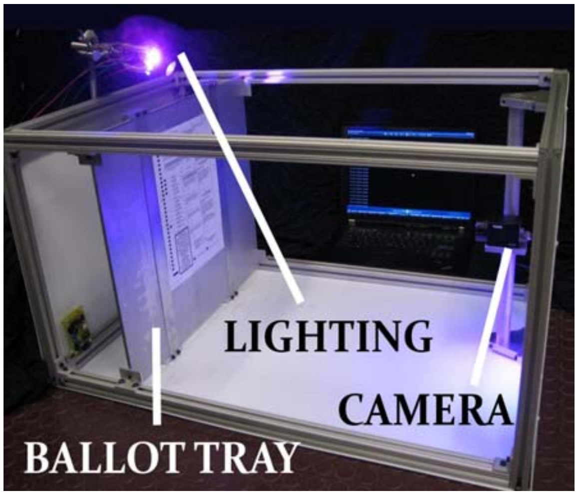

Figure 2. Diffuse light from behind the ballot is reflected from the sides, top, bottom and back of the partially shown white enclosure. The micro-processor behind the ballot tray will replace many functions of the laptop in the background.

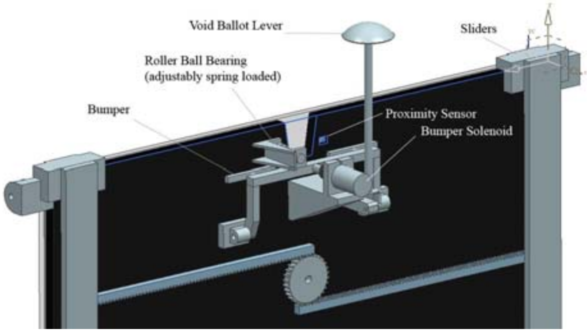

Figure 3. Drawing of back of the ballot feed tray with the rail guide adjustment mechanism and the cancellation punch (under construction).

The mechanical system controls the ballot's travel through the chute by means of a microprocessor that gathers data from optical sensors. The computer determines where the paper is located in the tray and triggers the camera and the lights when the system is ready. The mark extraction subsystem [7] will analyze the image and display the result for each race on a screen at the top of the box. During the development phase, the image will be transferred to a laptop for analysis and display of the results.

Should the voter notice an unintended overvote, undervote, or an error in mark interpretation, the voter can cancel the ballot and request a fresh ballot from election officials. Cancelled ballots are marked with a die that cuts a small circle out of paper when the Void Ballot Lever is pressed. Pressing the lever also triggers an input to the controller that is used to signal that the ballot has been cancelled and therefore should not be tallied. The physical marking of the ballot with a punched hole provides a visual signal that indicates to election officials which ballots should and should not be counted in the event of a recount. In an operational system, it may be necessary to route invalid ballots to a separate secure box below the ballot counter.

The shell of the enclosure is made from an opaque white polypropylene sheet designed for easy removal (not shown in Fig. 2). The shell provides mechanical protection, controls light infiltration into the imaging space and protects voter confidentiality by blocking cast and voided ballots from the view of nearby voters. The mechanical design of the prototype is primarily motivated by the need for flexibility that permits us to rapidly reconfigure lighting, imaging, ballot feeding and other subsystems throughout the course of the development. For this reason, the prototype is constructed using lightweight aluminum extrusions with multiple full-length mounting channels available from a number of manufacturers. The mounting channels allow us to reconfigure the lighting, sensors, switches, and other components inside the prototype enclosure simply and quickly. In contrast, we anticipate that a low-cost manufactured solution would be designed using fixed mechanical elements with adjustment features only where required.

3. Imaging considerations

As mentioned, the transparent cover plate that keeps the ballot flat in a near-vertical position gives rise to a novel twist in document photography. The production of ballot images suitable for mark extraction imposes the following constraints on the enclosure, the camera, and the source of illumination.

- The illumination must be uniform to preserve the contrast between the background of the ballot, the preprinted text and rulings, and marks recorded in arbitrary positions with an uncontrolled variety of writing instruments
- Glare (highlights) due to the geometry of the light sources, camera and cover plate must be avoided.
- The superimposition of a reflected image of the camera or of any part of the enclosure on the recorded ballot image must be prevented.

We have designed a systematic series of experiments to evaluate methods of illumination that have been suggested so far. We are now comparing various possible solutions. This is work in progress, and the results reported here are fragmentary and inconclusive. Several of the experiments require the acquisition or construction of components that will take place within the next month or two.

The evaluation of the illumination schemes is based on the camera images alone. The camera that we selected, based on its imaging properties, software library support, and moderate price (~$400), is a 15 megapixel Canon G10 Powershot. For mark extraction we use 127 µm (~200 dpi) images, but for imaging the slowly-varying illumination field 656 µm (~39 dpi) is sufficient. Although ballots may be as large as 400mm in either direction, for evaluating alternative illumination methods we settled on 261mm x 185mm, 380 x 270 pixel, 8-bit gray-scale images.

Our goal is to illuminate the target so as to obtain an array of nearly constant pixel values when we take a photograph of a white sheet of paper. We set the constant between gray values of 100 and 200 to avoid either overexposing or underexposing any part of the target. Setting the average gray value to between 100 and 200 is much easier with the camera controls than by controlling the current or voltage to the light source. The target and camera are fixed, the depth of field is shallow, and there are plenty of photons, so we typically use fairly low ISO-equivalent light sensitivity settings (80 to 200), long exposure (1/8 to 1 sec), and a wide lens aperture (f/3.2): For each photo, the camera records

FN	File_name (e.g. E1H1L2O2)
RN	“Resolution” (e.g. 200 dpi),
IM	Image Size (e.g. 640 x 480)
ET	Exposure_time (e.g. 100ms)
IS	ISO_number (e.g. 200)
AP	Aperture (e.g. f3.2)
CP	Compression (e.g. JPG_5)
FS	Focus Setting (e.g. 610mm)

Image formation is determined by the geometry of the optical axis of the camera, the wave-fronts of the illumination, and the plane of the target and cover plate. Photometric considerations include (a) the transmission coefficient and the bidirectional reflectance function (BRDF) of the cover plate, (b) the spectrum, coherence and polarization of the light, (c) the light collection properties of the camera lens and CCD sensors, and (d) the linearity and signal-to-noise ratio of the read-out, amplifier and analog-to-digital converter electronics.

We explore the imaging and illumination space in terms of the (1) color of the illumination, (2) diffuser, (3) enclosure, (4) cover plate,

(5) luminaire, and (6) polarizer. We list below the key to the options we intend to explore. The selections that we have not yet tried in some combination are italicized in the following list of experimental conditions.

Key to image designations in Table I. *All of images taken with blank (white) 8.5x11" paper used as the target/ballot. Categories in italics have yet to be tested.*

Color

C1	No filter
C2	*Red filter*
C3	Blue filter

Diffusion

D1	No diffuser
D2	White polypropylene mounted on lights
D3	White elastene/spandex mounted on lights
D4	*Optical diffuser mounted on lights*

Enclosure

E1	No enclosure
E2	Opaque box with white sides, top, & back
E3	Opaque Box with white sides, top, & black back
E4	*Box with tissue paper sides, white top, & back*
E5	Box with tissue paper sides, white top, & black back
E6	*Box with tissue paper sides, top, & white back*
E7	*Box with tissue paper sides, top, & black back*
E8	*CAT enclosure*
E9	*Elastene side panels*

Cover Plate

H1	None
H2	Acrylic thermoplastic (Plexiglas, Lucite, Perspex)
H3	Non-reflective acrylic
H4	*Non-reflective glass*
H5	*Coated museum glass (Mirogard)*

Luminary:

L1	Fluorescent ceiling lights
L2	30-LED light with reflector
L3	Pair of 3-LED fixtures
L4	Fluorescent circle light without reflector
L5	60W incandescent white bulb with reflector
L6	*High-frequency fluorescent ring light*
L7	*Grazing light*
L8	*Edge light / light guide*

Orientation

O0	Uncontrolled
O1	Direct illumination from behind the camera
O2	Direct, from above and behind the camera
O3	Direct, from sides
O4	Indirect, from above and behind the ballot
O5	Indirect, from the sides

Polarizer

P1	No polarizer
P2	Horizontal polarizer
P3	Vertical polarizer
P4	Horizontal and vertical polarizers at 90deg.
P5	*High-extinction linear glass polarizer*

Version

Vx	x is the version number

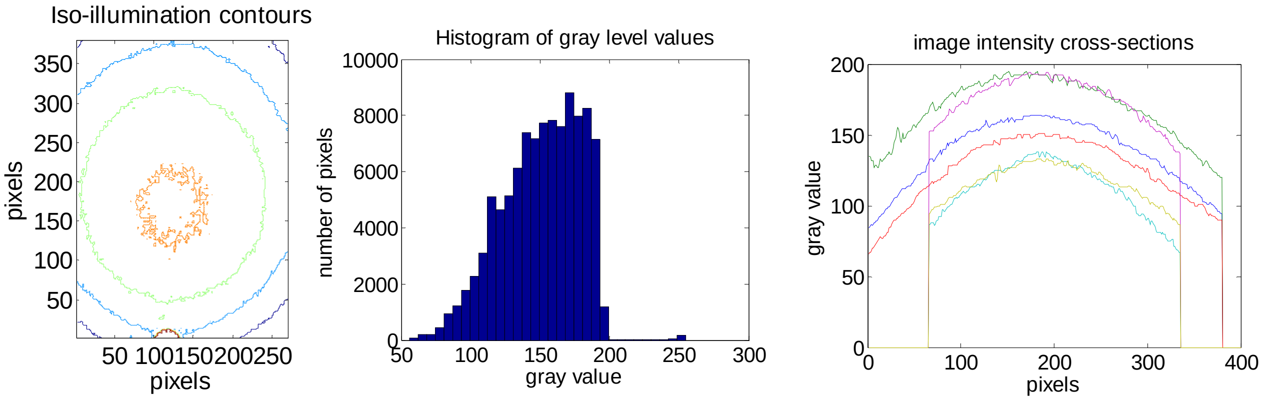

(a) (b) (c)

Figure 4. Matlab displays for observing the uniformity of illumination: (a) Five equally spaced illumination contours; (b) Thirty-two bin histogram of gray values. (c) Three horizontal and three vertical cross-sections of 380 x 270 pixel image. This example demonstrates the extent of non-uniformity if insufficient precautions are taken to prevent glare from the cover plate (Image #9 in Table 1, FM = 89).

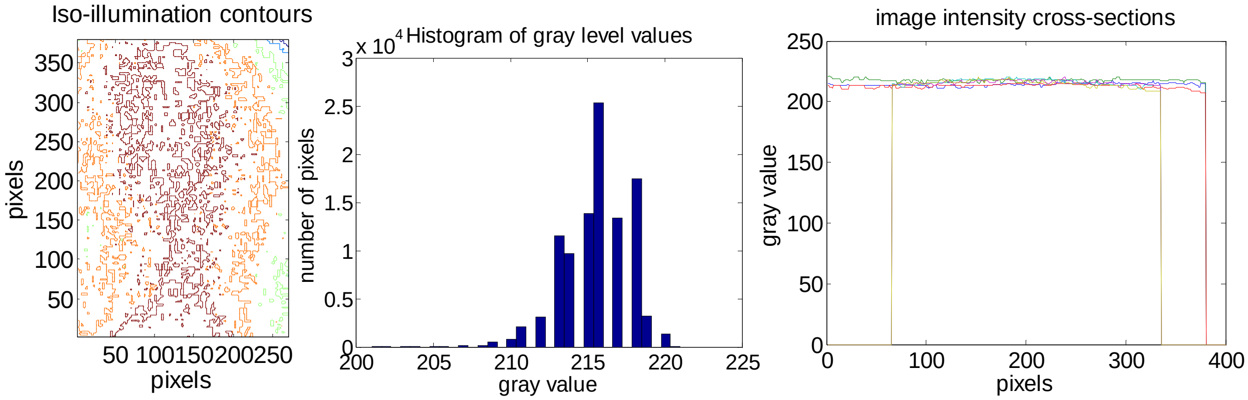

Figure 5. Contours, histogram, and cross-sections of the illumination for the most uniform lighting conditions that we have found so far (without the cover plate!). Some of the scales are different than in Fig. 4. Here the contours are only two gray levels apart. (Image #16 in Table 1, FM = 6).

4. Illumination

In order to maintain uniform spatial sampling and thereby avoid noise-prone and time-consuming resampling of the image, we align the optical axis of the camera perpendicular to the target plane. This leaves the location and orientation of the light sources as free geometric parameters.

We record the illumination field by photographing in color a matt white sheet of paper and analyzing the resulting image in Matlab. For most illumination conditions, we reduce the image to gray scale. We plot the contours of iso-illumination, examine the intensity histogram of the entire image or of parts of it, and trace horizontal, vertical and diagonal cross sections of the image at the center and near the edges (Fig. 4). We have recorded so far images under about 50 different lighting conditions.

We will eventually post our detailed evaluation on the PERFECT website [8], but here we provide only a simple Figure of Merit for the conditions we tested:

> *FM = number of gray levels between the 5% and the 95% quantiles of the measured reflectance.*

We normalize FM to NFM by dividing it by the average gray level of the image. Table I shows the Figure of Merit for various conditions. For the non-uniform conditions illustrated in Fig. 4, the 95 percentile is at gray level 189 and the 5 percentile is at 100, therefore FM = 89. Fig. 5 shows that creating uniform illumination is much easier without the cover plate. In the next paragraphs we describe the various factors that affect ballot illumination and the results of experiments conducted to date.

Reflection. Glare or highlights are caused by specular (mirror) reflection. Common glass or Plexiglas (H2 images) with a coefficient of refraction of ~1.5 and a critical angle of total reflection of 41˚ reflects about 8% of white light at 0˚(4% from each of its front and back surfaces). Non-glare acrylic (H3) or glass (H4) sold at framing shops reflects half as much. Museum glass reflects ~1% of the light, but generally requires additional support. Ultra-high transmission optical coated glass plates with less than 0.5% reflection are available only in small sizes and at high cost.

Polarization. The optical properties of uncoated paper depend on scattering inside the material. Ballots may be coated or uncoated, and even uncoated paper may acquire a certain gloss or sheen. Although uncoated paper is usually modeled as a Lambertian reflector, up to 5% of the light may undergo specular reflection at 45˚. Light reflected from a glass plate is polarized parallel to the plate and can be reduced by a polarizer. Some inexpensive sunglass lenses are surprisingly effective polarizers, with a coefficient of extinction of about 50:1 (therefore looking at a white light source through two lenses at right angles reduces the light by a factor of nearly 50 over that transmitted by both lenses oriented the same way). Optical linear glass polarizers have an extinction ratio of 10,000:1. Our experiments indicate, however, that with diffuse illumination very little light strikes the cover plate at large angles from the normal, so there is only minimal polarization (P1-P2-P3 images in Table I)

Diffusion. Diffuse light helps to eliminate both direct reflection and highlights. Portrait photographers often use white parasols. Small-objects are usually photographed (for instance for advertisements) inside a light box with matte or translucent white sides illuminated from within or from the outside. We experimented with both configurations.

A white box with light directed at the white wall behind the camera (E2) yields images of the bare target with near-uniform intensity. With the glass plate in place, however, the image contains a distinct picture of the black camera (Fig. 1). Even if the body of the camera is painted white or entirely concealed behind the white partition, the lens opening itself, which cannot be hidden, produces a dark spot on the image. This can be reduced by a black back wall (E3 and E5), or with diffuse illumination.

Large optics-quality diffusers are expensive. However, the inexpensive fabric spandex or elastene (a long-chain synthetic polymer fiber introduced by DuPont in 1958 under the trade name *Lycra*) has excellent light diffusion properties. In our application, even white paper impregnated with oil is sufficient to eliminate any sharp variation in light intensity, but Spandex is far more durable.

Luminaires. For the light source itself, LEDs that can be powered all day by a few dry cells are desirable. Even if 125V AC were available, incandescent lamps produce too much heat, and fluorescent lights flicker at the power frequency. High-frequency (~25kHz) linear and spherical illuminators are commonly used for diffuse fluorescent lighting in industrial computer vision applications, but they too require 125V AC.

LED light sources are typically housed in a metallic or acrylic reflector to increase light intensity. They act more like spot lights than flood lights. We therefore place a diffuser in front of every light fixture. Even with a diffuser, a single fixture cannot be placed far enough from the ballot to illuminate the entire area uniformly. We therefore experiment with various configurations of symmetrically placed clusters of three white LEDs, both inside the enclosure and outside the enclosure. In either case, the light is directed at the side walls, which diffuse the light further before reaching the target. This arrangement is inefficient in terms of the fraction of the total amount of light used for imaging, but the low duty cycle of ballot counters (<1000 images per day) and the sensitivity of the camera keep power consumption at negligible levels.

We are now placing orders for edge lights (light guides) and grazing lights. *Light guides* (Fig. 6a) are transparent, light diffusing acrylic sheets containing colorless diffusing particles. They accept light through their edges and redirect it to either or both surfaces. Strips of LED *grazing lights* (Fig. 6b) for surface illumination are available with a 10˚ beam perpendicular to the surface and a 60˚ beam parallel to the surface. We look forward to experimenting with edge illumination and grazing lights because they would provide a mechanically simpler and lighter configuration.

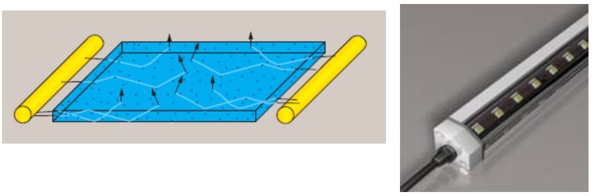
(a) (b)

Figure 6. (a) Edge illumination with light guide; (b) Architectural grazing lights consisting of a strip of LEDs.

5. Discussion

The goal of the experiments reported here was to find spatially smoothly varying and temporally constant illumination of the target by easily reproduced and relatively inexpensive means. Smooth spatial variation is important because during an election, high gradients in the illumination could move from pixel to pixel (CCD to CCD) under the influence of small physical changes due to vibration, stress or ambient temperature variation. These changes cannot be discriminated from image features. Temporal constancy is necessary to allow comparison of marked images to blank ballots photographed under identical conditions. Ballot counting devices must be inexpensive enough for large-scale purchase by cash-starved municipalities, and robust enough to be dragged out of storage and set up overnight once or twice a year by non-technical personnel.

Several of the configurations investigated keep the variability of the target to less than 15%, which corresponds to FM = 30 at average gray level 200 (Table I). We would like to reduce this to about 5% variability, well below the 25% recommended for exacting proofing applications (ISO 3664:2000). If the mark extraction experiments reveal that this is still too much, the uniformity can be increased by storing multiplicative and additive normalization factors for each pixel, as demonstrated in [1].

The best result that we have obtained so far with a cover plate and inside an enclosure is E5H3L3O5 (Fig.7). Here FM = 22 at an average gray level of about 165, external illumination with two LED light fixtures through

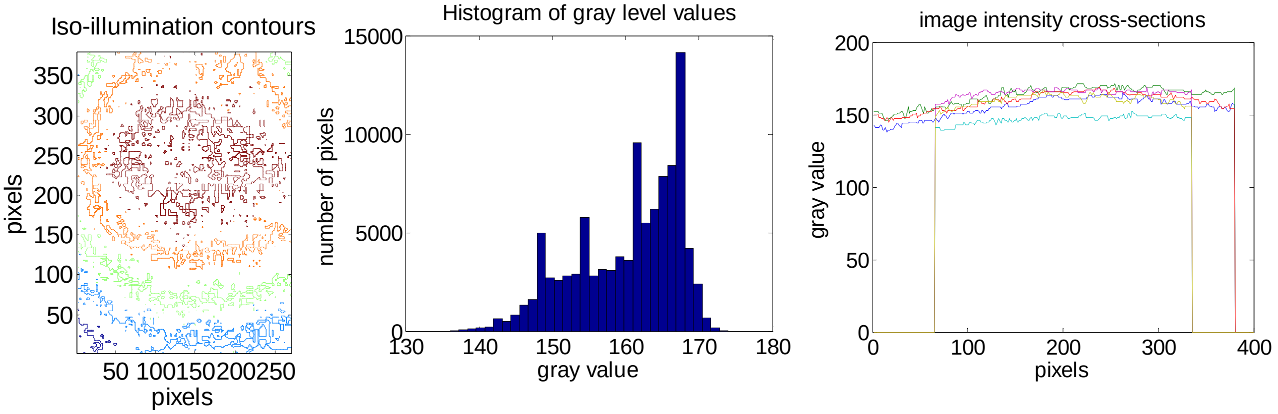

Figure 7. Contours, histogram, and cross-sections of the illumination for the most uniform lighting conditions that we have found so far *with* the cover plate (Image #31). The contours are 5 gray levels apart, FM = 22. The contours show that additional LEDs are necessary to increase the uniformity.

Table I. Figures of Merit (FM) & Normalized Figures of Merit (NFM) for various types of illumination

#	IMAGE (key in Section 3)	FM	NFM	#	IMAGE (key in Section 3)	FM	NFM
1	E1H1L1O0	11	0.06	24	E2H2L2O4P4*	3	0.99
2	E1H1L2O2	87	0.46	25	E2H3L2O4	40	0.24
3	E1H1L3O3	130	0.77	26	E3H1L2O4	43	0.27
4	E1H1L4O1	20	0.10	27	E3H2L2O4	32	0.22
5	E1H1L5O2	19	0.09	28	E3H3L2O4	32	0.22
6	E1H2L1O0	13	0.07	29	E5H1L3O5	25	0.15
7	E1H2L2O1	60	0.29	30	E5H2L3O5	23	0.14
8	E1H2L2O2	89	0.48	31	E5H3L3O5	22	0.14
9	E1H2L2O2P2	89	0.59	32	D2E1H1L2O2	15	0.11
10	E1H2L2O2P3	89	0.59	33	D2E1H1L3O3	39	0.25
11	E1H2L3O3	131	0.81	34	D2E1H2L2O2	16	0.12
12	E1H2L5O2	20	0.10	35	D2E1H2L2O2P2	17	0.09
13	E1H3L2O2	90	0.49	36	D2E1H2L2O2P3	18	0.10
14	E1H3L2O2P2	92	0.56	37	D2E1H2L3O3	41	0.28
15	E1H3L2O2P3	91	0.55	38	D2E1H3L2O2	17	0.12
16	E2H1L2O4V1	6	0.03	39	D2E1H3L2O2P2	19	0.10
17	E2H1L2O4V2	10	0.06	40	D2E1H3L2O2P3	18	0.10
18	E2H2L2O4V1	41	0.25	41	D3E1H1L2O2	36	0.24
19	E2H2L2O4V2	37	0.21	42	D3E1H1L3O3	32	0.24
20	E2H2L2O4P2V1	34	0.24	43	D3E1H1L5O2	17	0.11
21	E2H2L2O4P2V2	30	0.23	44	D3E1H2L2O2	34	0.25
22	E2H2L2O4P3V1	34	0.24	45	D3E1H2L3O3	35	0.28
23	E2H2L2O4P3V2	31	0.24	46	D3E1H2L5O2	22	0.13
				47	C3E2H2L2O4	99	0.73

Note: For simplicity, if the color, diffusion, or polarization classifications are “C1,” “D1,” or “P1”, then the category label is omitted from the image name. Similarly, if there is only one version of an image, the version category label is omitted
*Very dark image because of crossed polarizers

tissue paper sides, black back, and anti-reflective acrylic cover plate. The contours of Fig. 7 suggest that we can increase the uniformity with additional lights near the bottom. We would, however, much prefer to avoid having to construct a box-within-the-box. (Table I shows some images with even lower FM and NFM, but detailed examination indicates high spots too small to affect the figure of merit but large enough to cover a mark.) The imminent completion of the prototype shown in Fig. 2 will accelerate and improve our experiments, which have so far been conducted with ad hoc fixturing.

It may seem that we are going overboard in our attempts to design sound illumination for the relatively simple task of counting filled ovals, X's or checkmarks in or near known positions on a fixed form. Note, however, that according to the Voluntary Voting System Guidelines of 2005, "the system shall achieve a target error rate of no more than one in 10,000,000 ballot positions." No matter how this guideline is interpreted, or whether in fact error rates of 0.000001% can be verified at all, we must strive to minimize all avoidable sources of miscounts.

6. Acknowledgments

This material is based upon work supported by the National Science Foundation under grants Nos. CNS-0716647 (EBS), CNS-0716393 (GN), CNS-0716368 (DL), the Rensselaer Open Source Software Foundation, and the NYS/RPI Center for Automation Technologies and Systems (CATS). We gratefully acknowledge the suggestions of Architectural Lighting designer Naomi Miller, Professor Jean Paul Freyssinier of the Rensselaer Lighting Research Center, and Professor Ishwara Bhat of the Rensselaer Engineering Research Center for Smart Lighting.

7. References

[1] Nagy, G., Clifford, B., Berg, A., Saunders, G., Lopresti, D., Barney Smith, E.: Camera-based Ballot Counter. To appear in: Proceedings, ICDAR 2009.

[2] UPI http://www.upi.com/Top_News/2008/11/04/Rainy_weather_yields_wet_ballots_in_Va/ UP-22481225837607/ accessed January 11, 2008.

[3] "ES&S Voting Machines in Michigan Flunk Tests, Don't Tally Votes Consistently," K. Zetter, Wired blog. http://blog.wired.com/27bstroke6/2008/11/michigan-electi.html.

[4] "Chapter 11: Voting Equipment Malfunctions," 2004 Election Day Survey Report, Election Data Services, September 27, 2005. http://www.eac.gov/clearinghouse/docs/eds2004/eds-2004-part-2-chapter-11/attachment_download/file

[5] Hernson, P.S., Niemi, R.G., Hammer, M.J., Bederson, B.B., Conrad, F.C., Traugott, M.W.: Voting Technology – The Not-So-Simple Act of Casting a Ballot. Bookings Institution Press, Washington, D.C. 2008.

[6] Lausen, M.: Design for Democracy, Ballot + Election design, University of Chicago Press, 2008.

[7] E.H. Barney Smith, G. Nagy, D. Lopresti, Mark Detection from Scanned Ballots, Procs. Document Recognition and Retrieval, San Jose, SPIE, Jan 2009.

[8] PERFECT (Paper and Electronic Records for Elections Cultivating Trust), *htt://perfect.cse.lehigh.edu/,* accessed April 20, 2009.