

Adaptive Multispectral Illumination for Retinal Microsurgery

Raphael Sznitman, Diego Rother, Jim Handa, Peter Gehlbach,
Gregory D. Hager, and Russell Taylor

Johns Hopkins University
{sznitman,jthanda,pgelbach,hager,rht}@jhu.edu

Abstract. It has been shown that excessive white light exposure during retinal microsurgery can induce retinal damage. To address this problem, one can illuminate the retina with a device that alternates between white, and less damaging limited-spectrum light. The surgeon is then presented with a fully colored video by recoloring the limited-spectrum light frames, using information from the white-light frames. To obtain accurately colored images, while reducing phototoxicity, we have developed a novel algorithm that monitors the quality of the recolored images and determines when white light may be substituted by limited-spectrum light. We show qualitatively and quantitatively that our system can provide reliable images using a significantly smaller light dose as compared to other state-of-the-art coloring schemes.

1 Introduction

Retinal microsurgery is one of the most demanding types of surgery. The difficulty stems from the micron scale dimensions of tissue planes and blood vessels in the eye; the delicate nature of the neurosensory retina and the poor recovery of retinal function following significant injury. For many retinal surgeons the operative time is lengthy, increasing the risk of light-induced toxicity to the retina that can limit visual success despite achievement of all surgical objectives. Therefore, a central issue for the surgeon is balancing the need for adequate illumination of retinal structures with the risk of iatrogenic phototoxicity.

One of the principal sources for retinal light toxicity stems from the use of fiberoptic endoilluminators. Endoillumination is the primary means of directly illuminating surgical targets during posterior segment intraocular procedures and the primary source of retinal light toxicity during retinal surgery. Retinal phototoxicity from an operating microscope was first reported in 1983 in patients who had undergone cataract surgery with intraocular lens implantation [1] and has then been recognized as a potential complication of any intraocular surgical procedure. The frequency of these complications is reported to occur from 7% to 28% of patients undergoing cataract surgery [2,3]. Blue wavelength and ultraviolet light induce the greatest degree of retinal injury. In fact, in [4,5] it was found that commercially available light sources for endoillumination exceeded the International Commission on Non-Ionizing Radiation Protection guidelines

for retinal damage by visible light. In vitrectomy for macular hole repair, up to 7% of the patients have been reported to have experienced visually significant phototoxicity [6,7,8].

In [9], this issue was addressed by presenting a visualization system that can be used to significantly reduce the emission of highly toxic wavelengths compared to current practices. This was done by periodically illuminating the retina with white light and less damaging non-white light using a computer controlled rapidly switching multiple LED light source. A consequence of this is that images acquired are either fully colored (which we will refer to as white-light images) or monochromatic (or red-light images). To avoid visually straining the surgeon, monochromatic images were then recolored by using colors from white-light images.

In general, computer colorization schemes have existed since the 70's [10], for applications such as recoloring movies, and have since been further developed [11,12]. In general, however, such systems rely on a user to pre-select regions of the image that correspond to specific colors, making them ill-suited for this application. More recently, a time series analysis was proposed to model the retinal image scene [13]. This method however relies on having all visual cues (*e.g.* color and texture) available at all times to maintain an accurate retina model. In [9], a coloring scheme (ASR) was specifically developed for this setting. This algorithm estimates the different forms of motion which appear in the scene, taking this information into account to recolor the monochromatic images. An important limitation of ASR is that white illumination is used at regular intervals, regardless of the degree of change in the scene. This results in the use of excessive white light when no significant changes have happened in the scene, or conversely, poor recolorization (and image artifacts) when much has changed in the scene.

In order to further reduce phototoxicity and yet provide accurately colored images to the surgeon, we present a novel image acquisition and recoloring scheme for this setting. At each time frame, our algorithm determines what type of illumination should be used based on the estimated phototoxicity levels and the recolored image quality. Our method is simple and requires little parameter tuning, making it easy to use and flexible for surgical applications. We have experimentally shown that our method provides a quantitative improvement in coloring accuracy over state-of-the-art methods.

The remainder of this article is organized as follows: in section 2 our framework and colorization algorithm is presented; section 3 compares our novel approach with other methods; and finally, in section 4, we discuss some of the pitfalls of our system and future works.

2 Multispectral Illumination and Adaptive Color Fusion

In order to provide the surgeon with accurately colored images when using the light source from [9], we present an algorithm that *dynamically* chooses which illumination type to use at each time step, depending on estimates of the rendered image quality and phototoxicity levels induced. That is, the quality of the

recolorization and phototoxicity levels are continuously monitored, allowing us to estimate when it is appropriate to use white light illumination. In general, this occurs when the scene changes cannot be adequately “predicted” with the current available information.

The system we use consists of a device (as in [9]) capable of illuminating the retina using either white light, or less phototoxic red light. We define the sequence of images provided by the system as $I = \{I_1, \dots, I_N\}$ for N discrete time steps. Each image I_t is associated with a particular illumination L_t , where $L_t = 1$ means that white light was used at time t , and $L_t = 0$ means that red light was used. Consequently, when $L_t = 1$ all three color channels are available, $I_t = \{I_t^R, I_t^G, I_t^B\}$, whereas when $L_t = 0$ only the red channel I_t^R is available. We define the illumination history as $\mathbf{L}_t = \{L_1, \dots, L_t\}$. As in [9], the overall rate at which white light is flashed can be defined as $\varphi(\mathbf{L}_t) = \frac{\sum_{i=1}^t L_i}{t}$. We denote by F_t the final fully colored image rendered by our algorithm. To recolor the monochromatic images we maintain a color model of the scene for each time t , $M_t = \{M_t^R, M_t^G, M_t^B\}$.

Our goal then is to choose which illumination type, L_{t+1} , to use for the next time step. To do this, our criterion is to maximize a quantitative estimate of the patient’s wellbeing. This criterion combines the two costs incurred by the patient at time t : the *surgeon impairment cost* and the *phototoxicity cost*. The surgeon impairment cost, $S(\epsilon_t)$, is the cost of being accidentally harmed by the surgeon because of the error levels present in the recolored images, ϵ_t . The phototoxicity cost, $T(\mathbf{L}_t)$, is given by the damage to the patient produced by the illumination. In the next section we describe these costs in more detail. In section 2.2 we show how these costs are combined to select which illumination type to use at each time step.

2.1 Modeling the Cost Functions

As described in the previous section, there are two different costs incurred by the patient at time t during the procedure. The first cost is the *surgeon impairment cost*, $S(\epsilon_t)$. This is the cost (for the patient) of being accidentally harmed by the surgeon at time t . Clearly this risk (and hence the cost) increases as the recolorization error, ϵ_t (defined below) increases, since the surgeon is relying on poorer images to perform his job. The exact relationship between this cost and the error is unknown and depends, among many things, on the particular surgeon using the system. However we expect $S(\epsilon)$ to be an increasing function that levels off at a certain error, ϵ^* , at which stage the quality of the image is so poor that further deterioration does not result in additional risk. In practice, we will make sure that the system remains in the linear part of S , far from the critical value ϵ^* , where the surgeon is critically impaired. Based results from [9], we will model this relationship with the following function,

$$S(\epsilon) = \begin{cases} 1 & \text{if } \epsilon > \epsilon^* \\ \frac{\epsilon}{\epsilon^*} & \text{otherwise} \end{cases} \quad (1)$$

The recolorization error, ϵ , is due to the fact that the color model at any given time is not perfect, since the background scene changes due to the manipulations performed by the surgeon. In order to compute this error, we note that errors are only committed in the green and blue channels, since the red channel is observed at all times. We assume that the error committed in the green and blue channels at time t , $\epsilon_t^{G,B}$, is approximately equal to the error that would be obtained in the red channel, ϵ_t^R , if it were treated as the green and blue channels ($\epsilon_t^{G,B} \approx \epsilon_t^R$). Since the red channel is available at all times irrespective of the illumination type, ϵ_t^R can be directly computed as,

$$\epsilon_t^R = \|M_t^R - M_{t_w}^R\|_2 \tag{2}$$

where t_w is the last time step in which $L_{t_w} = 1$. Assuming further that the error does not change significantly in one time step, we approximate the error at time $t + 1$ by the error at time t , hence $\hat{\epsilon}_{t+1}^{G,B} \approx \hat{\epsilon}_t^{G,B} \approx \epsilon_t^R$.

The second cost, the *phototoxicity cost*, $T(\mathbf{L}_t)$, is the estimated damage at time t suffered by the patient because of the illumination used up to this point in time \mathbf{L}_t . It seems reasonable from the current literature [14] to relate the amount of phototoxic damage, T , to a function of the recent light exposure $\varphi(\mathbf{L}_t)$, where $\varphi(\mathbf{L}_t)$ is a function that models how the illumination history \mathbf{L}_t affects a cell at time t . We chose to define $\varphi(\mathbf{L}_t)$ as an exponential loss (approximated from [14]). That is, as time goes on, the influence of the past decreases exponentially fast. Hence, we approximate the phototoxicity cost by,

$$T(\mathbf{L}_t) = \begin{cases} 1 & \text{if } \varphi(\mathbf{L}_t) > L^* \\ e^{-\frac{(\varphi(\mathbf{L}_t) - L^*)^2}{2}} & \text{otherwise} \end{cases} \tag{3}$$

where L^* is some level of illumination at which irreversible damage to the patient (cell death) is produced.

It must be noted that while the choice of these functions is based on reasonable assumptions, these functions ultimately need to be empirically determined.

2.2 Choosing the Next Illumination Type

We can then formally define the estimated total cost for the patient at time $t + 1$ as the sum of the two costs described in the previous section,

$$E(\mathbf{L}_{t+1}, \hat{\epsilon}_{t+1}) = (1 - \lambda)S(\hat{\epsilon}_{t+1}) + \lambda T(\mathbf{L}_{t+1}). \tag{4}$$

where, $\hat{\epsilon}_{t+1}$, is the measure of the recolorization error defined in Eq. 2, \mathbf{L}_{t+1} is the history of illuminations at time $(t + 1)$ and λ is a tuning parameter which can be adjusted by the user (*i.e.* surgeon) to specify a bias for either image quality or phototoxicity. We select the next illumination type, by minimizing the patient wellbeing cost,

$$L_{t+1} = \arg \min_L E(\mathbf{L}_{t+1}, \hat{\epsilon}_{t+1}) = \arg \min_L \{(1 - \lambda)S(\hat{\epsilon}_{t+1}) + \lambda T(\mathbf{L}_{t+1})\}. \tag{5}$$

Notice that L can take only two values (0 or 1). Hence, this optimization reduces to

$$(1 - \lambda)S(\hat{\epsilon}_{t+1}) + \lambda(T([L_t; 1]) - T([L_t; 0])) \geq 0 \tag{6}$$

Since $\hat{\epsilon}_{t+1} = 0$ when $L_{t+1} = 1$, and $\hat{\epsilon}_{t+1} \approx \hat{\epsilon}_t$ when $L_{t+1} = 0$, all the quantities in Eq. 6 are known and choosing the next illumination type simply reduces to determining whether or not equation Eq. 6 is true.

2.3 Adaptive Active Scene Rendering

We now present the outline of our algorithm: Adaptive Active Scene Rendering (AASR). Fig. 1 provides a visual outline of AASR and associated images, respectively.

First for each image I_t , (Fig. 1 (Left)(a) and (Right)(a)) we detect and segment the tool in the image by using a 3D tool model (see [9] for more details). This provides us with a mask region for the tool, T_t (Fig. 1(Left)(b)). Then, in order to compute the new color model: if $L_t = 1$, M_t is computed by keeping pixel regions of M_{t-1} which appear where the tool is located and using I_t for regions where the tool is not present (Fig. 1(Right)(c)). This is done by using T_t to mask regions of the tool and allows for regions displaying the retina to be updated, keeping tool regions unchanged (similar to the work in [13]). If $L_t = 0$, then $M_t = M_{t-1}$ (Fig. 1(Left)(c)). Rendering the recolored image, F_t , is then done by combining $M_t^{G,B}$ and I_t^R on regions outside the tool, and using a tool color model to fill in the tool ((Fig. 1(Right)(b) and (Left)(d)). Having computed these, we can then estimate the error, ϵ_t , using M_t^R and $M_{t_w}^R$, as described by Eq. 2 (Fig. 1(Right)(d)), and choosing the following illumination type can be computed as in Eq. 6 (Fig. 1(Right)(e)).

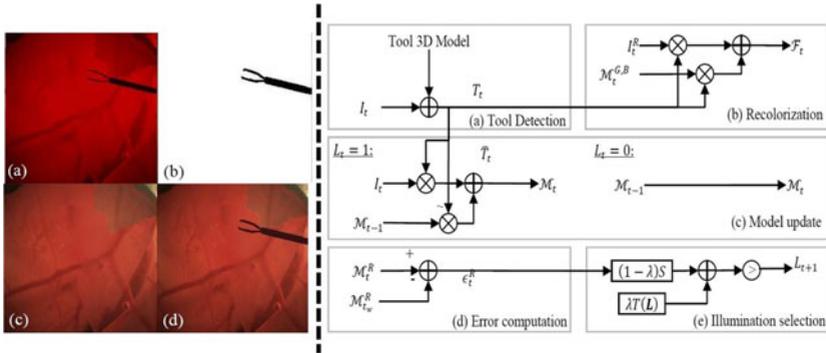


Fig. 1. (Left) Intermediate steps of the AASR algorithm: (a) non-white image provided by the device when $L_t = 0$, (b) tool segmentation, (c) representation of M_t and (d) recolored image by AASR. (Right) Block diagram of the proposed system. See section 2.3 for details.

3 Experiments

We now show how our system performs on image sequences from phantoms and from chorioallatonic chicken embryos. First, a quantitative comparison of AASR and a state-of-the-art method is presented, where it is shown that AASR surpasses ASR in a setting where ground truth is known. This is shown by measuring both image recoloring quality and quantity of white light used. We then show qualitative results of our algorithm on image sequences.

To validate the approach described in section 2, we recorded 5 image sequences of membrane peelings on phantom eyes using only white light. Each sequence consists of approximately 300 frames in similar visual settings. Doing so allows us to synthetically generate limited-spectrum images at any given time, by using only the red channel of white light images. This provides us with a way to quantitatively compare AASR and ASR, as ground truth is available.

For each image sequence we then ran AASR with three different settings: $\lambda = \{0.25, 0.5, 0.75\}$. This allows us to see results for cases where the surgeon applies a bias towards image quality, phototoxic levels, or no bias at all. For each image sequence, we also generated 4 recolored sequences using ASR, with different values of $\phi = \{1/2, 1/4, 1/8, 1/16\}$. As in [9], the L_2 (or mean squared error) norm is chosen to measure the error between the ground truth and the rendered images. In order to estimate phototoxicity levels, we observe the proportion of white-light images used.

In Fig. 2(*Left*) we show the results of this experiment by plotting the average recolorization error against the average estimated phototoxicity level. The dotted line (4 vertices; 1 for each value of ϕ) shows how ASR performs while the full line describes the performance of AASR (3 vertices; 1 for each value of λ). In general, we can notice that both methods displays a trade-off in accuracy: reducing one type of error induces the other and vice versa. We can also see that the AASR

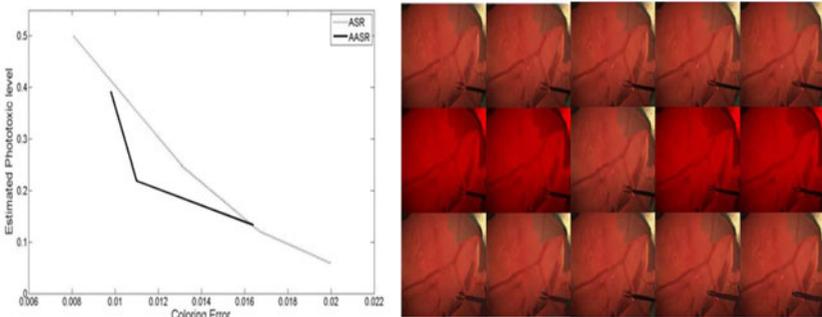


Fig. 2. (*Left*) Plot of estimated phototoxicity levels and recoloring error for both ASR and AASR. Notice that AASR is less phototoxic than ASR for every recoloring error level. (*Right*) Example image sequence of membrane peel. (top) Ground Truth, (middle) white and non-white illumination images triggered by AASR and (bottom) AASR image recolorization.

curve lies below that of ASR for every recolorization error level, hence achieving smaller total costs for the patient. In general, from our current experimental setup, AASR significantly outperform ASR, for the values of λ specified. Also, note that if all incoming images were registered to a reference frame (as in [9]) an additional reduction in colorization error would be expected.

Having observed that AASR provides a better way to model retinal-type scenes, we now present results on a typical image sequence of a chorioallatonic membrane peel from an 12 day old chicken embryo. In Fig. Fig. 2(*Right*) we show a small set of images from this sequence and the resulting recolorization using AASR ($\lambda = 0.5$). The original and recolored video sequence can be seen in a video included in the supplementary materials. In the video, the same peeling sequence is visible and each row corresponds to a different value for $\lambda = \{0.25, 0.5, 0.75\}$. The first column shows the original images. The second column displays the images provided by the device, while the third column shows the images recolored by AASR. The last column displays the retina color model over time. Other similar video sequences are provided in the supplementary materials.

Notice that in general, in image sequences which contain little membrane manipulations, few white light images are used. Since in this scenario our prediction model is capable of correctly estimating the colors of the retina, few white light images are necessary. Conversely, frames which show membrane peeling require more frequent white light illumination, in order to correctly render the colors. This indicates that the framework is able to choose which illumination type to use depending on the surgeon's actions.

4 Discussion and Conclusion

In this paper we have presented a novel algorithm that can be used to reduce toxic light exposure during retinal microsurgery. When used with the LED light source from [9], our recoloring scheme can dynamically choose the illumination based on the circumstances, reducing potential light induced retinal toxicity. Our algorithm balances the risks of phototoxic retinal damage with the illumination requirements of the operating surgeon to perform the surgical tasks. In this report we provide qualitative and quantitative evidence that this novel method reduces the dose of light, and hence retinal damage, while maintaining sufficient illumination to execute required surgical maneuvers safely.

While the results we have presented are in part dependent on the modeling choices of the cost functions, our framework is generic enough to accommodate a large number of functions. This being said, a natural future direction to improve the present work is to empirically determine the specific forms of the cost functions to use. Determination of these relationships would permit a truthful quantitative evaluation of the harm reduction. In ongoing and future work, we will be exploring these issues.

Acknowledgements

Funding for this research was provided in part by NIH Grant R01 EB 007969-01, a research subcontract from Equinox Corporation and by NSF Cooperative Agreement EEC9731478.

References

1. McDonald, H., Irvine, A.: Light-induced maculopathy from the operating microscope in extracapsular cataract extraction and intraocular lens implantation. *Ophthalmology* 90, 945–951 (1983)
2. Khwarg, S., Linstone, F., Daniels, S., Isenberg, S., Hanscom, T., Geoghegan, M., Straatsma, B.: Incidence, risk factors, and morphology in operating microscope light retinopathy. I. *Am. J. Ophthalmol.* 103, 255–263 (1987)
3. Byrnes, G., Antoszyk, A., Mazur, D., Kao, T., Miller, S.: Photic maculopathy after extracapsular cataract surgery. a prospective study. *Ophthalmology*, 731–737 (1992)
4. International Commission on Non-Ionizing Radiation Protection: Guidelines on limits of exposure to broad-band incoherent optical radiation (0.38 to 3). *Health Phys.* 73, 539–554 (1997)
5. van den Biesen, R., Berenschot, T., Verdaasdonk, R., van Weelden, H., van Norren, D.: Endoillumination during vitrectomy and phototoxicity thresholds. *Br. J. Ophthalmol.* 84, 1372–1375 (2000)
6. Poliner, L., Tornambe, P.: Retinal pigment epitheliopathy after macular hole surgery. *Ophthalmology* 99, 1671–1677 (1992)
7. Michels, M., Lewis, H., Abrams, G., Han, D., Mieler, W., Neitz, J.: Macular phototoxicity caused by fiberoptic endoillumination during pars plana vitrectomy. *Am. J. Ophthalmol.* 114, 287–292 (1992)
8. Banker, A., Freeman, W., Kim, J., Munguia, D., Azen, S.: Vision-threatening complications of surgery for full-thickness macular holes. *Ophthalmology* 104, 1442–1453 (1997)
9. Sznitman, R., Billings, S., Rother, D., Mirota, D., Yang, Y., Handa, J., Gehlbach, P., Kang, J., Hager, G., Taylor, R.: Active multispectral illumination and image fusion for retinal microsurgery. In: Navab, N., Jannin, P. (eds.) *IPCAI 2010. LNCS*, vol. 6135, pp. 12–22. Springer, Heidelberg (2010)
10. Museum of Broadcast Communication: Encyclopedia of Television, O.: <http://www.museum.tv/archives/etv/c/htmlc/colorization/colorization.htm>
11. Yatziv, L., Sapiro, G.: Fast image and video colorization using chrominance blending. *IEEE Transactions on Image Processing* 15(5), 1120 (2006)
12. Skora, D., Burinek, J., Zra, J.: Unsupervised colorization of black and white cartoons. In: *Int. Symp. NPAR, Annecy*, pp. 121–127 (2004)
13. Sznitman, R., Lin, H., Manaswi, G., Hager, G.: Active background modeling: Actors on a stage. In: *ICCV, Workshop on Visual Surveillance*, pp. 1222–1228 (2009)
14. Ham, W.J., Mueller, H., Ruffolo, J.J., Guerry, D., Guerry, R.: Action spectrum for retinal injury from near-ultraviolet radiation in the aphakic monkey. *Am. J. Ophthalmol.* 93, 299–306 (1982)