Design Considerations for a Computer-Vision-Enabled Ophthalmic Augmented Reality Environment

Jeffrey W. Berger[1,2,5] Michael E. Leventon[3,4] Nobuhiko Hata[4] William Wells[3,4,5]
Ron Kikinis[4,5]

September 5, 1996

Abstract

We have initiated studies towards the design and implementation of an ophthalmic augmented reality environment in order to allow for a) more precise laser treatment for ophthalmic diseases, b) teaching, c) telemedicine, and d) real-time image measurement, analysis, and comparison. The proposed system is being designed around a standard slitlamp biomicroscope. The microscope will be interfaced to a CCD camera, and the image sent to a video capture board. A single computer workstation will coordinate image capture, registration, and display. The captured image is registered with previously stored, montaged photographic and angiographic data, with superposition facilitated by fundus-landmark-based fast registration algorithms. The computer then drives a high intensity, VGA resolution video display with adjustable brightness and contrast attached to one of the oculars of the slitlamp biomicroscope. Preliminary studies with a modified binocular operating microscope interfaced to a Sun Ultra1 Workstation and an IBM-compatible PC demonstrates proof-of-principle. Robust, accurate fundus image montaging is accomplished with Hausdorff-distance-based methods. For photographic and angiographic data where the vessel gray levels vary from light to dark, and intensity-based correlation methods fail, image-preprocessing with smoothing, edge-detection, and thresholding facilitates registration. Non-real-time registration ($\sim 0.4 - 4.0$ CPU seconds) is achieved by non-optimized simple template matching (translation only, Matrox Inspector) or Hausdorff-distance-based (translation and scaling) algorithms performed on edge-detected fundus photographic and angiographic images, and on images of a model eye. Successful registration and image overlay of color, monochromatic, and angiographic images is demonstrated. To our knowledge, these studies represent the first investigation towards design and implementation of an ophthalmic augmented reality environment.

Key words: Augmented reality, eye surgery

Address correspondence to:
Jeffrey Berger, MD, PhD
Retina Service
Scheie Eye Institute
51 North 39th Street
Philadelphia, PA 19104
215-662-8675, Fax: 215-662-0133
Email: adinatamar@aol.com

[1]Department of Ophthalmology, Scheie Eye Institute, University of Pennsylvania, Philadelphia, PA
[2]Massachusetts Eye and Ear Infirmary, Boston, MA
[3]AI Lab, MIT, 545 Technology Sq, Cambridge MA
[4]Dept. of Radiology, Brigham and Womens Hospital, Boston MA
[5]Harvard Medical School, Boston, MA

Summary Page

a. **What is the original contribution of this work?**

To our knowledge, our investigations represent the first investigations towards the design and construction of an ophthalmic augmented reality environment

b. **Why should this contribution be considered important?**

Our application should allow for more efficacious treatment and evaluation of potentially blinding diseases.

c. **What is the most closely related work by others and how does this work differ?**

Similar investigations of image overlay/augmented reality in other surgical disciplines have be explored, but these technologies and methods have not been applied to diagnosis and treatment of eye diseases.

d. **How can other researchers make use of the results of this work?**

Our results demonstrate new methods for fast, two dimensional image registration, and we present a novel design for an image overlay applications.

e. **If this work extends or relates closely to some other work you have published or submitted, please state precisely how it differs from that work?**

Not applicable

f. **Select the categories which characterize your work**

An integrated system, in part tested in phantom studies.

1 Introduction

Diabetic macular edema and age-related macular degeneration (AMD) are the two major causes of visual loss in developed countries. While laser therapy for these and other diseases has prevented loss of visual function in many individuals, disease progression and visual loss following suboptimal treatment is common. For AMD, there is unambiguous evidence that incomplete laser photocoagulation of the border of a choroidal neovascular lesion is associated with an increased risk for further visual loss, while treatment beyond the borders unnecessarily destroys viable, central photoreceptors, further degrading visual function.

As a concrete example, in eyes with juxtafoveal choroidal neovascularization (CNV) secondary to ocular histoplasmosis, only 5% of eyes with laser treatment that covered the foveal side of the lesion with a narrow ($< 100\mu$) treatment border suffered severe visual acuity loss, while $\sim 25\%$ of eyes with either some of the foveal side untreated, or a wide border of treatment on the foveal side suffered severe visual loss (MPS, 1995). Similar results have been reported for AMD. Building on the recommendations proposed in Macular Photocoagulation Studies, clinicians generally attempt to correlate angiographic data with biomicroscopic images using crude, time-consuming, potentially error-prone methods. In a recent practical review (Neely, 1996), the author suggests that "...to assist you in treatment (of neovascular AMD), project an early frame of the fluorescein angiogram onto a viewing screen. Use the retinal vessels overlying the CNV lesion as landmarks. I suggest tracing an image of the CNV lesion and overlying vessels onto a sheet of onion skin paper. It takes a little extra time, but I find it helps to clarify the treatment area." Accordingly, precise identification of the treatment border during laser therapy by correlating the biomicroscopic image with fluorescein angiographic data (where the lesion extent is better delineated, see for example Figure 1) should be beneficial for maximizing post-treatment visual function. Diagnosis and treatment relies on synthesizing clinical data derived from fundus biomicroscopy with angiographic data, but methods for correlating these data, and for direct guidance of laser therapy, are not well-developed.

We are exploring techniques to overlay angiographic data on the real-time biomicroscopic slitlamp fundus image in order to guide treatment for eye disease (for example, to better define and visualize the edges of choroidal neovascular membrane, and improve identification of focal areas of leakage in diabetic macular edema). The biomicroscopic fundus image will be "augmented" in real time with available angiographic data. Text display and a "virtual pointer" will be incorporated into the augmented reality display to facilitate teaching, telemedicine, and real-time measurement and image analysis. Moreover, image superposition will allow for direct comparison with previous images to judge disease progression (for example, to judge progression or stability of AMD or cytomegalovirus retinitis—a common, blinding disease afflicting patients with acquired immunodeficiency syndrome) and allow for real-time identification of prior treatment areas. This technology is straightforwardly extended to an indirect-ophthalmoscope-based or operating-microscope-based system to facilitate correlative, teaching, and telemedicine applications in these environments.

The availability of formidable computation power has facilitated real-time image processing, allowing for investigation of useful augmented reality (AR) applications. Our interest in developing a slitlamp-biomicroscope-based AR environment are twofold. First, there is considerable interest in medical applications of AR for surgical planning and execution. The slitlamp biomicroscope is an ideal model system for medical AR applications since a) the slit-lamp fundus image is quasi-two-dimensional requiring less computational power than a three-dimensional application, and b) the fundus landmarks are suitable targets for machine-vision-based recognition and tracking. Second, "augmenting" the fundus biomicroscopic view will allow for a) more precise laser treatment for retinal diseases such as diabetic macular edema and age-related macular degeneration, b) teaching, c) telemedicine, and d) real-time image measurement and analysis.

In this report, we describe considerations for the design and implementation of an ophthalmic augmented reality environment and report on our preliminary studies in model systems.

2 Methods

2.1 Engineering Considerations

Design and implementation of the system described requires solution to a number of challenging problems. First, the system must be robust, tolerate a small area of fundus illumination, and respond rapidly to moderate changes in eye position during evaluation and treatment. We approach this by acquiring multiple, partially overlapping photographic images and montaging these images in non-real-time. Available angiographic and photographic data will be registered with the montaged data set to allow for rapid rendering. Since photographic and angiographic data may vary considerably in intensity, e.g. the blood vessels are dark in a monochromatic image and in the early angiographic phases but bright in

Figure 1: Monochromatic photographic (left) and fluorescein angiographic (right) image of an eye with age-related macular degeneration. The optic nerve is at the far right of each photograph, with the fovea located centrally. Note that the angiogram conveys additional information regarding areas of leaky blood vessels.

the late angiographic phase, intensity based registration algorithms are unsuitable. We are exploring edge-based registration algorithms, and have demonstrated successful image registration of color, monochromatic, and angiographic data following image pre-processing with smoothing, edge detection, and thresholding. The real-time fundus image will then be registered in near-real time with the montaged data set. Automated montaging of fundus images and fast fundus image registration are non-trivial problems, and appropriate, application-specific distance metrics must be deduced.

Next, the stored data must be "overlayed" on the biomicroscopic image. The computer rendering must be fast, and ergonomically well-tolerated. Lastly, technologies permitting interactivity will be developed. A remote observer (e.g. an expert supervising the local examiner, or a trainee "observing" a diagnostic or therapeutic maneuver) will view a real-time graphical display of the biomicroscopic image, and will be able to communicate questions or recommendations by text or voice. The remote observer will control a mouse-based "virtual pointer" to unambiguously identify regions of interest. The local user will similarly control a pointer to enable distance and area measurements, the results of which will be displayed in the text window. With both the local and remote user in control of a pointer, and able to communicate by voice or text, interactivity is facilitated.

An appropriate similarity metric for registration and montaging functions must be chosen. The "Hausdorff distance" (Huttenlocher et al., 1993; Rucklidge, 1995) is well-suited for our purposes since it a) runs well for edge-detected images, b) tolerates errors as well as the presence of extra or missing data points between data sets, and c) operates on an arbitrary, user-defined transformation function. For matching the fundus images, we are searching only over translation and scale, since there is very little rotation between images ($< 5°$) and since searching over rotation is much more computationally expensive.

The Hausdorff distance is computed only for positively thresholded (e.g. edge-detected) points and is defined by

$$H(A,B) = \max(h(A,B), h(B,A))$$

$$h(A,B) = \max_{a\in A} \min_{b\in B} ||a-b||$$

where $||a-b||$ represents Euclidean distance.

The Hausdorff distance identifies the point $a \in A$ that is farthest from any point of B, and measures the distance from a to its nearest neighbor in B. Equivalently, if $H(A,B) = d$, then

$$\forall a \in A, \exists b \in B \ni ||a-b|| \le d$$

or all points in A must be within a distance d from some point in B, with the most mismatched point at exactly a distance d from the nearest point in B.

Due to occlusion and outliers, not all points in A will have a meaningful correspondence with a point in B. We therefore use the partial Hausdorff distance measure:

$$H_K(A,B) = \max(h_K(A,B), h_K(B,A))$$

$$h_K(A,B) = \mathrm{K}^{\mathrm{th}}_{a\in A} \min_{b\in B} ||a-b||$$

An extension of the Hausdorff distance by Huttenlocher et al. (1993) defines a fraction of points in set A that lie within ϵ of some point in B:

$$F_\epsilon(A,B) = \min(f_\epsilon(A,B), f_\epsilon(B,A))$$

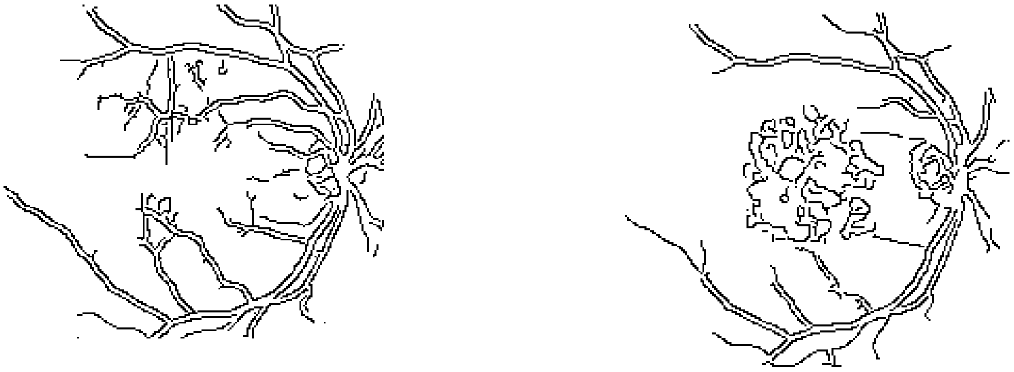

Figure 2: Binarized images of the monochromatic (left) and angiographic (right) images depicted in Figure 1. Edges were found using the Canny edge detector.

$$f_\epsilon(A,B) = \frac{\#(A \wedge B^\epsilon)}{\#(A)}$$

where B^ϵ is the point set B dilated by ϵ, or the Minkowski sum of B with a disk of radius ϵ. In other words, $f(A,B)$ is the fraction of points in A that lie within ϵ of some point in B. Instead of fixing K and minimizing $H_K(A,B)$, we fix the dilation ϵ and maximize $F(A,B)$.

2.2 Montaging

Fundus photographic and angiographic data will be superimposed on a real-time slit-lamp biomicroscopic image. The algorithm will initially be developed for macular applications. Specifically, images of the posterior fundus will be acquired and montaged into a single data set. Available angiographic data will be registered (off line, i.e. not in real time) with the montaged data set prior to biomicroscopic examination. Welch and coworkers (Markow et al 1993; Barrett et al 1994, 1996) and Becker et al. (1995) describe algorithms for real-time retinal tracking for automated laser therapy, however, the tracking algorithms require multiple templates to follow retinal vessels at several fundus locations, and therefore requires a large illumination area; tracked vessels must be visible at all times. This requirement is not consistent with our goal to allow for narrow beam illumination, and is more suitable for a fundus camera with monocular, wide angle viewing. Further, their tracking algorithms limit the search to a small area surrounding the previously tracked position, and is not tolerant of large changes in fundus position as might be encountered during a slit-lamp fundus, biomicroscopic examination.

We have implemented a modified, bidirectional Hausdorff-distance-based (Huttenlocher et al. 1993; Rucklidge, 1995) approach for image montaging. Multiple, partially overlapping fundus images are acquired. The images are then smoothed and edge detected (Canny, 1986). To match two images together, the maximum Hausdorff fraction over translation and scale is found. Given that the fundus images only partially overlap, if all points in both images are used in the Hausdorff fraction calculation, then the resulting fraction will be very low. In fact, it is likely that a random transformation where the fundus images have more overlap will have a larger fraction if all points are considered.

Therefore, when computing the Hausdorff fraction $F(A,B)$ for a given transformation, A and B are the sets of points from the two images that are in the overlap region of the images. All points that fall outside the overlap region for a certain transformation are ignored. To prevent a false match with a high fraction where only a very small percentage of points lie within the overlap region, a minimum percentage of image points (currently 10%) must lie in the overlap region for the transformation to be considered a possible match.

To build the montage, the position and scale of all images must be determined relative to a fixed coordinate system (that of the montage). Under the current design, a "central" image is chosen a priori by the user, and the coordinate system of this image is used for the montage. Generally, the "central" image is easily identifiable by the user because it contains both the fovea (which should be approximately centered) and the optic nerve. All the other images are then matched to the central image, yielding a transformation that maximizes the Hausdorff fraction.

All images that match the central image with a Hausdorff fraction above some threshold (currently 70%) are considered to be correctly placed in the montage. Any images that do not match the central image are then matched against the other images. Although worst case,

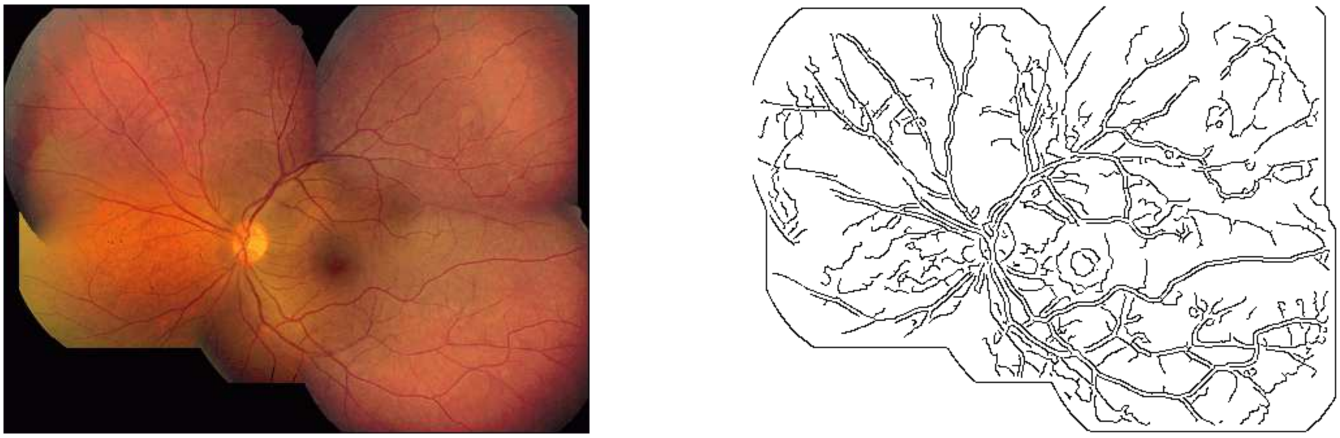

Figure 3: Montage of five distinct, partially overlapping fundus images as determined by the Hausdorff distance-based methods. The right image is the result of applying the Canny edge detector on the left image.

this is $O(n^2)$ matches, in practice, most if not all of the images match the central image.

Once the positions and scales of all the images are known relative to the central image, the montage image can be built. One simple method of building the montage is to average the intensity values of the overlapping images at every pixel. The montage image, I, is computed as follows:

$$I(x,y) = \frac{\sum_{i=1}^{N}(\delta_i(x,y) \times I_i(x,y))}{\sum_{i=1}^{N}\delta_i(x,y)}$$

where $\delta_i(x,y)$ is 1 if (x,y) lies inside image I and 0 otherwise.

However, in general, the average intensities of the images are not equal, so edge artifacts are introduced at the image borders. Furthermore, the average intensity in any one image generally varies over the image. The fundus images are often brighter near the center and lose brightness and contrast near the image borders. The vessels we are using as landmarks disappear near the image borders. Thus, in addition to the edge artifacts, simply averaging images also degrades the visibility of the vessels in the montage near image borders.

Therefore, when blending all the images to form the montage, we use a convex combination that puts greater weight on pixels closer to the center of an image:

$$I(x,y) = \frac{\sum_{i=1}^{N}(d_i^2(x,y) \times I_i(x,y))}{\sum_{i=1}^{N}d_i^2(x,y)}$$

where d_i is the distance from the point (x,y) to the nearest border of image i if (x,y) is inside the image, and 0 if (x,y) is outside the image. This blending removes many of the artifacts due to varying contrast and intensity. Figure 3 shows the result of montaging five distinct, partially overlapping fundus images.

2.3 Registration

Our preliminary studies have demonstrated that preprocessing with edge-detection algorithms allow for registration of photographic and angiographic data where the vessel gray levels may vary from black to white (Figure 2). Early non-real-time studies demonstrating the applicability of edge-detection-based algorithms for image registration were performed on a PC compatible computer with Matrox Inspector (Quebec, Canada) software.

For potential near-real-time application, the biomicroscopic image will be acquired, undergo edge detection and thresholding (binarization). A pruned search tree will then be searched extending from no displacement to the maximum displacement as determined by the frame rate and image size. The Hausdorff distance will be minimized, with computational frugality achieved by a) the small size of the illuminated biomicroscopic image, b) calculations with binarized images (permitting logical operations), c) restricting our template search to translational displacements with correction for scaling, and d) limiting the search area. Crucial to efficient application is the off-line montage creation, allowing for instant access to photographic and angiographic data. Optimum temporal performance may be achieved with a small illuminated fundus area, i.e. a small template, with the remainder of the image presented efficiently from previously stored data. Along these lines, the influence of the size of the illuminated biomicroscopic image on the speed and fidelity of registration will be explored.

2.4 Image Overlay

Ultimately, a standard ophthalmic slitlamp biomicroscope will be used together with either a fundus con-

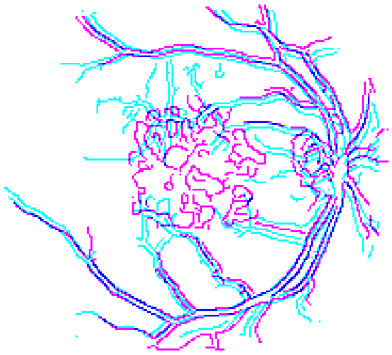

Figure 4: Registration and superposition of the images in Figure 2 by a Hausdorff-distance-based algorithm incorporating translation and scaling parameters. The angiographic and photographic data are depicted in magenta and cyan, respectively.

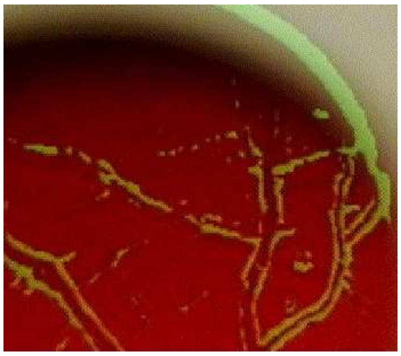

Figure 5: Image overlay of previously stored edge-detected image of a model eye onto a real-time biomicroscopic image of the model eye. The model eye vessels are not visible in this photograph, but are identified by the image overlay.

tact lens or a hand-held 60-90 diopter non-contact lens to view the ocular fundus. As an analog to the operating microscope, the slit-lamp biomicroscope is perfectly suited to serve as an imaging and display conduit for AR applications. We are using a modified Zeiss OPM1-DFC operating microscope that permits image overlay of previously stored data on the real-time biomicroscopic image. The microscope is interfaced to a CCD camera, and the image is sent to a frame-grabber and digitizer. The image is processed on a Sun Ultra I workstation where registration calculations are performed. The workstation drives a custom-made high-intensity video display with VGA resolution and adjustable brightness and contrast, and presents images in a false color (e.g. green) to maximize visibility. Algorithm development has utilized both Sun and PC compatible workstations.

3 Results

Modified Hausdorff-distance-based methods are effective for rapid, accurate montaging of multiple, partially overlapping fundus images. Transformation functions as derived from edge-detected, binarized images allow for construction of a single data set (Figure 3).

For photographic and angiographic data where the vessel gray levels vary from light to dark, and intensity-based correlation methods fail, image-preprocessing with smoothing, edge-detection, and thresholding facilitates registration. Following image pre-processing, Matrox inspector software (using a simple template matching algorithm) allows for identification of corresponding regions in photographic and angiographic images. Translation-only transformation functions are thus defined, and the utility of edge-based registration methods is demonstrated.

Non-real-time registration ($\sim 0.4 - 4.0$ CPU seconds) is achieved by non-optimized Hausdorff-distance-based (translation and scaling) algorithms performed on edge-detected fundus photographic and angiographic images, and on images of a model eye. Figure 4 depicts edges corresponding to vessel skeletons demonstrating high-fidelity registration of photographic and angiographic data (corresponding to the images in Figures 1 and 2).

Image overlay is demonstrated by simple image superposition on a computer monitor. In addition, image overlay is accomplished on the operating-microscope-based system described. Photographic images of a model eye were acquired and edge detected. The edge-detected images were then rendered in green, and overlayed on a real-time biomicroscopic image of a model eye (Figure 5). The real-time view through the binocular objectives was then a merged image of the green edges and the real-time model eye fundus image. Accordingly, proof-of-principle for an ophthalmic augmented reality environment is demonstrated.

4 Discussion

Our preliminary studies represent the first investigations towards the design and implementation of an ophthalmic augmented reality environment. As an analog to the operating microscope, the slit-lamp biomicroscope is perfectly suited to serve as an imaging and display conduit for AR applications. Further, since our application is quasi-two-dimensional, computational demands are far less than in three-dimensional image registration tasks, potentially allowing for efficient real-time performance.

Image registration has played a major role in facilitating augmented reality (AR) applications. As described by Feiner (1993), an AR system utilizes computer generated graphics, such that the virtual world is superimposed on, and enriches, the real world. Whereas virtual reality (VR) is the graphical construction of a synthetic environment, AR extracts information from the real world and augments it (Bowskill and Downie, 1995; Caudell, 1994).

Although there has been an explosive development of investigation into VR applications in medicine, AR applications might have a far greater utility, but have received much less attention (O'Toole et al. 1995). An early study described the superposition of ultrasound images on the abdomen, using a position-tracked, see-through head mounted display (Bajura et al. 1992). Recently, there has been great interest in neurosurgical applications of AR. Specifically, the intraoperative registration of CT, MR, and PET images on a living patient in the operating room may greatly facilitate surgical planning and execution (Gleason et al. 1994; Edwards et al. 1995; Grimson et al. 1996). Typically, registration of the patient with the radiographic data is accomplished by tracking fiducial markers placed on the skin surface and tracking the position of the operating microscope. Our application does not require fiducial markers or image tracking. The real-time image is registered with a previously stored, montaged data set.

We present design considerations for an ophthalmic augmented reality environment, and demonstrate the technical feasibility for each element of the proposed system. In subsequent investigations, angiographic images will be registered off-line to create a three-dimensional vector containing time-dependent angiographic data. Specifically, the individual frames from previously acquired angiographic studies will be registered with each other with high precision, and subsequently registered with the montage. Therefore, the angiographic data will be presented in cine form in the augmented reality environment, offering dynamic visualization of (for example) leakage from a CNVM, diabetic microaneurysm, or central serous retinopathy. Further, image registration and overlay will allow for real-time image comparison to judge disease progression and guide treatment. Finally, technologies permitting interactivity as described above will be incorporated to allow for telemedicine applications.

References

[1] M. Bajura, H. Fuchs, R. Ohbuchi, "Merging virtual objects with the real world: Seeing ultrasound imagery within the patient", in *Computer Graphics*, **26**:203-210, 1992.

[2] S.F. Barrett, M.R. Jerath, H.G. Rylander, A.J. Welch, "Digital tracking and control of retinal images", in *Optical Engineering*, **33**:150-159, 1994.

[3] S.F. Barrett, C.H.G. Wright, M.R. Jerath, S. Lewis, B.C. Dillard, H.G. Rylander, A.J. Welch, "Computer-aided retinal photocoagulation system", in *J Biomedical Optics*, **1**:83-91, 1996.

[4] D.E. Becker, J.N. Turner, H. Tannenbaum, B. Roysam, "Real-time image processing algorithms for an automated retinal laser surgery system", in *Proc IEEE 2nd Int Conf Image Proc*, 1995, p 426-429.

[5] J. Bowskill, J. Downie, "Extending the capabilities of the human visual system. An introduction to enhanced reality", in *Computer Graphics*, **29**:61-65, 1995.

[6] J.F. Canny, "A computational approach to edge detection" in *IEEE Trans. Pat. Anal. Mach. Intel.*, **8**:34-43, 1986.

[7] T.P. Caudell, "Introduction to augmented and virtual reality", in *Proc SPIE (Telemanipulator and Telepresence Technologies)*, **2351**:272-281, 1994.

[8] P.J. Edwards, D.L.G. Hill, D.J. Hawkes, R. Spink, A.C.F. Colchester, A. Strong, M. Gleeson, "Neurosurgical guidance using the stereo microscope" in *Proc First Int Conf Computer Vision, Virtual Reality and Robotics in Medicine*, Nice, France, p 555-564.

[9] S. Feiner, M. Macintyre, D. Seligmann, "Knowledge based augmented reality" in *Communications of the ACM*, **36**:53-61, 1993.

[10] P.L. Gleason, R. Kikinis, D. Altobelli, W. Wells, E. Alexander, P. Black, F. Jolesz, "Video registration virtual reality for nonlinkage stereotactic surgery", in *Sterotactic Funct Neurosurgery*, **63**:139-143, 1994.

[11] W.E.L. Grimson, G.J. Ettinger, S.J. White, T. Lozano-Pérez, W.M. Wells III, and R. Kikinis, "An Automatic Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization", in *IEEE Trans. Medical Imaging*, **15**:129, 1996.

[12] D.P. Huttenlocher, G.A. Klanderman, W.J. Rucklidge, "Comparing images using the Hausdorff distance" in *IEEE Trans Patt Anal Mach Intell*, **15**:850-863, 1993.

[13] M.S. Markow, H.G. Rylander, A.J. Welch, "Real-time algorithm for retinal tracking", in *IEEE Trans Biomed Engineering*, **40**:1269-1281, 1993.

[14] Macular Photocoagulation Study Group, "The influence of treatment extent on the visual acuity of eyes treated with krypton laser for juxtafoveal choroidal neovascularization" in *Arch Ophthalmol*, **113**:190-194, 1995.

[15] P. Nagin, B. Schwartz, K. Nanba, "The reproducibility of computerized boundary analysis for measuring optic disc pallor in the normal optic disc" in *Ophthalmology*, **92**:243-251, 1985.

[16] K.A. Neely, "How to be more successful in laser photocoagulation", in Ophthal Times, 1996, pp 103-108.

[17] R.V. O'Toole, M.K. Blackwell, F.M. Morgan, L. Gregor, D. Shefman, B. Jaramaz, DiGioia, T. Kanade, "Image overlay for surgical enhancement and telemedicine", in *Interactive Technology and the New Paradigm for Healthcare*, K. Morgan et al., eds, IOS Press, 1995.

[18] W.J. Rucklidge, "Locating objects using the Hausdorff distance" in *Proc IEEE Int Conf Computer Vision*, 1995, p 457-464.