THE THEORY AND DEVELOPMENT OF

A NONINVASIVE RETINAL FLUORESCENCE SCANNER

WITH APPLICATION TO EARLY DIAGNOSIS OF DIABETIC RETINOPATHY

by

JONATHAN MARC TEICH

B.S., California Institute of Technology (1976)
S.M. and E.E., Massachusetts Institute of Technology (1979)
M.D., Harvard Medical School (1983)

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN ELECTRICAL ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

July 1985

(c) Jonathan M. Teich 1985

The author hereby grants to M.I.T. and to Boston University Medical Center permission to reproduce and to distribute copies of this thesis document in whole or in part.

Signature Redacted

Signature of Author_____
Department of Electrical Engineering and Computer Science
July 2, 1985

Signature Redacted

Certified by_____
William M. Siebert
Thesis Supervisor

Signature Redacted

Accepted by_____
Arthur C. Smith
Chairman, Departmental Committee on Graduate Students

THE THEORY AND DEVELOPMENT OF A NONINVASIVE RETINAL FLUORESCENCE

SCANNER WITH APPLICATION TO EARLY DIAGNOSIS OF DIABETIC RETINOPATHY

by

JONATHAN MARC TEICH

Submitted to the Department of Electrical Engineering and
Computer Science on June 27, 1985 in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Electrical Engineering

ABSTRACT

A noninvasive method of assessing tissue oxygenation in the human retina in vivo by observation of natural reflectance and fluorescence in the eye is discussed. The method makes use of the fact that flavin adenine dinucleotide (FAD) and other components of the mitochondrial electron-transport chain change between their oxidized and reduced forms depending on the local oxygen supply, and that these two forms have different absorption and fluorescence properties which can be measured. Thus by measuring these properties in a given local tissue, one may extract information about the local oxygen adequacy. Determination of oxygen by this principle could have important consequences for the early diagnosis and treatment of diabetic retinopathy and other metabolic eye diseases, because of the possible role of retinal hypoxia in the early stages of these diseases.

The light levels produced by retinal FAD fluorescence are extremely low, and the signal is altered by the absorption and fluorescence of the cornea and lens of the eye, as well as by the other chemicals of the retina. A theory of ocular reflectometry and fluorometry is advanced in order to understand what light emission would be expected under different metabolic conditions, and to optimize the measurement conditions for maximum sensitivity to oxygen changes. Experiments to test the theoretical findings are described.

The complexity of the light path to and from the retina, the dimness of the FAD fluorescence, and exposure-time and safety requirements have necessitated the development of a rapid, ultrasensitive ocular fluorometer and reflectometer and a base of image-processing algorithms to process raw data. The development of this scanner and the image processing techniques used in the analysis of retinal metabolism from FAD fluorescence images are described.

Thesis Supervisor: Dr. William M. Siebert

Title: Professor of Electrical Engineering and Computer Science

CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

Many persons have given guidance, ideas, and assistance during the course of this project. Certainly foremost among these is Prof. Jerrold Shapiro, director of the Ocular Physiology Laboratory at Boston University Medical Center, who initially conceived the idea of measuring native retinal fluorescence and who has provided technical, organizational, and (with Jane, Jody, and Jeremy) personal support over the entire course.

Prof. William Siebert ably served as thesis supervisor for the latter part of the project, and provided much helpful advice. Dr. John Loewenstein's ophthalmologist's knowledge and viewpoint were essential for the conceptual and experimental phases of the work. All of the above also served as readers of this thesis, as have Dr. John Tole and Dr. Walter Olson.

Amal Jeryes' friendship and help in the engineering and experimental work are much appreciated. Doug Philpott and Vince Pawlowski provided additional technical assistance. David Israel maintained the operating system on which program development and text writing was done. And Dana Klein's artistic talent will be obvious to any who view the illustrations and diagrams herein.

I gratefully acknowledge the financial assistance of the Fannie and John Hertz Foundation, the National Institutes of Health and the Harvard-MIT Division of Health Sciences and Technology over the course of this work.

While such scientific, technical and financial contributions are sincerely appreciated, the greatest measure of gratitude is owed for the companionship and personal support given by those above, by my family and by the many friends I have known from the MIT Biomedical Engineering Center,

the MIT Chorallaries, Harvard Medical School, the Brigham and Women's Hospital, Caltech, and anywhere else good people walk. This thesis could never have been completed without their continual gifts of friendship and patience over the very long course of the work.

CHAPTER 1

INTRODUCTION

This thesis describes the development and theory of a noninvasive means of observing tissue oxygenation via natural reflectance and fluorescence in the eye, particularly the fluorescence of flavin adenine dinucleotide (FAD). Determination of oxygen by this principle, which has been used in the past to assess hypoxia in exposed tissue preparations, could have important consequences for the early diagnosis and treatment of diabetic retinopathy and other metabolic eye diseases, because of the possible role of retinal hypoxia in the early stages of these diseases. The ultimate role for this new method is to study the relationship between retinal reflectance and fluorescence, tissue hypoxia, and the progression of these disease states.

The eye is a complex optical structure, with multiple components in cornea, lens, and retina which absorb light and fluoresce. The correct detection of changes in retinal oxygen status amid these varied interfering signals requires advances in the theory of fundus reflectometry, so as to more precisely understand what light emission would be expected under different metabolic conditions, and so as to better set up experimental parameters for maximum sensitivity to oxygen changes. The existence of these interfering signals, the dimness of the FAD fluorescence, and exposure-time and safety requirements have necessitated the development of a rapid,

highly sensitive ocular fluorometer and a base of image-processing algorithms to process the raw data. This theoretical study and development work constitute the major portions of this dissertation.

The thesis has 11 chapters. This first chapter provides an introduction to the pathology of diabetic retinopathy and explores the role of hypoxia in the disease process. I will describe other methods of studying oxygenation in biological systems, with their advantages and their failings, and present a rationale for the use of FAD fluorescence as a possibly superior method.

Chapters 2 and 3 present a mathematical theory of ocular fluorometry and reflectometry, with special attention to the problem of sorting out FAD fluorescence from the reflections and fluorescences due to other molecular phenomena in the eye. Chapter 4 describes the experiments which supported the theory, demonstrated the feasibility of ocular FAD fluorescence measurement and which provided the first evidence of fluorescence changes in the eye with hypoxia. This chapter also describes the need for, and the parameters of, a sensitive, flexible means of measurement with image processing capabilities.

Chapters 5-6 describe the development of an ultrasensitive scanning ocular reflectometer/fluorometer which attempts to meet these needs. The novel optical system designed to suppress artifact from other eye structures while providing precise retinal images is described first, followed by a description of the specialized data acquisition, scan control, real-time and digital-processed display techniques developed for imaging and analysis.

Chapter 7 describes some of the key algorithms which drive the system and allow it to perform image acquisition and storage, variable scanning,

and multiple-wavelength operation. The interactive processing necessary for proper adjustment and for monitoring of the quality of data while the experiment is in progress is discussed, as are some of the system utilities built to allow easier use and rapid development and testing of new software. Chapter 8 is concerned with the image processing algorithms themselves and their use in displaying, enhancing, and combining pictures as required by the mathematics of the model. Chapter 9 describes some safety features developed for this imaging device and includes a short presentation of light-damage hazards present in fundus cameras and other devices which send light into the eye.

Chapter 10 is devoted to a description of some of the experimental procedures that have been performed on chemical samples, humans, rabbits, and monkeys with the instrument, with interpretation of the results of some of these experiments which have brought us toward more reliable observation of human native retinal fluorescence. Pitfalls that were encountered and changes made as a result are discussed. The closing section outlines the most productive possibilites for future experiments, for potential improvements to the instrument and for expansion of the image processing procedures. It discusses the possibilities of ongoing clinical research which will use the fundus imaging instrument in a prospective study to observe the correlation between fluorescence changes and later disease in the local area.

1.1. Pathophysiological considerations

DIABETIC RETINOPATHY (DR) is one of the leading causes of new cases of blindness in the United States[1]. Diabetes mellitus is found in approxi-

mately 1% of the overall population[2], and diabetics have 25 times the risk of developing legal blindness compared to nondiabetics[3]. Over 50% of persons who have had diabetes mellitus for ten or more years develop retinopathy, which begins insidiously and can progress to severe visual loss. This implies that nearly 1 million people in this country alone have DR or are at substantial risk(Fig. 1.1)[4,5]. Studies have indicated that from 3 to 8 percent of ALL diabetics (including those who have had the disease less than 5 years) are legally blind[3]. The risk of severe retinopathy is especially high in those diabetics who were diagnosed at an early age, who have higher levels of glycosylated hemoglobin (hemoglobin A_1C), and who have higher systolic and diastolic blood pressure.

1.1.1. Pathology. It is widely suspected that diabetic retinopathy, along with other proliferative retinopathies -- sickle cell retinopathy, some forms of retinal vein obstruction, hypertensive retinopathy, Takayasu's arteritis -- has a basis in the development of retinal ischemia[6].

The human retina has two blood supplies: the retinal circulation, which supplies the inner (anterior) retinal layers and the choroidal circulation, which supplies the photoreceptors and some of the outer retinal connecting cell layers. The choroidal circulation lies behind the retinal pigment epithelium (RPE) and is thus not seen in typical fundus photographs (Fig. 1.2), except in albinos. The first step in the pathogenesis of DR, based on histologic studies of early cases, appears to be an increase in the thickness of the basement membranes of small capillaries in the retinal circulation. Other changes in the small vessels occur, including loss of the pericytes which support the endothelial cells, leading to the formation of microaneurysms, and possibly to micro-hemorrhages ("dot and blot"

-4-
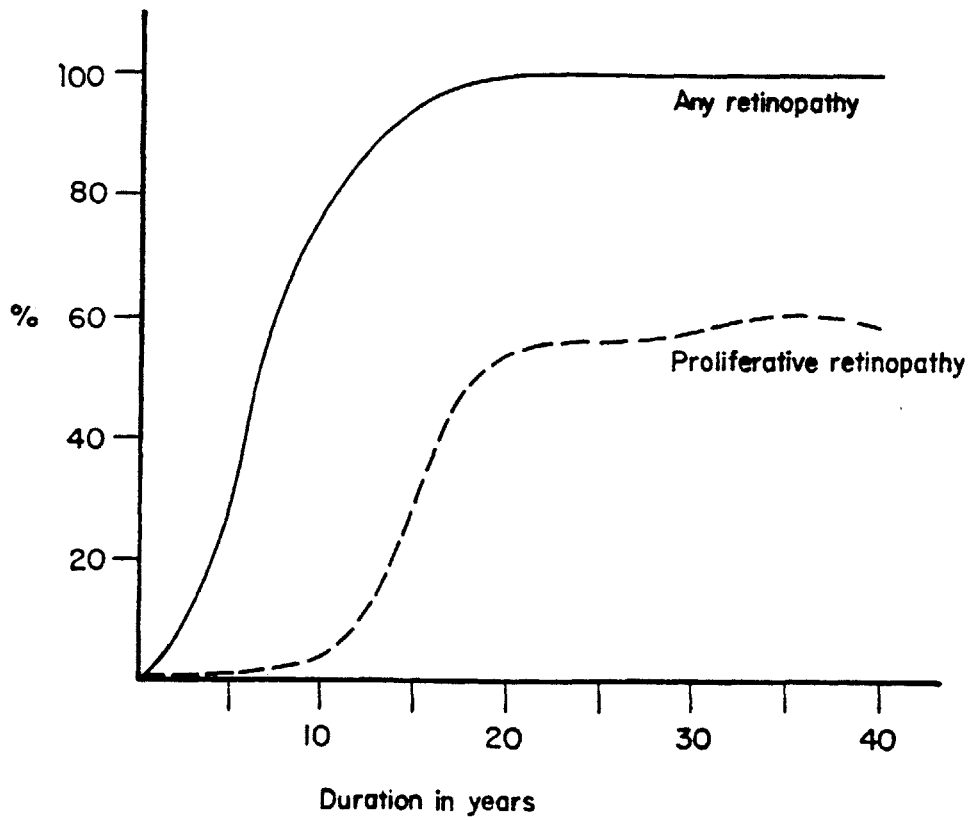
Figure 1.1. The relative risk of developing any DR (solid line) or prolif-
erative DR (broken line) as a function of duration of diabetes [3].

Figure 1.2. The retina of the eye seen through a fundus camera. The veins
are seen as solid vessels, while the arteries have a reflective
streak when seen in photographs such as this. In this fundus photo
several "dot and blot" hemorrhages can be seen (the most prominent
is in the lower left-hand corner of the field).

hemorrhages)[8]. There may actually be initial vasodilatation as a response to hypoxia[9].

These changes are characteristic of the background stage of DR. The junctions between cells cannot be maintained, possibly due to lack of metabolic energy secondary to hypoxia. The breakdown of these junctions causes leakage of serum from capillaries; higher viscosity material such as lipids can be clumped at the border of the vessel and lead to the hard exudates sometimes seen in early retinopathy.

If this leakage of fluid (edema) occurs over the macula where central vision takes place, visual blurring can occur; otherwise there are no visual symptoms at this point, nor is there any useful treatment available at this stage. It is more likely that if retinal edema occurs, it will be diffuse; however, a minority of practitioners will perform photocoagulation of edematous areas (see below).

Is hypoxia the primary underlying cause of the changes of DR? The current popular wisdom is that it is. Reduced oxygen supply can be responsible for the breakdown of pericytes and cell-cell junctions, and for the infarction of neural fibers in the next stage of DR (see below). Ditzel[9] has suggested that the hypoxia stems from two factors: the increased amount of glycosylated hemoglobin (hemoglobin A_1C) which releases oxygen to tissues more poorly than normal hemoglobin, and increased levels of 2,3-diphosphoglycerate (2,3-DPG) which also adversely affects the delivery of oxygen by hemoglobin. It has also been proposed, however, that the hypoxia is secondary to the capillary closure itself; it had even been suggested[10] that the initial event was hyperoxia due to the effect of hyperglycemia on blood vessels, followed by the excessive response of a normal physiologic feedback mechanism to bring about the secondary hypoxia.

Although this now appears to be a less likely model, the controversy between these two hypotheses cannot be resolved in the absence of direct measurement of the oxygenation of the tissue itself.

In the pre-proliferative stage of DR, tissue ischemia may be worsening. Seen in a funduscopic (retinal observation) exam are "cotton-wool spots"(Fig. 1.3), cloudy white densities which histologically are areas of infarction (cell death) secondary to ischemia of the nerve-fiber layer of the retina. Capillaries may close completely in this stage as their breakdown continues; behind them, small arteries deprived of an outlet may form fistulas to nearby veins; these can also leak[11,12].

The proliferative stage of DR is characterized by formation of new blood vessels. These vessels can appear anywhere on the surface of the retina; they often can be found on the optic disk, especially, and are sometimes but not always found on the borders of the non-perfused areas of the preproliferative stage[6]. These new vessels are believed to be formed in response to the demand for oxygen by the starved tissue, perhaps by means of a humoral "angiogenesis factor" such as is being postulated for vascular tumors. These vessels grow into the vitreous cavity. Blood flow through leaky new vessels leads to preretinal hemorrhage and vitreous hemorrhage. Hemorrhage is the second way in which DR affects vision (macular edema is the first).

There is also fibrous proliferation in this stage. Possibly in a manner akin to scarring in non-neural tissue, glial cells come in and proliferate along with a matrix of collagen. This fibrous growth often leads to contracture of the site, and can cause retinal detachment because the forces of contracture are greater than the very light force holding the retina onto the retinal pigment epithelium posteriorly.

Figure 1.3. Preproliferative retinopathy. Cotton-wool spots (CWS) and hard exudates (HE) are evident.

Proliferative retinopathy is frequently associated with significant visual loss. Visual loss due to preretinal hemorrhage does not correlate with earlier development of macular edema in the same patient. Hemorrhage is a more severe way for DR to impair sight.

1.1.2. Treatment The treatment of choice for severe proliferative DR is laser panretinal photocoagulation, which is the use of a high-intensity argon laser to burn sections of the retina. The term "coagulation" is not used in the sense of preventing bleeding from leaky vessels; rather, the tissue itself is destroyed. The philosophy is to reduce the metabolic load (oxygen demand) on the impaired vasculature, and so to reduce the drive for neovascularization. In addition, by destruction of the photoreceptors, the main sink for choroidal oxygen is removed, and the choroid is permitted to help nourish the inner retinal tissue.

The treatment, usually applied in two sessions, is applied so as to burn a "checkerboard" out of the usable retinal surface; between two burned regions there is preserved tissue. The macula and optic disk are not treated. The patient is expected to scan around the fields of lost vision, with the brain filling in the gaps.

The bleeding from new, leaky vessels passes into the vitreous humor in front of the retina. If the blood collection in the vitreous becomes significant, either before or after laser treatment, vitrectomy may also be required. In this procedure the opacified vitreous is aspirated and cut by a mechanized cutting device, replaced by a saline solution of appropriate density. This procedure must be done carefully, as the pressure of the vitreous plays a major role in holding the tissue-thin retina in its place against the retinal pigment epithelium of the eye.

Photocoagulation is performed in a set geographical pattern because

-10-

there is no reliable way to tell where proliferation will occur next. The hope is that by treating a large segment of the entire retina, that excess oxygen will diffuse over all potential sites of neovascularization, including the optic disk. However, such large-scale treatment carries a price, not only in amount of retinal tissue loss but in large amounts of scarring due to laser burns. The Diabetic Retinopathy Study Group, in its second report of findings of a study of photocoagulation, noted that inhibition of progression of retinopathy and lessening of the severity of diabetes-induced visual loss occurred in all stages of DR which were treated. However, only in patients with moderate to severe proliferative retinopathy did this benefit clearly outweigh the deleterious effects, which include loss of central acuity, peripheral visual field loss, and night vision deficits. Because of these harmful side effects the Study Group [14] could not clearly recommend treating other than already severe cases. The alternative -- waiting to see if severe retinopathy develops -- reduces the side effects but leaves the patient open to severe visual loss.

Photocoagulation does nothing for the damage caused by fibrous proliferation in the retina; if anything, the scarring caused by photocoagulation exacerbates that problem. This is another reason to limit the extent of laser treatment. Nor does photocoagulation itself necessarily stop the process of progressive retinal vascular change -- at a later time even the reduced metabolic load is too much for the deteriorating circulation, so that some patients require repeat laser treatment.

Presumably, if, as is widely believed, the neovascularization is a response to local tissue hypoxia, then knowledge of which areas are hypoxic could allow treatment over a more limited range. The treatment would then reduce local metabolic load and prevent vessel proliferation while reducing

harmful side effects. Although further treatments might be necessary later if hypoxia develops elsewhere, the field loss and other effects would be delayed significantly; the saved regions could provide valuable islands of vision if retinal ischemia spreads over the eye. Further, many diabetic patients may develop only limited regions of hypoxia and thus would have less unnecessary laser scarring over a lifetime.

This latter case may become more important if other efforts in the control of DR, particularly attempts at ever more precise and continuous blood glucose control, are successful. Recent studies[15,16] on blood glucose control suggest that some of the basement membrane changes which may be associated with the progression of diabetic microangiopathy are delayed by fine blood-glucose control; however, studies showing whether DR itself is actually delayed by better control will not be available for some time yet. If neovascularization occurs in more limited regions of the eye in the future, then the advantages of spotting it early and preventing vitreous bleeds will be even more important, as limited and localized treatment may be all that is necessary.

Thus there is a need for a means of assessing tissue oxygenation and determining the relation of local hypoxia (which would occur earlier than visible changes of DR) to future progress of retinopathy in the region.

If it happens that hypoxia is not a primary cause of the neovascularization of DR, then perhaps the whole rationale for laser photocoagulation must be rethought. That the technique of photocoagulation has some success does not indicate that the basis for it is sound, and it is dangerous to become complacent about a treatment without a solid pathophysiological base.

1.2. Measuring Oxygenation

Among the more significant methods of assessing oxygenation in vivo are oxygen electrodes[17], transcutaneous oxygen measurements, pyrene butyric acid infusion[18] and hemoglobin reflectometry[19].

Transcutaneous oxygen measurements are good on a crude scale for determining whole-body oxygenation, and are used, for example, in the monitoring of infants rather than placing an arterial catheter to measure blood gases; however, they are not precise, nor can they be used for local tissue measurements. A transcutaneous monitor placed on the eye would at best measure corneal oxygenation, which is a function more of the atmosphere than of the blood supply.

Oxygen electrodes, which are based on a half-cell electrochemical reaction of platinum, permit direct measurement of local oxygen. They are, of course, extremely invasive, and suffer in that they themselves cause consumption of oxygen to a greater extent than the other methods.

Jobsis[18] developed the idea of pyrene butyric acid (PBA) infusion for in situ oxygen concentration in the isolated perfused heart because, when excited with light at 328-340 nm, PBA fluoresces (in the 375-400 nm band) and the fluorescence is quenched quantitatively as oxygen concentration increases. Although this is invasive in that one must introduce a foreign chemical in the bloodstream, it did demonstrate the principle of a fluorescent marker for oxygen measurement. In fact, PBA is related to the pyridine nucleotides which form the basis of our proposed solution (see below).

Hemoglobin has a well-established spectral absorption curve with varying oxygen concentration. This is the principle used in arterial blood gas monitoring -- the oxygen concentration is determined by the response of the

-13-

sample to light of specific wavelengths. This is a simple, effective method when only blood is present; it has also been used in a few instances to measure blood oxygenation in the eye. The oxygen content in the major vessels of the retina is not useful in studying DR and other diseases because the blockage and hypoxia are due to changes in the arterioles and capillaries which carry blood from the major vessels to the tissue, and it is in these capillaries that the O_2 gradient occurs. It is possible to use this method to look at the capillary layer itself, which will give a combined picture of the pre- and post-blockage oxygen concentration. If all the oxygen in this layer mixed directly with the tissue (by diffusion), this would be an effective way to demonstrate tissue oxygenation. However, especially in disease states, much of the blood is blocked from the diffusion capillaries and passes back out via the small venules while maintaining high PaO_2 -- thus presenting a false picture.

In addition, workers in this area note "interfering signals"[20] caused by other tissue components which can alter the hemoglobin spectrum. Some of this interference is due in fact to the respiratory-chain (electron-transport-chain) enzymes which we will study, and the interference is therefore itself oxygen-dependent -- a fact not accounted for in most explanations of the hemoglobin reflectometry method.

What is really needed is a marker which 1) already exists in the eye, 2) responds to light which may be noninvasively applied to the eye, and 3) describes the oxygenation of the retinal tissue itself, which is the tissue at risk in retinopathy.

1.3. Fluorescence in the Electron Transport Chain

In the normal aerobic conversion of fats, proteins, and sugars to energy in biological systems, each of these three sources has a specific metabolic chain which breaks down the molecule, yielding ATP and reducing power in the form of reduced nicotinamide adenine dinucleotide (NAD) (a member of the class of pyridine nucleotides) and flavin adenine dinucleotide (FAD), the latter bound to proteins known as flavoproteins (Fp). NAD and FAD (their reduced forms are NADH and $FADH_2$) from all three energy sources are funneled into the electron-transport chain (respiratory chain) where their excess electrons are transferred to oxygen in a stepwise procedure yielding ATP. In fact, the electron-transport chain is the site of production of the lion's share of ATP produced in the breakdown of food for energy. (Fig. 1.4) (32 out of 36 molecules of ATP produced per molecule of glucose; 17 out of 18 of the ATP molecules produced per fat breakdown step). This is the major role of oxygen in the biosystem -- to be reduced in the electron-transport chain with concomitant ATP production[21].

In the absence of oxygen, the workings of the chain back up. First all available cytochromes are reduced, then all of the flavins, finally all of the pyridine nucleotides are trapped in the reduced form, bearing electrons supplied by the Krebs and fatty-acid cycles and having no receptor molecule to pull the electrons off again. Certain poisons, such as the barbiturates, cyanide, and antimycin A can also block the chain in various sites, causing total reduction of all molecules before the block, and oxidation of all molecules following the block.

The workings of the electron-transport chain were deduced by Chance[22]. One of his major forms of study was the absorption, and later the fluorescence, of the components of the chain. He noted that each of

Figure 1.4. The electron-transport (respiratory) chain, showing the role of the pyridine nucleotides (NAD) and the flavoproteins (FP), which contain FAD, in transport of electrons to oxygen. The energy levels of successive elements in the chain drop (the scale shows energy in volts of a half-cell potential), thus energy can be released as ATP at the steps indicated. Note the role of amytal (a barbiturate) antimycin A and cyanide in poisoning the chain. From [18].

the molecules had absorption and fluorescence spectra that differed between the oxidized and reduced state. When mitochondria (the intracellular sites of the electron-transport chain) were incubated with and without oxygen, with or without substrates (succinate, ADP, e.g.), with or without various toxins, the state of the chain could be assessed by observing these optical properties[23]. In particular, it was clear that hypoxia in the presence of normal substrates led to reduction of all of the components of the chain, whereas under conditions close to normal healthy in vivo states the components were nearly 100% oxidized. There is a monotonic and quantifiable relationship between fluorescence/absorption of these molecules — particularly NAD and FAD — and the oxygen supply. In fact, these measurements produce a measure of adequacy of oxygenation as much as level of oxygen, since under conditions of lesser substrate formation, the need for oxygen is reduced and that is reflected in the presence of less reduced NAD and FAD.

In recent years several researchers have been trying to use this relationship to make devices which can assess hypoxia in exposed rat heart[24], and exposed brain [25] with encouraging results. Of course, these experiments are possible only at surgery or in a perfused organ preparation; still it appears to be possible, for example, to locate the ischemia produced by an experimentally blocked coronary artery.

Given these results, it stands to reason that in any tissue which can be directly illuminated and visualized it may be possible to isolate the fluorescence from the electron-transport chain, and thus to directly measure the oxygenation in vivo, without heavily invasive procedures. The eye, which is, of course, basically transparent to the retina over the visible spectrum, is one such tissue (other tissues are those which can be seen

with fiber-optics, e.g. the lung during bronchoscopy; given the current hypothesis that infantile respiratory distress syndrome is a problem of mitochondrial metabolism, this idea is more than just casual). Wexler[26] tried looking at the NAD fluorescence in the rabbit eye, but the NAD excitation peaks at 360 nm, which is ultraviolet and poorly transmitted to the retina; it became necessary to remove the cornea and lens from the rabbits before making the observations. However, the FAD fluorescence has both its excitation and emission in the visible (Fig. 1.5). The excitation peak is approximately 460 nm, and the emission is at 520 nm -- wavelengths at which the ocular transmission exceeds 80%.

FAD fluorescence, if it can be measured quantitatively and sorted out from the other materials in the eye which absorb and fluoresce, could prove to be a noninvasive means of assessing the true degree of hypoxia in the retina.

1.4. Goals for Determination of Retinal Hypoxia

In the overall effort of the Ocular Physiology Laboratory at the Boston University School of Medicine, along with the Biomedical Engineering Center for Clinical Instrumentation at MIT and the Department of Ophthalmology at University Hospital, to develop noninvasive methods of assessing retinal hypoxia, there are a number of major goals:

1. First, to establish the feasibility of the in vivo use of retinal measurements of respiratory-chain enzyme states (see below under "Testing and Verification"), and to establish the link between these direct measurements and oxygen supply to the retina;

2. Foremost, to establish the role of hypoxia in DR, whether it is a
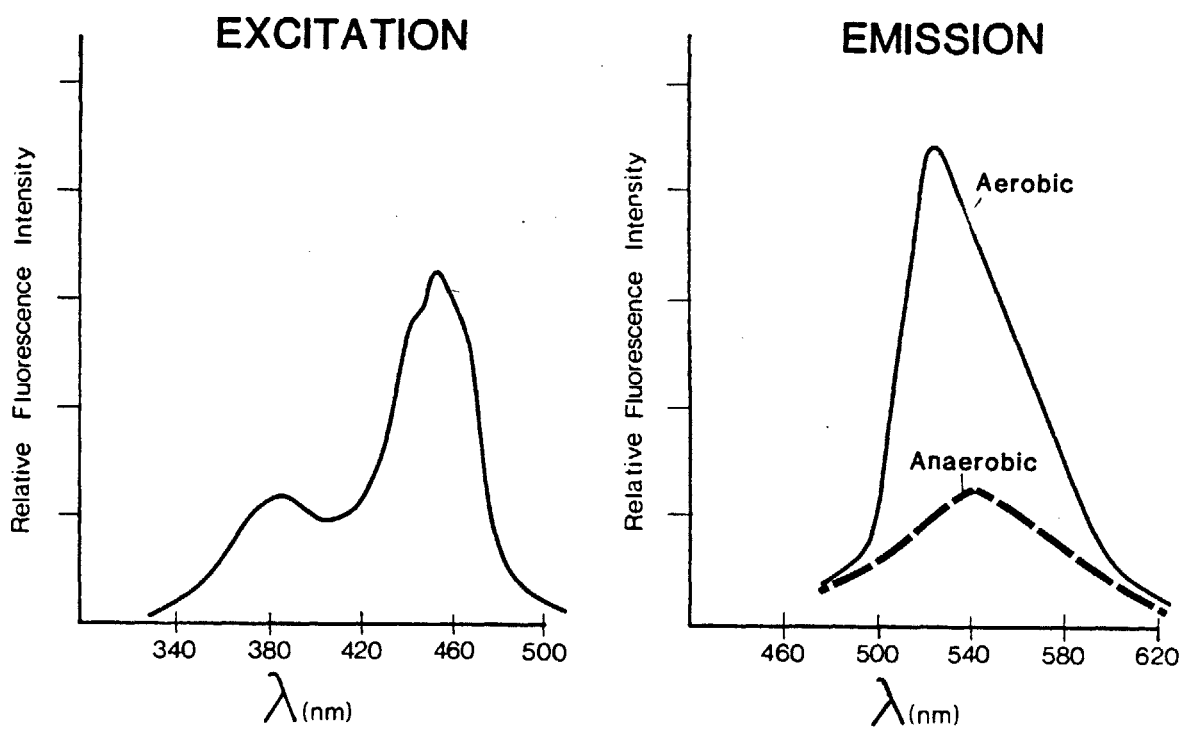
Figure 1.5. Excitation spectrum of FAD (emission at 520 nm) and emission spectrum of FAD (excitation at 436 nm) [20].

primary event, a secondary result from another initial event as described above, or not present at all;

3. To assess the progression of hypoxia in the diabetic retina over time; in particular, to correlate the changes of retinal oxygenation to the development of neovascularization in particular regions of the retina;

4. If development of local hypoxia appears to be predictive of later neovascularization in the same locale, to consider the feasibility of restricting the area of laser photocoagulation to those locales, and to determine the risks and benefits of this procedure instead of panretinal photocoagulation;

5. To determine the effect of photocoagulation itself on the oxygen status of the retina; in particular, to see whether overall tissue oxygen demands are better or more poorly met after photocoagulation compared with the same retina prior to treatment. This should be assessed for the local area of treatment as well as for the retina in general, as a function of distance from the photocoagulation spot.

CHAPTER 2

OCULAR REFLECTOMETRY AND FLUOROMETRY

We seek to develop a means of learning about oxygen status in the retinal tissue by measuring reflected and fluorescent light emerging from the eye. In order for this light to have any useful meaning to us, we have to have a fair idea of what we would expect to measure in normal and hypoxic states, and whether there are other factors which can affect our ability to distinguish between the two. The next sections address some questions about the predicted reflectometric and fluorometric properties of the eye under various conditions.

It would be nice to have a system to study where the full and final answer to the question of oxygenation could be reached by taking a single optical measurement, secure in the knowledge that that measurement carries all the information of the FAD oxidation state and no other information. Thus, if FAD in a cuvette has excitation at 430–460 nm and emits at 520–540 nm, then our device should be able to provide excitation in this former band, observe emission in the latter, and declare that all captured light is due to FAD fluorescence and provide a figure of the current redox state.

Unfortunately, the intact human eye is nowhere near that simple. The fundus must be illuminated and observed through a narrow pupil, at most 9 mm in diameter. Between the light source and the retina lie the cornea, aqueous humor, lens, and vitreous, each of which may alter the excitation light through absorption and scattering, and may cause emission in the form of reflectance and fluorescence on their own. Furthermore, the FAD redox state itself is not determinable with a single measurement because the total FAD concentration itself is a variable. While changes in FAD redox state over time in the same site might be revealed in a single measurement,

pencil beam

reflection
scatter
fluorescence

aqueous humor
cornea
tear layer
iris
lens
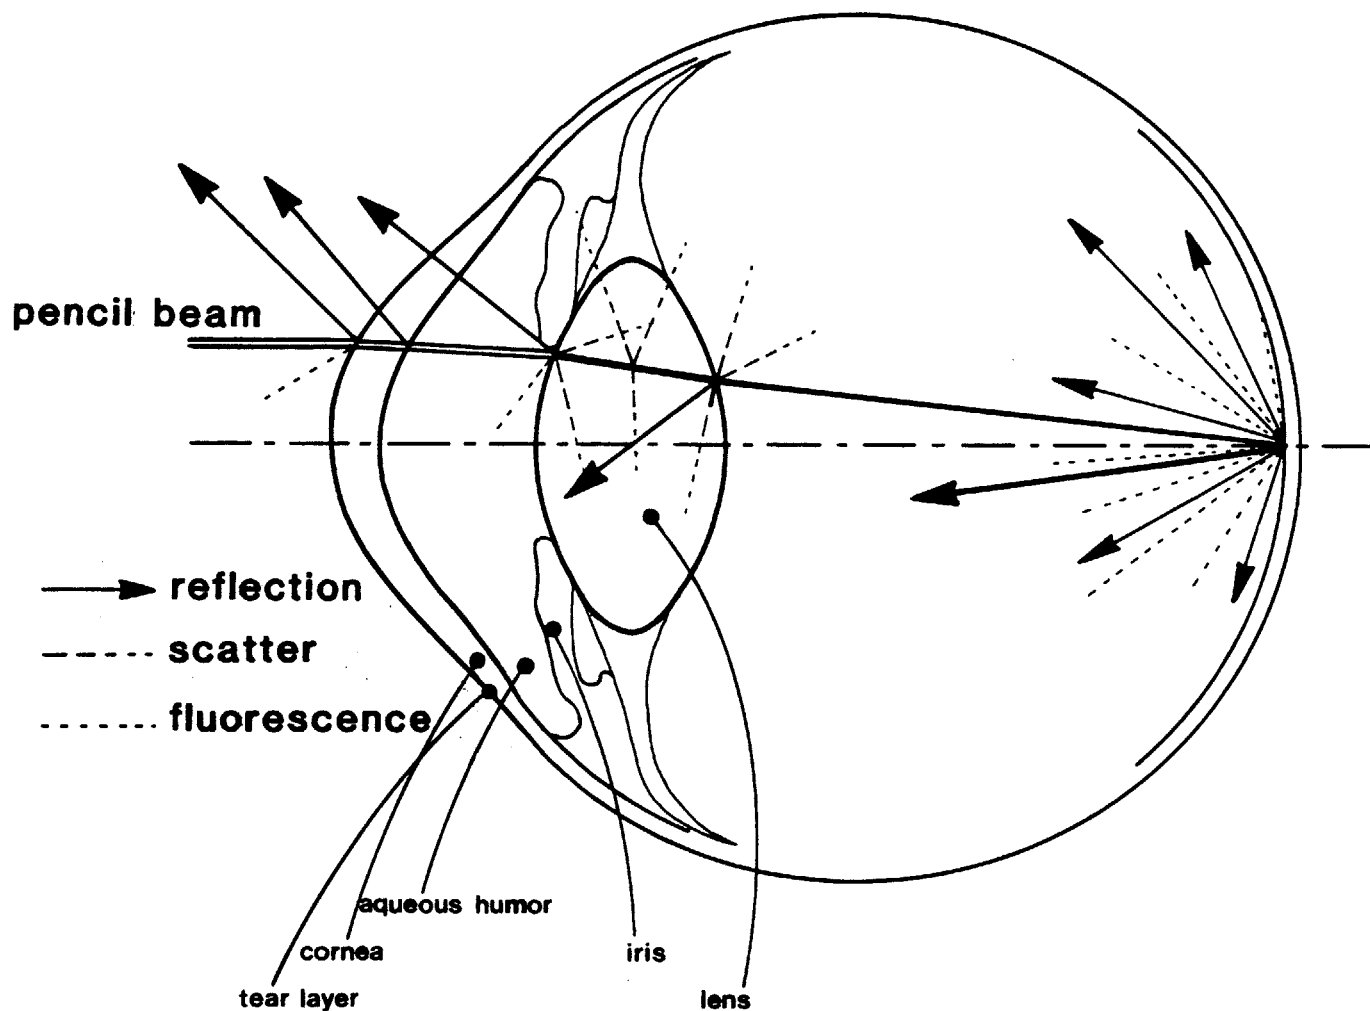
Figure 2.1. Structures in the eye and their possible effects on transmission of light to the retina and on the emitted light signal.

the absolute redox state is not known from one measurement (see below).

Clearly there are many possible ways for the retinal FAD signal to be altered. Which of these is potentially significant, and how to sort through them to the FAD itself, is the subject of the following discussion.

2.1. Retinal structure

A schematic drawing of the retinal cell layers is shown in Fig. 2.2 [7].
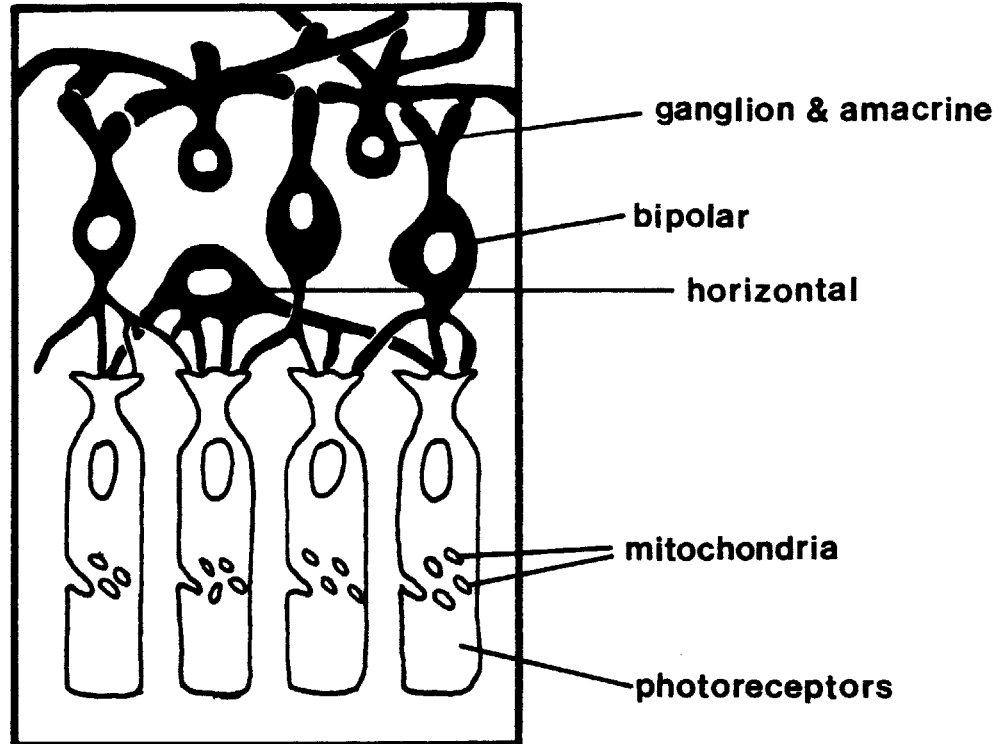
Figure 2.2. The cell layers in the human retina. Note the location of the photoreceptors' mitochondria. There are also mitochondria in the inner retinal cells. After Young[7].

The photoreceptors (rods and cones) lie most posterior in the retina, and the light-sensitive photopigment, rhodopsin, is found in the most posterior segment of the photoreceptor cell. In the inner portion of the photoreceptor is found a large collection of mitochondria, which provide energy for the regeneration of cis-rhodopsin after light has changed it to the all-trans-form. These mitochondria also provide energy for the production of new rhodopsin as needed[7].

Anteriorly to the photoreceptor layer are layers of other cells -- bipolar, amacrine, horizontal, and ganglion cells which do some of the

early processing and which transmit visual information to the brain. These cells also contain mitochondria which provide energy for their functions.

As seen in Figure 2.3 , there are two major blood supplies to the retina: the retinal circulation, which supplies the inner retina primarily, and the choroid, which supplies primarily the photoreceptors[27].

Figure 2.3. The retinal (above) and choroidal circulation of the retina seen in cross section[27].

The choroid is a fixed-flow, high-flow system which provides a base-line level of oxygen for the outer retina, whose large metabolic demands rapidly use it up[27]. The retinal circulation, which includes the vessels commonly seen in fundus photographs, can be regulated to change the blood flow and oxygen delivery to the tissue.

Behind the retina is the retinal pigment epithelium, which has high absorbance for visible light wavelengths (Fig. 2.4).

In non-albinos, the transmission of visible light through the retina and through the RPE is negligible.

FAD is found almost entirely in the mitochondria, attached to the membranes[23]. Thus, the mitochondrial zone of the photoreceptors, and the

Figure 2.4. Spectral absorption of melanin, found in the pigment epithelium.

connecting-cell layers, are the sites where FAD may be detected. Other compounds found in important quantities include blood products, particularly hemoglobin, in the vascular layers and also intermixed with the cell layers; rhodopsin itself, which is found in the outer segments of the receptors; xanthophyll, a yellow pigment found in the macula (which has no overlying blood vessels). We need to consider these compounds and others in the following discussion.

2.2. Reflectometry of the eye

Reflection from the eye can arise from various sources:

1. Reflections can occur at any transition between two areas of different refractive index; the larger the difference in index, the more significant the reflection. Thus the interface between the successive chambers -- air, cornea, aqueous, lens, vitreous -- are sites of reflection. There may also be transitions between the different layers of the

cornea and lens, and transitions due to blood vessel walls, which give off lesser reflections. The corneal and lens surfaces in particular give rise to specular (Purkinje) reflections which form a source of interference to our study.

2. Light that passes through the retina is reflected by the pigment epithelium and sclera (if it reaches that layer), then back through all of the structures heading anteriorly. It is this light that is reflected into the physician's ophthalmoscope, and it is this light that forms the basis of fundus reflectometry. The light is altered by absorption and scattering in both directions.

3. Backscatter from surfaces such as erythrocyte walls in the anterior blood vessels can theoretically give rise to reflections. This light will have a variable effect on fundus reflectometry depending on how far through the vascular layer it traveled before reflection. Bakker[28] measured this effect and found that it contributed only 1% to the total returning light for a vascular layer of 100 μm thickness. Our calculations of reflection, therefore, are similar in form to a calculation of transmission through the eye and back out again, subject to certain constraints caused by the limited angles of incidence and reflection.

2.3. Building up the model

We will build up the model of the eye's reflection and fluorescence from the simple to the more complex, starting with a cuvette containing only FAD and building up the other components as we go along.

2.3.1. FAD reflectometry. Consider a system such as shown in Figure
2.5
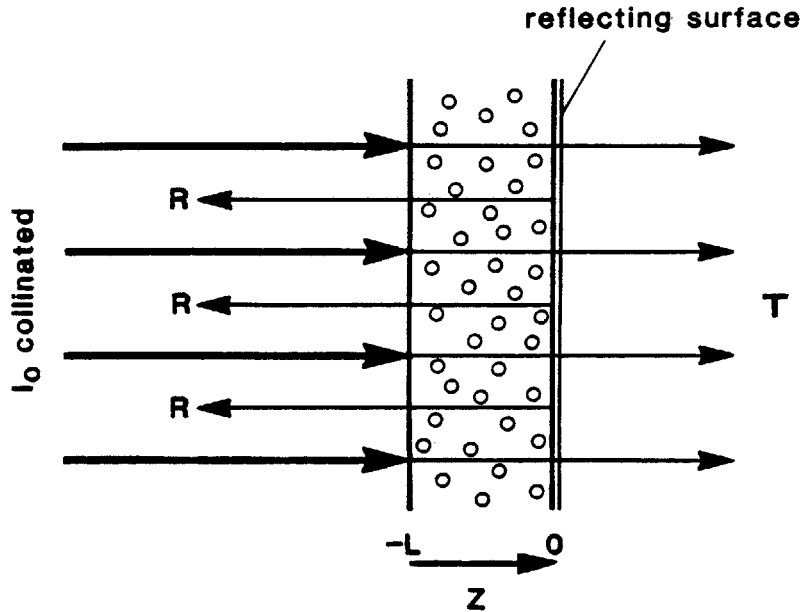
Figure 2.5. System containing light incident on FAD alone.

In each incremental thickness dz there is absorption, scatter, and
fluorescence taking place. The thickness is small enough so that only one
such interaction will occur within the volume element. When collimated
light of intensity I hits a unit area of the layer, the contribution to the
continued transmission T is (we are considering only the incident, not the
reflected, path)

(1) $dT = -\alpha T\ dz - \rho T\ dz$

where α is the fraction of incident light T that is absorbed in this layer,
and ρ is the fraction of light that is scattered[29].

The probability that a certain amount of light $\alpha T\ dz$ is absorbed is
the product of the number of photons incident (i.e., T) and the number of
targets to be hit and absorb light (concentration, C, times unit area --
usually 1 cm^2 -- times dz) times a normalized efficiency of absorption (the
molar extinction coefficient ε) which is the chance of an absorptive

interaction in such an environment at a given wavelength. Thus, for absorption alone, the change in transmittance is

(2) $dT = -\varepsilon CT dz$

which when solved for a medium of depth (z-dimension) $\underline{L}$ gives us

(3) $T = I_0 e^{-\varepsilon CL}$

the familiar Beer-Lambert Law.

The scattering factor depends on the size of the molecules involved, ranging from the Rayleigh equation for scatterers much smaller than the wavelength[30] to the more complex Mie model for large scatterers. For sufficiently small concentrations, however, ρ is a constant $\underline{a}$ (dependent on λ but not on C) times $\underline{C}$. (For hemoglobin, a molecule of the same order of size as the flavoproteins, the concentration criterion is that C is much less than 5 millimolar; thus inside blood vessels for which hemoglobin is normally 1-4 millimolar this does not hold. For the flavoproteins which are typically 30 micromolar the criterion holds). The scattering factor varies inversely with the wavelength (by a factor λ^{-4} for Rayleigh scatter) so that the scatter is more prominent in the blue end of the visible spectrum[30]. At any rate, ρ is independent of T, being dependent only on C and λ[29]. Thus the amount of light lost from the direct transmission is

(4) $dT = -\varepsilon CT dz - \rho(C,\lambda)T dz$

By the time we go through a layer (depth L) of FAD, assuming all of the scattered light is lost, and reflect specularly off a surface (perpendicular to $\underline{z}$) of reflectivity $\underline{R}$ back through the layer, we find a total reflectance intensity of

(5) $I = I_0 R \, e^{-\varepsilon'(2L)}$

(5) $I = I_0 R \, e^{-\varepsilon'(2L)}$

where $\varepsilon' = (\ \varepsilon C + \rho(C)\ ) = (\varepsilon + a)C$ for sufficiently small concentrations.

In fact, all the scatter is not lost; some of the scattered light bounces back and may strike the detector; the rest continues at oblique angles through the medium and can still contribute to the reflectance (or fluorescence, if it excites FAD molecules and we are set up to detect it). In a scanning system where all of the emitted light at a given moment is detected and the image is created by time-separation, scatter can contribute significantly to blurring if it has its effect away from the desired spot in the medium.

2.3.2. Angle effects. With some malice aforethought, let us consider a few optical variants from this simple model. In particular, (1) the excitation is not parallel but converges in a cone with plane angle θ, i.e. solid angle $2\pi(1 - \cos \theta)$ (θ is the angle between the midline and the boundary of any plane projection of the cone); (2) the reflectance is not specular but diffuse; (3) we capture reflected light over a different cone which shares its apex with the first cone at the reflecting surface.

The effect of a non-collimated excitation is that the path length of the light is longer by $(1/\cos \theta)$ for any ray coming in at angle θ to the normal to the medium (see figure). Furthermore, we can no longer think of the beam as going through a unit area, for the area decreases as we converge to a point at some focus plane. We will assume that we focus the beam on the last layer of FAD, just before the reflector. For a cone of solid angle $2\pi (1 - \cos \theta)$ the absorption through medium of depth $\underline{dz}$ is

(6) $dT = -T(\varepsilon + a)C \dfrac{dz}{\cos\varphi}$

for each hollow cone; for the whole volume element of normal distance dz,

Figure 2.6. Effects of noncollimated excitation, diffuse reflectance, and capture with a separate exit pupil

$$(7) \quad = \frac{-T(\varepsilon+a)C}{2\pi(1-\cos\theta)} \int_0^\theta \frac{dz}{\cos\varphi} \, 2\pi\sin\varphi \, d\varphi$$

$$(8) \quad = \frac{-T2\pi(\varepsilon+a)C \, dz}{2\pi(1-\cos\theta)} \ln(\sec\theta) = -T \, k(\theta) \, (\varepsilon+a)C \, dz$$

Thus, compared to the case of collimated excitation, our attenuation $-dT$ is multiplied by the normalized area times a factor

$$(9) \quad k = \frac{\ln(\cos\theta)}{\cos\theta - 1}$$

This can be considered to be a constant which may be combined into ($\varepsilon+a$)

(As θ approaches 90 degrees, In order to keep to our criterion of no more than one event per light ray passing through dz, we have to keep scaling

down dz by multiplication by cos θ. Then our integration over dz to find the total transmission goes from 0 to L/cos θ, rather than 1).

For example, a cone of half-angle 12 degrees effectively multiples ε+a by 1.011. (Values of ε and a taken from the literature should also be multiplied by 2.303 as they are usually based on common logarithms). Thus total transmission for a 12 degree half-angle cone would be (one direction only)

$$(10) \quad \ln \frac{T}{I_0} = \int_{-L}^{0} -1.011 \ (\varepsilon+a)C \ dz$$

$$(11) \quad T = I_0 \ e^{-1.011 \ (\varepsilon+a)C \ L}$$
$$= I_0 e^{-BL}$$

where

$$(12) \quad B = k(\theta) \ (\varepsilon+a)C$$

The effect of a diffuse reflector is similar in that it represents a divergent rather than collimated beam. It is slightly different because instead of the beam having equal power in all angles, a Lambertian reflector provides a power which varies as the cosine of the angle. Thus the equivalent to Eq. 7 is

$$(13) \quad dT = \frac{-T(\varepsilon+a)C}{\pi \ \sin^2\theta} \ dz \int_0^{\theta} 2\pi\sin\phi \ d\phi$$

$$(14) \quad = \frac{-2T(\varepsilon+a)C \ (1-\cos \ \theta)}{\sin^2\theta} \ dz$$

(The normalizing denominator term $\sin^2\theta$ is the total light in the cone,

i.e. $\pi \sin^2\theta = \int_0^{\theta} 2\pi\sin\phi \ \cos\phi \ d\phi$).

If the full diffuse reflectance would be distributed over a hemisphere, ($\theta = \pi/2$) the net result would be to multiply ε + a by 2. (We are

ignoring for the moment that the incident beam is still not collimated, so that the angle distribution is slightly skewed.) The third consideration, that we pick up the emission over a limited cone as well, implies that θ above has a limitation, and thus the total reflectance would be multiplied by the reduced value of this integral. If the second cone is additionally skewed at an angle to the first, then to a first approximation we multiply the integrand by the additional factor $\cos\chi$, χ being the angle between the centerlines of the excitation and emission cones.

Appendix A contains the full derivation which accounts for a limited emission cone with half-angle ϕ skewed at a certain angle χ from an excitation cone of half-angle θ and exhibiting properties of diffuse reflectance. The fraction of light striking the base (z=0) which heads toward the exit pupil is (R_0 is the intensity at the base):

$$(15) \quad \frac{R}{R_0} = \frac{\sin^2\phi \, \sin^2\theta \, \cos\chi}{2 \, (1 - \cos\theta)}$$

so the full expression for reflected light which is absorbed on the way in and on the way out of the retina is

$$(16) \quad R = I_0 e^{-(B_\theta + B_\phi)L} \frac{\sin^2\phi \, \sin^2\theta \, \cos\chi}{2 \, (1 - \cos\theta)}$$

where the subscripts on $\underline{B}$ denote which angle is to be used in the calculation of $\underline{B}$.

2.3.3. <u>Fluorescence</u>. Fluorescence is a process which is proportional to the <u>absorption</u>, rather than the transmission, through a medium. In general fluorescence by a given compound is the product of the light absorbed by that compound times the quantum efficiency η. Thus for an excitation beam of the type shown in Equation 8, the fluorescence at each layer dz is

$$(17) \quad dF = \eta \, T(z) \, k(\theta) \, \varepsilon C \, dz$$

-32-

For a layer of FAD extending from −L to 0 in depth (i.e. focused at the back of the layer) the total fluorescence generated from the excitation beam is

$$(18) \quad F = \int_{z=-L}^{z=0} dF$$

$$(19) \quad = \eta \; k(\theta) \; \varepsilon C \int_{-L}^{0} I_0 e^{-BL} \; e^{-Bz} \; dz$$

$$(20) \quad = I_0 \eta \; k(\theta) \; \varepsilon C \; e^{-BL} \int_{-L}^{0} e^{-Bz} \; dz$$

where

$$(21) \quad B = k(\theta) \; (\varepsilon + a)C$$

The term from I_0 on in the integral in Equation 19 is the transmission $T(z)$ from −L to z. The final result turns out to be

$$(22) \quad F = I_0 \eta \; k(\theta) \; \varepsilon C \; e^{-BL} \frac{e^{BL} - 1}{B}$$

$$(23) \quad = \eta I_0 \frac{\varepsilon}{\varepsilon + a} \; (1 - e^{-BL})$$

The only variable in this equation for study of tissue like the retinal layer is the concentration, which is a linear factor in B. Note that, for wavelengths where ε is not too small (i.e. wavelengths where absorption and hence fluorescence are appreciable), the fluorescence does not go up linearly with concentration; for small values of B it does, but as B increases (because of increased concentration) the fluorescence peaks and further concentration will increase the signal only minimally. This is due to increased absorption of the light which would otherwise be transmitted to later layers for fluorescence. Because of the exponential dependence of transmission on (ε + a) and C, we can define a mean free path, being the

−33−

average depth of tissue to which a photon may be exposed. Quistorff and Chance noted experimentally[31] that the penetration of 436 nm light into frozen liver sections which had been emptied of blood was about 100 μm, using collimated light. The concentration of FAD in the retina may be expected to be of a similar order of magnitude to the liver. (See chapter on monkey experiments).

Figure 2.7. Plot of $(1 - e^{-BL})$ vs. B. B is linear in C and ε+a.

As mentioned before, this is the isotropic fluorescence generated by the excitation path through the sample. If we extract information only over a limited exit cone, we must multiply the result above by the solid angle of that cone divided by 4π. Also, fluorescence may be generated in the return path (from the reflected light); it will be reduced in intensity due to the non-perfect reflectivity R and due to the reduced amount of light available for reflection, the light having been absorbed and scattered once already by the sample from z= -L to 0. For the human eye

-34-

reflectivity is approximately 0.01, thus fluorescence on the reflected pathway from the same sample of FAD in the retina will be small in comparison to the direct fluorescence.

The fluorescent light itself may be absorbed or scattered on its way back through the sample to the exit pupil. We can change Equation 17 to consider the small exit cone and the absorbance of the fluorescing light by figuring in the solid angle limitation and combining with Eq. 11 for the back transmission:

$$(24) \quad dF = \eta \ T(z) \ k(\theta) \ \varepsilon C \ dz \ \tfrac{1}{2}(1-\cos \varphi) \bullet$$
$$e^{-k(\phi) \ (\varepsilon_2+a_2)C \ (z + L)}$$

Here φ is the half angle of the output cone from any fluorescent absorber (for a narrow and thin layer of FAD, relative to the distance to the exit pupil, the angles will be extremely similar for any fluorescence in the layer) and the use of subscripts on ε and a take account of the different wavelength of the fluorescent light.

The result for the total fluorescence emission is:

$$(25) \quad F = I_0 \eta \ k(\theta) \varepsilon C \ \tfrac{1}{2}(1-\cos\varphi) \ e^{-(B+B_2)L} \int_{-L}^{0} e^{-(B+B_2)} \, dz$$

where

$$(26) \quad B_2 = k(\phi) \ (\varepsilon_2+a_2)C$$

By comparison with Eq. 20 through 23 this becomes

$$(27) \quad F = \eta I_0 \ \frac{\varepsilon}{\varepsilon+a} \ \tfrac{1}{2}(1-\cos \varphi) \ \frac{B}{B+B_2} \ (1 - e^{-(B+B_2)L})$$

if the whole surface of excitation at all z can send fluorescent rays to the exit pupil. It is also worth remembering that if the exponent is small compared to 1 the expression becomes

(28) $F = \eta I_0 \frac{\varepsilon}{\varepsilon+a} \frac{1}{2}(1-\cos \varphi)$ B L

that is, fluorescence is proportional to concentration for low concentrations.

2.3.4. FAD and FADH. We now introduce to the medium another compound — the presence of FADH, into which FAD readily converts in a reducing environment. Our determination of retinal oxygenation is derived directly from the FAD redox ratio, which is defined as

(29) $\sigma = \text{redox ratio} = \dfrac{[FAD]}{[FAD] + [FADH]}$

Over a short period of time the total concentration

(30) $[FAD]_{total} = [FAD] + [FADH] = \text{constant}.$

Since these two compounds are readily interchangeable, it is fair to assume that they exist arranged randomly in the same compartment; that there is no singular area of oxidized or reduced flavin in the local environment (except as occurs as a random occurrence).

Wodick and Lubbers[32] have addressed the problem of multicomponent systems by the concept shown in Figure 2.8(B). They speak of "additive" and "multiplicative" components: additive components are essentially side-by-side facing the light, and the interaction of light with one will not affect the probabilities of interaction of light with the other; the output light $I(\lambda)$ is the sum of the interactions of additive components I_1, I_2, etc. The factors ϕ_i represent the relative amount of space occupied by light paths of type I_i; the ϕ_i's sum to 1.

Multiplicative combination occurs when one component is in front of the other, closer to the light source and potentially blocking light from the second component. In this case, especially if there is fairly high concentrations of the components, the contribution to I_i is the
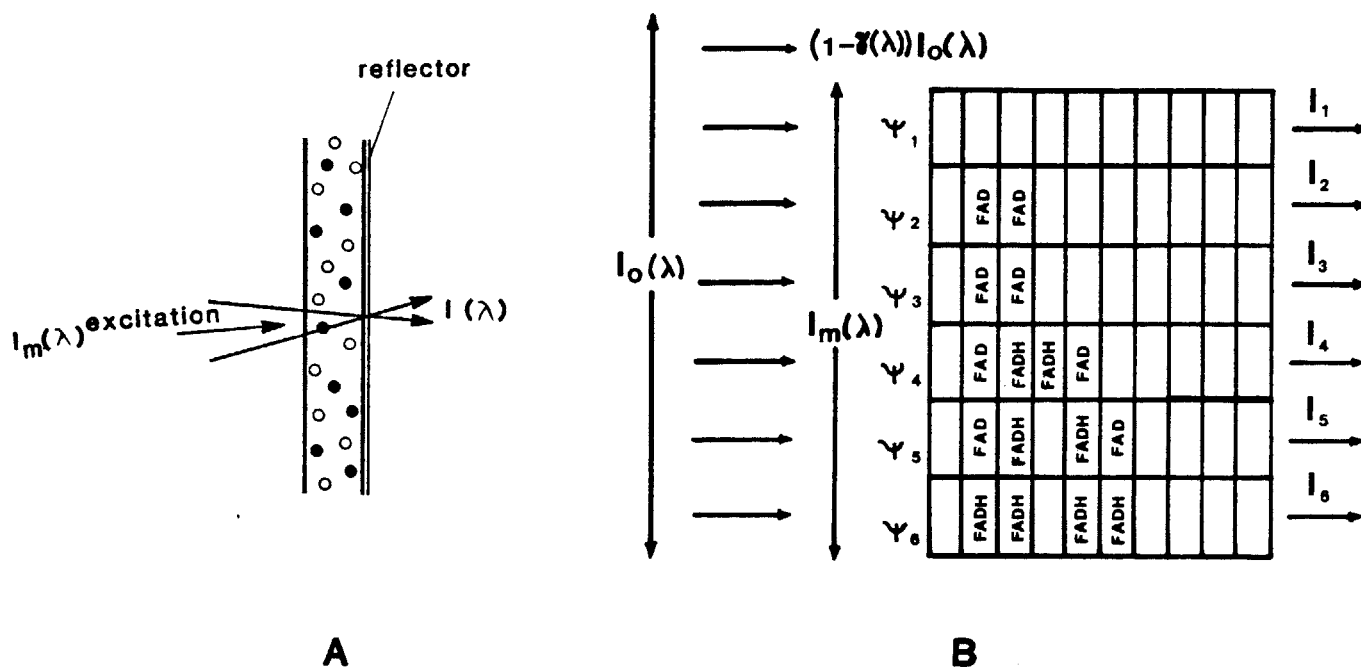
Figure 2.8. (A) The system with both FAD (open circles) and FADH (closed circles) randomly arranged. (B) Wodick's model of "additive" and "multiplicative" compartments[32].

transmission through the first component, which is then available to the second component and thus is multiplied by its transmission to give the total I_i. The fact that $I(\lambda)$ carries a wavelength dependence acknowledges, as we did above, that transmission is different at different wavelengths.

Between the incident light $I_o(\lambda)$ and the available light for interaction $I_m(\lambda) = \gamma(\lambda) \, I_o(\lambda)$ there are other interactions and absorbances, unspecified, which remove a fraction $(1 - \gamma(\lambda))$ of the light.

A more precise treatment of the FAD/FADH problem can be found by allowing more than one component in each of the cells, and bringing the depth of the cells down to the infinitesimal (i.e., dz). Let us assign C_T to mean the total concentration of FAD and FADH, which we will hold to be constant for the time of an experiment. If the redox state is given by σ, then

(31) $[FAD] = \sigma C_T$ and $[FADH] = (1-\sigma)C_T$

We will assign $\varepsilon_o(\lambda)$ and $\varepsilon_r(\lambda)$ to be the molar absorption coefficients of oxidized (FAD) and reduced (FADH) flavins at a given wavelength. Similarly with $a_o(\lambda)$ and $a_r(\lambda)$ for scattering coefficients, assuming we can approximate $\rho(C) = aC$.

Now in the first layer dz, we have incident light $I_m(\lambda)$, and as we derived above (Eq. 4), the extinction in the layer is

(32) $\dfrac{dT}{dz} = -\sigma TC_T (\varepsilon_o(\lambda) + a_o(\lambda)) - (1-\sigma)TC_T (\varepsilon_r(\lambda) + a_r(\lambda))$

We can create parameters:

(33) $\varepsilon_T(\sigma,\lambda) = \sigma\varepsilon_o + (1-\sigma)\varepsilon_r$

(34) $a_T(\sigma,\lambda) = \sigma a_o + (1-\sigma)a_r$

and then, assuming that the relative concentration of the two stays the same throughout the sample on a macroscopic scale, we can use ε_T, a_T, and C_T in all of our previous equations for reflectance and fluorescence, i.e.

(35) $\dfrac{dT}{dz} = -\varepsilon_T TC_T\, dz - a_T TC_T$

In the model above, even if we stick to one component per compartment, there is FAD only in the first compartment of row 1, with a width $\phi_1 = \sigma$, and FADH only in the first compartment of row 2, $\phi_2 = (1-\sigma)$; behind them lie statistically identical cohorts of more molecules to follow. The total transmittance through the medium in this fashion, after Equation 11, is

(36) $I(\lambda) = T = I_m(\lambda)e^{-k(\theta)(\varepsilon_T + a_T)C_T L}$

(37) $ = I_m(\lambda)e^{-B_T L}$

B_T has the same form as B but substitutes ε_T, a_T, and C_T. For small concentrations,

-38-

(38) $I(\lambda) = I_m(\lambda)(1 - B_T(C_T,\sigma,\lambda) L)$

which is directly (though negatively) proportional to

(39) $C_T (\sigma\varepsilon_o(\lambda) + (1-\sigma)\varepsilon_r(\lambda) + \sigma a_o(\lambda) + (1-\sigma)a_r(\lambda))$

In the biosystems we are considering, C_T and σ are unknown and the ε's and a's are known, thus the measurement I(lambda) is a measurement with two variables. One problem with this linear approximation is that when B_T is small then the term in parentheses in Equation 38 is very close to 1. If we tried to take two measurements at different wavelengths, solve each for $B_T(\lambda)$, and from those two values solve for the two unknowns C_T and σ, we would be solving equations of the form

(40) $B_T(\lambda_1) = (1 - \dfrac{I(\lambda_1)}{I_m(\lambda_1)})/L$

which involves the difference of two very close numbers, and thus a noise problem.

Observing fluorescence, with a converging beam, we have, after Equation 17,

(41) $dF = T(z,\varepsilon,a,C) \; k(\theta) \; \varepsilon C \; \eta \; dz$

The appropriate transformations for us to use are ε_T, C_T and a_T as above for determining the transmission $T(z)$, and for the fluorescent absorption / quantum emission $\eta\varepsilon C$ we use

(42) $\sigma \; C_T\varepsilon_o\eta_o + (1-\sigma) \; C_T\varepsilon_r\eta_r$

This shows that the oxidized form absorbs at a rate ε_o and emits fluorescence with quantum efficiency η_o; similarly for the reduced form. In the layer dz the relative frequency of such interactions are σ and $1-\sigma$.

The same derivation continues and we reach, like Eq. 25,

$$(43) \quad F = I_m \eta \, k(\theta) \varepsilon C \frac{1}{2}(1-\cos\phi) \, e^{-(B_T+B_{T2})L} \int_{-L}^{0} e^{-(B_T+B_{T2})z} \, dz$$

In this equation B_{T2} is the value for B_T at the fluorescence emission wavelength. Performing the substitution for $\eta \varepsilon C$ and solving gives us

$$(44) \quad F = \frac{B_T}{B_T+B_{T2}} \, I_m \, \frac{\sigma\varepsilon_o\eta_o + (1-\sigma)\varepsilon_r\eta_r}{\varepsilon_T + a_T} \, (1 - e^{-(B_T+B_{T2})L}) \cdot$$
$$\frac{1}{2}(1 - \cos\phi)$$

For small concentration C_T, the quantity $(B_T + B_{T2})L$ is insignificant compared to 1. In this case we obtain :

$$(45) \quad F = \frac{I_m}{2} \, \frac{B_T}{\varepsilon_T + a_T} \, L \, (\sigma\varepsilon_o\eta_o + (1-\sigma)\varepsilon_r\eta_r) \, (1 - \cos\phi)$$

$$(46) \quad = I_m \, K \, C_T(\sigma\varepsilon_o\eta_o + (1-\sigma)\varepsilon_r\eta_r)$$

where

$$(47) \quad K = \frac{L}{\varepsilon_T + a_T} \, \frac{B_T}{C_T} \, \frac{1 - \cos\phi}{2}$$

($B_T/(C_T(\varepsilon_T + a_T))$ is constant at a given wavelength, length L and excitation and emission cones). Finally, let

$$(48) \quad K_o(\lambda) = K \, \varepsilon_o\eta_o$$

$$(49) \quad K_r(\lambda) = K \, \varepsilon_r\eta_r$$

Then we find that for small concentrations,

$$(50) \quad F(\lambda) \quad = \quad I_m(\lambda) \quad (K_o(\lambda) \quad [FAD] \quad + \quad K_r(\lambda) \quad [FADH])$$

This fluorescence measurement does not suffer from the "small difference of large numbers" problem mentioned above. If we measure F at two different wavelengths, we can solve for [FAD] and [FADH] and hence, for the redox state, as Figure 2.9 shows.

NORMALIZATION

$$I_{\lambda 1} = K_{ox_{\lambda 1}} \left[Ox \right] + K_{red_{\lambda 1}} \left[Red \right]$$

$$I_{\lambda 2} = K_{ox_{\lambda 2}} \left[Ox \right] + K_{red_{\lambda 2}} \left[Red \right]$$

$$\left[Ox \right] = \frac{K_{red_{\lambda 2}} I_{\lambda 1} - K_{red_{\lambda 1}} I_{\lambda 2}}{K_{red_{\lambda 2}} K_{ox_{\lambda 1}} - K_{red_{\lambda 1}} K_{ox_{\lambda 2}}}$$

Figure 2.9. Use of two wavelength measurements to solve for the concentrations of oxidized and reduced flavins.

Studying this equation, one notes that it no longer has terms in ε_2 or a_2.

Thus we see that our condition, $(B_T + B_{T2})L \ll 1$, has given us the fluorescence measurement when absorption is insignificant in either the forward or backward path.

2.3.5. Hemoglobin and other absorbers. Now that we can find the redox ratio in a sample of FAD and FADH mixed together, we are faced with the next problem -- that of the other chemical components found in the retina.

Among the most important are hemoglobin, which is found both anterior to the retina and within its layers in diffused form; rhodopsin, which is the visual pigment and which is found in the outer retina; xanthophyll, which is the chief pigment of the macular area, where there are no major blood vessels. Our model now looks (in Wodick's form) as follows:

Figure 2.10. Schematic representation of the relation between various molecules in the retina.

A complicating factor of hemoglobin is that its light absorption pro- perties differ markedly in oxygenated and deoxygenated environments; just as FAD changes state in tissue hypoxia, so hemoglobin changes state (transforming from oxyhemoglobin to deoxyhemoglobin). For the low- concentration, "additive" part of the hemoglobin which lies within the

-42-

retina, this is a minor problem only in that there are two more components (Hb and HbO_2) to deal with rather than one. The "multiplicative" component, due to the overlying blood vessels, is more of a problem, especially since the intravascular pO_2 is not the same as that of the tissues where the FAD lies.

The multiplicative (i.e. intravascular) hemoglobin is not uniformly distributed across the surface of the retina; there are large vessels and spaces between them. The smallest blood vessels may be smaller than the size of our scanning spot, although not so finely distributed as to be considered uniform. These small vessels contribute to image _clutter_ and will cause erroneous readings if we do not correct for them.

The transmission of a given beam through a certain area of the vascular layer (ranging from depth $-L_v$ to $-L$) over an area element dA (measured at $z = -L_v$; remember, these are convergent beams) in which the hemoglobin concentration _is_ uniform is

$$(51) \quad dT = I_m(\lambda) \, e^{-B_{HbT}(L_v - L)} \, dA$$

where $\quad B_{HbT} = k(\theta) \, (\varepsilon_{HbT} C_{HbT} + \rho_{HbT}(C))$

and ε_{HbT} has the same form as ε_T except the two forms are oxy- and deoxy-yhemoglobin. Recall that for the intravascular space ρ is not linear with concentration (in two different papers the dependence of ρ on C is given as $\rho = aC \exp(-bC)$ and $\rho = aC (1-C)$; see [29,33]). There is a dependence of B_{HbT} on the relative proportions of oxy- and deoxyhemoglobin, and hence, on the intravascular pO_2. We cannot use a linearization approximation because the concentrations of hemoglobin in the vessels is too high. Further, different classes of vessels (arteries and veins, for instance) which may both be in our excitation cone, have different pO_2. The best we can do is to say that the light supplied to the FAD/FADH/Hb layer is

$$(52) \quad T = \frac{1}{\pi L_V{}^2 \sin^2\theta} \iint I_m(\lambda) \, e^{-B_{HbT}(dA)} \, (L_V - L) \, dA$$

(integrating over the area illuminated by the incident beam) and we would have to know the statistics of vessel thickness and pO_2 to get the complete answer.

If we choose an excitation cone wide enough so that the blood layer appears to have uniform hemoglobin content and oxygenation compared to other cones (in a scanning instrument), then the integral above will be the same for all regions scanned, and will depend only on the mean intravascular pO_2. Furthermore, if we can find a wavelength (called an _isobestic_) at which $\varepsilon_{oxyHb} = \varepsilon_{deoxyHb}$ and the scattering factors are also equivalent, then the integral above is a constant. There are a number of isobestic points for ε in the hemoglobin absorption spectrum -- or isobestic regions, which are the local neighborhoods of isobestic points (see Fig. 2.11) and observations taken with excitation at these wavelengths and a cone following these constraints allows us to measure the composition of the FAD layer with only a constant attenuation factor preceding it.

Dealing with the hemoglobin found within the FAD layer is an easier problem, especially if we can assume that the hemoglobin is either randomly distributed or is in capillaries much smaller than the size of the cone at z for a large fraction of the depth of the layer, so that different scanning cones are uniform. Then we treat it as two more absorbers and

$$(53) \quad \frac{dT}{dz} = - \sigma TC_T \, (\varepsilon_o(\lambda) + a_o(\lambda)) - (1-\sigma)TC_T \, (\varepsilon_r(\lambda) + a_r(\lambda))$$

$$- \sigma_{Hb} TC_{HbT} \, (\varepsilon_{oxyHb}(\lambda) + a_{oxyHb}(\lambda))$$

$$- (1-\sigma_{Hb})TC_{HbT} \, (\varepsilon_{deoxyHb}(\lambda) + a_{deoxyHb}(\lambda))$$

(Here where the concentrations are small we may use the linear dependence of ρ on C.) The same model applies if we set

Figure 2.11. Absorption spectrum of oxyhemoglobin and deoxyhemoglobin.

$$(54) \quad \varepsilon_{TT} = \frac{[FAD] \, \varepsilon_o + [FADH] \, \varepsilon_r + [HbO] \, \varepsilon_{oxyHb} + [HbD] \, \varepsilon_{deoxyHb}}{C_T + C_{HbT}}$$

and do the same for $\underline{a}$ and use $C_{TT} = (C_T + C_{HbT})$ in place of C_T. We will

then have an expression just like Eq. 37 once more, only this time B_T has

four variables. For the fluorescence, even though only FAD and FADH

fluoresce in the appropriate band (and hoping we can select an isobestic

wavelength of hemoglobin for our fluorescence) all four variables still

appear in the equation.

It would seem that for every new absorbing or fluorescing component we

must make an additional measurement in order to extract the FAD redox ratio

from the signal. We shall first add the lens and cornea onto the system

and then return here in the next chapter and try to outline ways to save equipment, energy and measurement error by cutting down on the required number of measurements.

Rhodopsin should be mentioned as a special case of a substance which falls behind the FAD layer. Clearly the absorption of rhodopsin will not determine how much light is available to excite the FAD molecules. However, in a pure reflectance measurement, the light will pass through the rhodopsin layer before it reflects; thus it must be treated much as the hemoglobin layer did, as a source of constant absorption (there is only one main type of rhodopsin, its spectrum does not vary with oxygen in the dark-adapted eye, and its concentration is relatively constant except near the macula and the optic disk) in both directions. One way to compensate for this if necessary -- only if it is reducing the light return to intolerable levels -- is to bleach the photoreceptors with a pulse of bright light, thus temporarily destroying their ability to absorb light, and then performing the reflectometry. More on this later. For fluorescence, we do not need to include the rhodopsin as a factor since the mitochondria which contain the FAD are all anteriorly situated and the emission cone will pick up only fluorescence that is heading generally anteriorly from the fluorescent sources. (Good thing, too, because rhodopsin has an absorption peak at the same place as FAD's fluorescent emission peak.)

2.3.6. Other parts of the eye. The cornea and lens in particular are significant absorbers of light -- their structure is the reason why the visible spectrum is as limited as it is, for the retina left uncovered can absorb farther into the UV than the cornea and (especially) lens allow to reach it. The transmittance of the eye's structures is shown in this diagram from Boettner[34] It is plain that observation with excitation

Figure 2.12. Transmittance to various anterior surfaces of the human eye. From Boettner[34]. (1)Aqueous; (2)Lens; (3)Vitreous; (4)Retina.

wavelengths less than 400 nm in the intact eye will pay a severe penalty in light return (and will probably damage the eye first; see chapter on Safety). This is due to the lens itself; thus the use of long UV wavelengths in aphakes (persons without lenses, usually due to cataract surgery) is a consideration if it is needed, e.g. to measure NAD and NADH. The transmittance rises rapidly, so that excitation with a 436 nm or 458 nm source incurs only a small penalty.

There is significant scatter in the lens, probably of the Rayleigh kind. Mellerio[35] plotted the transmittance for young and old persons' lenses (Figure 2.13) and deduced from the relatively linear slope that Rayleigh scatter was probably the major process involved, although the data was not complete. As with the anterior vascular layer, if this scatter is uniform for all scan regions and is not affected by oxygen tension, it will be at least a uniform nuisance, a constant factor which attenuates the direct light transmission. It is problematic in that unlike absorbed light, light scattered by the lens will still excite reflectance and
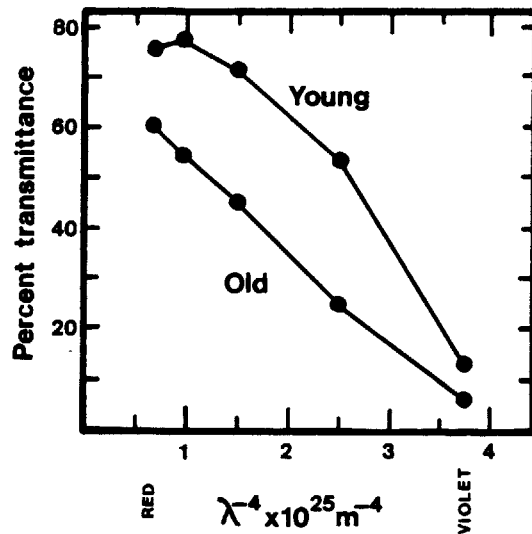
Figure 2.13. Transmission through the lens of the eye vs. λ^{-4}. From Mellerio[35]

fluorescence -- from parts of the retina other than the desired area.

The fluorescence of the cornea and lens has been demonstrated[36] and may be related to the oxygenation of these tissues -- it in fact is probably the fluorescence of respiratory-chain components. Since the overall fluorescence of the lens is perhaps 10 times that of the retina, and that of the cornea about equal to that of the retina, this is a severe limitation on performance of any instrument that tries to detect retinal fluorescence by passing through the cornea and lens. However, it should be possible to solve most of this problem by optically masking the system so that fluorescent light from the area of the cornea and lens struck by excitation light is stopped from reaching the detector. Then the only source of corneal or lens fluorescence is from light which has passed through the eye and reflected back, which at least reduces the problem by a factor of 100.

Finally, there are four main specular reflections in the eye, -- the four Purkinje reflections -- one each at the anterior and posterior surfaces of the cornea and the lens. These are far brighter than the diffuse

reflectances which have come from the retina and would easily swamp out any data which is in the same region of the image. In particular the first Purkinje image, which is the reflection at the air-cornea interface, is by far the brightest (the other three Purkinje reflections together are only 0.025 the brightness of the first image[37]). Fundus cameras deal with this problem with careful use of optical stops.

CHAPTER 3

PRACTICAL CONSIDERATIONS FOR RETINAL FAD MEASUREMENT

Now that we have the basic calculations of the light intensity we should observe on fluorescence and reflectance in the eye, we must turn these equations into a form that allows us to extract the redox state of FAD from one or more measurements. How easy this is to do, and how many measurements are necessary, is a function of the concentrations and absorption spectra of the component molecules mentioned last chapter. The choice of wavelengths to use also influences the degree to which these components confuse our measurement.

There are other practical considerations in the measurement of retinal oxygenation. Observing over a long time allows more light to be collected, and thus the signal-to-noise ratio may be improved; however, artifacts such as eye movement may render this method useless if the time is too long. Similarly, observation with low resolution may allow us to negate the effects of variations in the vascular layer by averaging them out; however, we may also lose definition of small areas of hypoxia in this fashion.

In this chapter, I will show how to combine multiple readings to extract the redox state, and then proceed to try and find optimum wavelengths which allow the best signal to noise characteristics and the fewest necessary measurements, given an estimate of the actual conditions of mitochondria and FAD in the eye.

3.1. Relevant Parameters in the Eye

3.1.1. Cone Angle. According to the classic Gullstrand model of the human eye[37] ,

Figure 3.1. The constants of the Gullstrand eye focused at infinity.

parallel light reaching the eye will be refracted to the retina in a cone approximately 22.5 mm in length (measured from the principal plane). The iris can open to a maximum of 8-9 mm in most persons; this would be an aperture of 9-10 mm at the principal plane. Thus the widest possible cone would have a half-angle of $\theta = \tan^{-1}(5/22.5)$ or $12.5°$.

3.1.2. FAD and other Chemical Components in the Retina. The retinal thickness (exclusive of the pigment epithelium) ranges from 100 μm at the ora serrata to 350 μm around the macula (dropping to a minimum of 90 μm at

the center of the fovea)[38]. As mentioned before, the major components we are interested in are the fluorochromes FAD and FADH, and the major absorbers: hemoglobin, lipofuscin, rhodopsin, and xanthophyll (which is found mainly in the macula). Spectra of some of these compounds are shown in Figure 3.2

Figure 3.2. Absorption spectra of hemoglobin, rhodopsin, and xanthophyll.

3.1.2.1. FAD. The fluorescence of FAD in well-oxygenated pigeon-heart mitochondria at pH 7 corresponds to the fluorescence of a 1.6 µM solution of riboflavin[23] ; since the concentration of FAD has been given as 20.5 µg/g for this tissue[39] , or (dividing by its molecular weight of 786) 26 µM, we can estimate that the value for $\eta_o \varepsilon_o$ for FAD is about 1.6/26 of riboflavin, or about 300 mmolar^{-1}cm^{-1}. The peak fluorescence of FAD actually approaches that of riboflavin, but this occurs at pH 2-3. The decrease of fluorescence with rising pH (toward typical intracellular range

of 7.0[40]) is probably due to intramolecular complexes in the flavoproteins, which unfold at low pH.

Figure 3.3. Dependence of FAD fluorescence and riboflavin fluorescence on pH. From[41].

FADH has a value of $\eta_r \epsilon_r$ which is only 0.25 that of oxidized FAD[23]. The quantum efficiency of FAD has been estimated as 0.3 at pH 3[42] and 0.025 at pH 7[41] so we may estimate a value of 12 $M^{-1}cm^{-1}$ for ϵ_o at 458 nm at pH 7. This is consistent with other reports[43]. At 436 nm the fluorescence drops to about 60% of this value. These data are summarized in the table

below.

	η	ε	$\eta\varepsilon$	C
		$(M^{-1}cm^{-1})$	$(M^{-1}cm^{-1})$	$(\mu g/g)$
FMN	.25	12000	3000	1? (retina)
FAD pH 3	.3			
FAD pH 7	.025	12000	300 (458 nm)	20.5 (heart)
				5? (retina)
			180 (436 nm)	
Riboflavin	.25	12000	3000 (458 nm)	
	.21	12000	2600 (436 nm)	

As for the concentration of FAD in the retina, no hard data is available. The concentration in brain tissue is only 14% of that in the heart[39] but the concentration in the retina is probably much higher because of its large metabolic rate; a concentration of 25% of that in the heart is suggested in [44] . Of concern is that FAD can split under various conditions into the mononucleotide FMN, which fluoresces in the same bands as FAD, but without a dependence on oxygen because its conformation is not altered by oxidation in the same way. FMN may have fluorescence efficiency up to 10 times that of FAD, and is found in tissues at about 1/10 the concentration -- i.e. it yields a fluorescence signal equal to the FAD signal, but does not have its oxygen variation -- thus, this provides a baseline of equal amplitude to the maximum possible FAD variation itself[41].

The value of B for FAD can be calculated to be .176 cm^{-1} for $\theta = 12°$ (.175 for $\theta = 3°$) at 458 nm. (This value was multiplied by 2.303 as described in Chapter 2). For a 300 μm thick retina, this gives a value for BL of .0050.

The fluorescence emission of FAD shifts by 10-30 nm towards the red when the free flavin is bound into the flavoproteins. Thus free FAD exhibits fluorescence peaks around 500 nm, for 450 nm excitation. However, the key respiratory-chain flavoproteins -- succinate dehydrogenase (from the

Krebs cycle), NADH dehydrogenase, dihydrolipoyl dehydrogenase and the fatty-acyl oxidases -- fluoresce with a peak in the 520-530 nm range, as the pigeon-heart data showed.

3.1.2.2. Hemoglobin. Hemoglobin has an extinction coefficient ε of 16(oxy-) or 12(deoxy-) $mmolar^{-1}cm^{-1}$ at 458 nm. At 436 nm (the wavelength of a mercury arc line) the extinction coefficient increases to 40 and 140 for oxyhemoglobin and deoxyhemoglobin respectively. In the vessels, the concentration of hemoglobin is typically 150 g/L (2.3 mM). Thus, over a major vessel with thickness 150 μm, BL for hemoglobin with 450 nm excitation is equal to 0.95-1.25 for $\theta \leq 12°$.

Note that the isobestics of hemoglobin occur at 448 nm, 505-525 nm (very close through this range) and at 549 nm. The 520 nm region is also the region where FAD has a fluorescence emission maximum in pigeon-heart mitochondria. Furthermore, in the 430-448 nm range, oxyhemoglobin absorbs less than deoxyhemoglobin. This means that as the pO_2 increases the flavin fluorescence increases but also the hemoglobin in the tissue passes more light (the opposite is true between 448 and 500 nm).

3.1.3. Total Retinal Transmission. If the total transmission of the retina approaches 100% (exclusive of the retinal vasculature), then the mathematics of solving for the redox state from fluorescence measurements are much simpler. Geeraets[45] and Bakker[28] have both studied this issue. Geeraets found that the bloodless retina of rabbits (of any pigmentation) without the pigment epithelium never transmitted less than 99% at any visible wavelength. With the RPE intact, transmission was between 6 and 60% depending on the rabbit's pigmentation. Human eyes appeared to have similar characteristics. Bakker recorded only the relative spectrum,

which he found to be without peaks and valleys.

3.1.4. Flavins in Other Eye Structures. The concentration of FAD in the cornea is given as 1 μg/g wet weight of tissue. Oxygenation of the cornea is provided by the atmosphere and shows little dependence on vascular changes, although corneal hypoxia and swelling is common in wearers of certain types of contact lenses[46].

The lens also shows fluorescence in the 450 to 520 nm band. Some of this is due to flavins (the lens flavin concentration is given as 0.14 μg/g or 1.7×10^{-7}M. and the lens is approximately 5 mm thick, so one might expect a similar order of magnitude to the retinal fluorescence; in fact it is considerably higher) and some due to other proteins of the lens crystalline matrix.

3.1.5. Time Course. Eye movement causes major problems in recordings which take longer than 1 second; this is the mean time between saccadic movements which will blur any image of the eye[48]. There are fine ocular microtremors in the 80-120 Hz range[49] of very small amplitude (less than 1 minute of arc) which probably do not cause significant blurring.

The time for acute hypoxia to the body to become known to the respiratory chain appears to be 20-40 seconds[50]. This agrees well with the finding[17] that the sodium-potassium pump, which is ATP-driven, starts to fail 15 sec after hypoxemia is started. The time course is similar for recovery upon reoxygenation in tissue which is undamaged by the hypoxia: the cytochromes and other components of the chain nearest to oxygen recover (in a perfused organ whose perfusate is altered) in 20 seconds; NAD, farther up the chain, takes over 30 seconds to recover to a state of 50% oxidation.

3.1.6. Rhodopsin. Rhodopsin in the photoreceptors has been measured to have an overall optical density of 0.20 (37% absorption) at its maximum of 520 nm[51]. The rhodopsin absorption can be bleached out by flash[52] or prolonged light[53] with a drop of over 0.15 log units of optical density. With exposures up to 30 seconds (using a wide-band xenon arc lamp) up to 90% of the photopigment can be bleached. Rhodopsin lies exclusively in the external segments of the photoreceptors, behind all retinal mitochondria, so it presents no obstacle to the passage of FAD fluorescence. It exhibits mild fluorescence itself, at 450 nm and 620 nm emission peaks[54].

Xanthophyll is mainly a problem at the fovea, and we will not deal with it in the general solution at this time.

3.2. Extracting the Redox State

Now with the numerical data guiding our approximations, we return to the problem of solving for the redox state of FAD given fluorescence and/or reflection measurements.

3.2.1. Solving the Equations. In chapter 2 we derived the basic equations for reflection and fluorescence from two species (FAD and FADH) in a medium which contained other nonfluorescing but absorbing species (e.g. oxy and deoxyhemoglobin), possibly located behind a vascular layer, and with light illumination and detection via two (possibly distinct) cones separated in angle by χ. To summarize, the final results were:

$$(55) \quad R = V^2 \, I_0 e^{-(B_T(\theta) + B_T(\phi))L} \frac{\sin^2\phi \, \sin^2\theta \, \cos\chi}{2 \, (1 - \cos\theta)} \cdot (\text{reflectivity})$$

$$(56) \quad F = \frac{V^2 I_0}{2} \frac{\sigma \varepsilon_o \eta_o + (1-\sigma) \varepsilon_r \eta_r}{\varepsilon_T + a_T} \frac{C_{FAD+FADH}}{C_T}$$

$$\cdot \frac{B_T(\theta)}{B_T(\theta) + B_{T,\lambda_f}(\phi)} (1 - e^{-(B_T(\theta) + B_{T,\lambda_f}(\phi))L}) (1 - \cos \phi)$$

where $\underline{V}$ is the effect of the vascular layer, θ and ϕ are the half-angles of the excitation and emission (input and output) cones, and ε_T and a_T are the absorption and scattering coefficients weighted over all components

$$(57) \quad \varepsilon_T = \frac{\sum_{i=1}^{N} \varepsilon_i C_i}{C_T} \quad \text{where} C_T = \sum_{i=1}^{N} C_i$$

and B_T is given by

$$(58) \quad B_T(\theta) = \frac{\ln(\cos \theta)(2.303)(\varepsilon_T + a_T) C_T}{(\cos\theta - 1)}$$

taken at the excitation wavelength λ (except where λ_f, the fluorescence emission wavelength, is specified).

Note that the fluorescence signal increases as ϕ increases; for the same intensity I_0, we will get more light out as we increase ϕ.

We did note that if the total argument to the exponential is much less than 1 then the fluorescent light is passed through with only minimal absorption and that we could create constants independent of any concentration such that

$$(59) \quad F = V^2 I_0 (K_o [FAD] + K_r [FADH])$$

The error associated with trying to do this linearization, as a function of the fraction of light absorbed as it travels through the retinal FAD layer is shown in the Table 3.2.

Table 3.2

$(B_T(\theta)+B_{T,\lambda_f}(\phi))L$	Absorption	Error
0.20	18%	10%
0.10	9%	5%
0.04	4%	2%
0.02	2%	1%

A little later on we will see if the biosystem we have to deal with and the signal-to-noise requirements we are faced with allow us to do this.

If we can get away with this linearization, then the solution is fairly easy: if we do our excitation and emission at isobestic wavelengths, and use a wide enough cone so that the vascular layer appears homogeneous, then the vascular layer presents only a constant attenuation.

3.2.2. Reflection vs. Fluorescence. It is clear from the above numbers that the major impediment to reflection is the vascular component, whereas the major source of fluorescence (not counting the lens and cornea, which can possibly be shielded optically) are the flavins — FAD, FADH, FMN. Since the flavin signal contributes so little to the reflectance, comparatively, and since we have already noted that the solution of the reflectance problem will be difficult for small numbers, we are directed to make fluorescence the main measurement. If we can take the vascular component to be primarily in front of the mitochondria, then we can isolate the term V^2 and have a very straightforward solution as in Figure 2.10. If we cannot, then we can use the form of Equation 56 with hemoglobin being among the nonfluorescent components C_i (with or without an additional term for anterior concentrations of blood).

The result for the first case is as follows: if we take two fluorescence measurements, the first from $\lambda_1 \rightarrow \lambda_2$ (excitation -> emission) and the second from $\lambda_3 \rightarrow \lambda_4$:

$$(60) \quad F_{1\to2} = V_{1,\theta}V_{2,\varphi}I_0 \ (\ K_{1o}\ [FAD] + K_{1r}\ [FADH] + K_{1x}\ [FMN]\ )$$

$$(61) \quad F_{3\to4} = V_{3,\theta}V_{4,\varphi}I_0 \ (\ K_{3o}\ [FAD] + K_{3r}\ [FADH] + K_{3x}\ [FMN]\ )$$

$$(62) \quad [FAD] + K'[FMN] = \frac{1}{K_{3r}K_{1o} - K_{1r}K_{3o}} \ (\frac{K_{3r}F_{1\to2}}{V_{1,\theta}V_{2,\varphi}I_0} - \frac{K_{1r}F_{3\to4}}{V_{3,\theta}V_{4,\varphi}I_0})$$

$$(\ \text{where}\ K' = \frac{K_{1x}K_{3r} - K_{3x}K_{1r}}{K_{3r}K_{1o} - K_{1r}K_{3o}}\ )$$

and similarly for [FADH]. The K_{nr}, K_{no}, K_{nx} are all known constants for fixed θ, φ, L, λ_i. The only variables are the concentrations of the flavins and the hemoglobin concentrations and vascular layer depth which are part of the variables V_i. In order to handle this,

1. If the fluorescence measurements can be done at isobestic or near—isobestic wavelengths on excitation and emission, the vascular contribution has only one variable (the amount of total hemoglobin in the total excitation and emission paths) so that a single reflectance measurement made at an isobestic wavelength will yield the variable $\underline{V}$.

2. If the fluorescence must be made at wavelengths removed from isobestics of hemoglobin, two measurements of reflection may be necessary, which can then be combined to yield total oxy- and deoxy-hemoglobin to substitute into Equation 8.

As far as the "unchanging" fluorescence of FMN and other unreactive flavins, they will contribute a fixed source of "noise" to our measurement. However, this will only dampen the ratio of [FAD] to [FADH], and not change its sense. This contribution could be isolated with a third wavelength, or with a forced change in oxygenation (subject breathing 100% oxygen) for which the added fluorescence (after vascular contributions are removed)

will be due solely to increased oxidation of FAD.

3.3. The Biological Signal and Noise; Choice of Wavelengths

It is clear from the above discussion that, irrespective of noise due to the process we use to detect these signals, that there is inherent biological "noise" which is obscuring our detection of the true redox state. Absorbing components and oxygen-insensitive fluorescences, plus the fact that tissues never become completely reduced, are major factors in the inability of researchers to see fluorescence changes greater than 10–20% in exposed tissue preparations even though FAD fluoresces four times as much as FADH. The eye's anterior structure and limited pupil size further limit the amount of information we can truly say is from the retina. We can, however, choose our experimental parameters in such a way as to give us as much advantage as possible.

The "signal" we are trying to measure is the FAD redox state. In fact, since it is not just the existence of FAD fluorescence but rather its value that concerns us, the most sensitive measurement would be one which makes it as easy as possible to distinguish different values.

1. To this end, note that if we continue from Equation 8 to get an expression for the redox state ([FAD])/([FAD]+[FADH]), and if we are able to subtract the FMN emissions from the measured fluorescence $F_{a\rightarrow b}$ to give $F'_{a\rightarrow b}$, that we will have as an answer (multiplying top and bottom by $V_{1,\theta}V_{2,\varphi}V_{3,\theta}V_{4,\varphi}$):

(43)

$$\frac{[FAD]}{[FAD]+[FADH]} = \frac{K_{3r}V_{3,\theta}V_{4,\varphi}F'_{1\rightarrow2} - K_{1r}V_{1,\theta}V_{2,\varphi}F'_{3\rightarrow4}}{(K_{3r}-K_{3o})\,V_{3,\theta}V_{4,\varphi}F'_{1\rightarrow2} - (K_{1r}-K_{1o})\,V_{1,\theta}V_{2,\varphi}F'_{3\rightarrow4}}$$

If $\lambda_{1\to2}$ is a wavelength pair at which $(K_{1o}-K_{1r})$ is large, and $\lambda_{3\to4}$ is a pair at which $(K_{3o}-K_{3r})$ is small, (or vice versa), then the magnitude of the right-hand side of Equation 9 will be as large as possible for a given change in redox state, i.e. the system will have greater sensitivity to small redox changes. Thus we would ideally choose, for our two fluorescences, wavelength pairs for which dF/d(redox state) is as disparate as possible. Note that in fact K_{1o}/K_{1r} must not equal K_{3o}/K_{3r} or the two equations will not solve.

The FMN signal is removed in the general case by using a third wavelength, thus creating three equations for the three unknowns [FAD], [FADH], [FMN]. If FMN can be considered to be a known constant, or if we can choose our wavelengths so that $K_{1x}/K_{3x} = K_{1r}/K_{3r}$, we do not require the extra wavelength (in the latter case, the FADH and FMN will be indistinguishable, however, so we would determine only the fraction of total flavins that is FAD.)

2. Also noticeable from the above equation is that if $V_{1,\theta}V_{2,\phi} = V_{3,\theta}V_{4,\phi}$, the vascular effect is cancelled out in the determination of the redox state. This is somewhat looser than the requirement for isobestic wavelengths. By the same token, any absorbing component for which a similar equality holds will cancel out of the equation.

3. Subject to the above conditions, we must realize that the fluorescence signal will be quite small and that the relative statistical noise of the measurement will go inversely as the square root of the fluorescence measurement. Thus we want to choose wavelengths

which allow maximal fluorescence and minimal absorption.

The peak fluorescence excitation of flavoproteins is in the 455-460 nm range; this also is an area where hemoglobin has little absorbance. As we go farther into the blue, we are faced with less ocular transmission, greater hemoglobin absorption (and a greater absorption difference) and less potent excitation of the fluorescence. Thus the 455-460 nm excitation band appears to be excellent from a theoretical standpoint (going from 436-458 nm we gain a factor of approximately 2 in the eye's transmissivity, 1.6 in fluorescence efficiency and up to 5 times in transmission through the hemoglobin -- the latter assuming 300 μm combined blood vessel thickness, hemoglobin concentration of 1.2 mM, and change in ε from 40 to 16 $mM^{-1}cm^{-1}$, so the relative fluorescence light output for 458 nm excitation would be up to 16 times that for 436 nm). Emission will peak at 520 nm. This would be one wavelength pair; the second pair would be chosen subject to the conditions above. It might be possible to choose a different excitation band (either the 436 nm mercury line or the 441 helium-cadmium laser, while not as efficient as the argon 458 nm line, would provide high light levels) so that the emission may be measured with a broad-band filter, which will allow as much light detection as possible. Or, we may choose to keep a single excitation band and divide the emission band into two regions which satisfy the conditions. As we will see, the bands centered on 520 nm and 535 nm appear to have suitably distinct values of K.

If a monochromator may be used and can make readings rapidly enough, we can actually plot a curve of fluorescence versus λ to which we may then fit curves of the fluorescence and absorption of our major components. Most likely, because of the time it would take, such a technique would be most useful as a wavelength-series measured at one particular spot on the

eye to which the experimenter/ophthalmologist is pointing. Lubbers and Wodick[55] used such a curve-fitting technique for fluorescent materials in cuvettes.

3.3.1. The effect of scatter.

We have been neglecting the effect of scatter in the discussion up to this point. Scatter is significantly more of a problem in the elderly, especially from the lens[56]. Boettner[34] gives measurements for both total and straight-line transmittance through the ocular media; at 436 nm or 458 nm, 50% of the transmitted light is not direct transmission in young eyes; this figure may increase to as much as 70% in the elderly. For 520 nm the figure is approximately 45%. This scatter should equally affect fluorescence and reflectance measurements at similar wavelengths; it would mean that a certain percent of the light detected is from the local area, and another fraction is an average from a 5-10° area of the retina centered on the desired point, which must be subtracted. It is for this reason that the average fluorescence brightness over a local area is not nearly as important as the variations in fluorescence through the area.

Methods which can be used to reduce scattered light are discussed further in Chapter 5.

CHAPTER 4

PRELIMINARY EXPERIMENTS; FORMULATION OF INSTRUMENT REQUIREMENTS

Now that we have a theoretical basis for finding the fluorescence and for processing multiple-wavelength images to extract the redox state, we shall describe some of the experiments which justified the conclusions of the last two chapters and which established the feasibility of flavoprotein fluorescence measurement.

4.1. Goals

In the long run, we seek to develop an instrument or method capable of safely reading FAD fluorescence in animals and humans, including all of the compensations described in the preceding chapters. We also seek to produce an image or map of the retina which shows geographical variations in this fluorescence. Before we can design such an instrument or method, we must have an estimate of the real signal and noise levels inherent to the biosystem (in this case, the eye's optics and chemistry, circulation, FAD concentration, non-respiratory FAD, etc.). Once we know these levels, we can determine if we can use camera and film to record images from a fundus camera, or if more sensitive techniques are necessary. We can also learn for the first time whether native retinal fluorescence can be seen at all, and whether there are variations with hypoxia. In particular, these experiments were designed with the following goals in mind:

1. First, to repeat the determinations published by Chance[23] of the FAD fluorescence spectrum found in mitochondria, and observe its variation with oxygenation and other environmental conditions;

2. Second, to see if we can observe retinal fluorescence in animals

at all, and whether we could observe a variation _in vivo_ when oxygenation was changed;

3. To estimate the fluorescence intensity, and the variation in fluorescence with oxygenation, for a given level of illumination upon the retina;

4. To estimate the intensity of noise (statistical, biological and optical) present in particular spectral bands;

5. To determine, from the above, how much signal intensity would be necessary to have a particular signal(oxygen-dependent FAD fluorescence variation) to noise ratio. Then, since we have data about maximum safe illumination levels to the eye (see chapter 9), and about the ratio of emitted fluorescence intensity to illumination intensity (from these experiments), we can determine how sensitive light detection must be to maintain safety and effectiveness.

6. Finally, to use the above parameters to specify instrument performance requirements, and to design the instrument accordingly.

4.2. Mitochondrial experiments

A set of experiments was performed to determine the optical and fluorescence properties of mitochondria, including full spectra, by observing them with a cuvette spectrofluorophotometer.

4.2.1. Materials and methods. Rat liver mitochondria from white Norway rats were prepared by a method modified from that of Chance and Schoener[57], and that of Sanadi:[58]

For each day's experiment, 10 g of fresh rat liver was minced on a Petri dish over ice, then suspended in 50 ml of a solution of 0.2M mannitol and 0.016M Tris base to keep proper osmotic and pH balance. This solution was kept in a homogenizer tube in ice and homogenized until fibrous liver tissue was no longer discernible. The homogenate was adjusted to 80 ml with mannitol/Tris solution, and centrifuged at 600G for 10 minutes. The overlying "skin" was discarded from the centrifuge tube. The supernatant was saved; the pellet was washed with 20 ml mannitol/Tris solution, then resuspended in 20 ml of solution, and spun again at 600G for 10 minutes. This pellet, containing nuclear material, was discarded. The supernatants were combined and centrifuged at 15000G for 5 minutes. From this, the supernatant was discarded, as well as the loosely packed "pink microsomes" which overlie the tan mitochondrial pellet. The pellet was resuspended in 20 ml solution, and spun at 15000G for 5 minutes once again. The supernatant was discarded; the pellet was washed with mannitol/Tris solution. The pellets, containing mitochondria, were used as soon as possible to prevent decomposition and changes due to lack of nutrients.

For experiments, the pellets were resuspended in the following test medium:

Mannitol	225 mM
Sucrose	75 mM
N-morpholinopropanesulfonic acid(MOPS)	50 mM
Malate	6 mM
Glutamate	6 mM
monobasic potassium phosphate	2 mM

adjusted to pH 7.4 except as noted below. The medium provides electrons (although, as Fig. 1.4 shows, the transfer is not as good in the absence of NAD). nutrients, and buffering.

For experiments, the mitochondria were suspended in 80 ml of the test medium; 1 ml of this suspension was aliquoted to the cuvette for each experiment. The final concentration of mitochondrial protein was determined by absorption spectrophotometry (A_{280}).

Mitochondria could be kept in a deoxygenated environment by perfusing the cuvette with nitrogen through a two-hole stopper, or reoxygenated by opening the cuvette to air and stirring. A more efficient source of electrons could be obtained by adding succinate, 6 mM, to the medium. Adenosine diphosphate (ADP), 1 mM, could be added to act as cofactor in the oxidation of FADH by oxygen, as shown in Fig. 1.4. Thus we could achieve the five classic redox states for mitochondria[59] :

State	Estimated % oxidation
1. No ADP or substrate	80%
2. No substrate	100%
3. Active (ADP + substrate)	80%
4. No ADP	60%
5. Anaerobic	25-50%

4.2.2. Pitfalls. Other factors could have enhanced or reduced the fluorescence in this experimental setup. These mitochondria are outside of cells and thus in a different environment from in vivo mitochondria; we did try to match the nutrients and the pH to reduce this artifact. If the solution becomes too concentrated, the mitochondria will absorb the incoming light and will absorb some of the emissions from each other, as described in Chapter 3; to reduce this, the concentration of mitochondria in the spectrofluorophotometer was checked along with the results, and a range in which fluorescence variation with concentration was linear was sought. The reaction medium, ADP, succinate, were checked for spurious fluorescence.

The mitochondria came from rat liver, not human eye. The relative concentration can be assessed (see below) but it is possible that there are variant types of mitochondria with different fluorescence properties. In general, different tissues' mitochondria tested in the literature have similar (though not exactly identical) spectra, including liver, heart, and brain tissue.

4.2.3. Results. The results of the experiments are summarized in the following figures. Fig. 4.1 shows the emission spectrum of mitochondria with excitation at 455 nm. This wavelength shows the greatest emission and the greatest O_2-N_2 difference. The peak emission from the oxygenated specimen, in repeated trials, is 20%-25% higher than that of the same cuvette perfused with nitrogen; the shape of the curves is consistent with

Figure 4.1. Emission spectrum of rat liver mitochondria; (a)nitrogen environment; (b) oxygen environment.

published rough curves.

Of further interest is the fact that the peak emission wavelength shifts from 520 nm for the oxygenated to 528 nm for nitrogenated mitochondria. There is a potential advantage in our instrument design if we choose wavelengths such as these, in which the fluorescence coefficients for oxidized and reduced FAD are moving in different directions (see discussion).

These curves really represent only a transition from State 5 to State 1, as there is no ADP to serve as cofactor in the reaction, nor is there succinate to provide reducing power. When ADP is added, we should achieve the higher oxidation of State 2; the experimental result bears this out, as Figure 4.2(b) shows a further 21% increase in fluorescence. Addition of succinate (State 3) brings levels back to that of the original oxygenated cuvette. There is also a significant effect of pH on the fluorescence, as

Figure 4.2. Fluorescence of mitochondria in oxygenated environment. (a) no
ADP or succinate; (b) ADP added; (c) ADP and succinate added.

suggested in Table 3.1. Fluorescence rose 4-5x at pH 3 compared to pH 7.4,
due both to conformational change in the flavins themselves and to the
unfolding of intramolecular complexes at low pH. The pH of mitochondria in
vivo has been measured at 6.3-7.6[40], so the pH 7.4 value is more
appropriate for our calculations.

4.2.4. Discussion. It was noted that the peak emission wavelengths
for oxygenated and deoxygenated mitochondria were different. Between these
wavelengths is a region where the fluorescence coefficients for the two
species are varying (with wavelength) in opposite directions; more specifi-
cally, the values of $dK(\lambda)/d\lambda$ for the two species differ significantly in
this band. We want our measurement to be as sensitive as possible, i.e. to
show a large change in intensity for a small change in $FAD_{ox} - FADH_{red}$.
Looking back at Figure 2.10, this value is given by

-70-

$$(64) \quad \frac{(K_{red}(\lambda_2) - K_{ox}(\lambda_2)) \, F(\lambda_1) - (K_{red}(\lambda_1) - K_{ox}(\lambda_1)) \, F(\lambda_2)}{K_{red}(\lambda_2) \, K_{ox}(\lambda_1) - K_{red}(\lambda_1) \, K_{ox}(\lambda_2)}$$

Assuming that $F(\lambda_1) > F(\lambda_2)$ (an arbitrary setting of subscripts), the sensitivity increases as $K_{ox}(\lambda_1) - K_{ox}(\lambda_2)$ increases, and also as $K_{red}(\lambda_2) - K_{red}(\lambda_1)$ increases Thus this is a region of high intensity change with small changes in species fraction. These two wavelengths may not represent the largest sensitivity, but certainly both of the above expressions are increasing in this region.

The experiments confirm that mitochondrial fluorescence appears to vary with oxygenation, ADP, and succinate in the predicted fashion, and gives us some indication of the amount of variation. We can use the measurements made above, along with a number of assumptions, to calculate the expected light level in a retinal-fluorescence experiment:

The actual ratio of I_{out}/I_{in} for the mitochondria was found by placing a mirror in the spectrofluorophotometer in place of the cuvette (see figure 4.3) and adding absorption filters until the emitted light was equal to that emitted by the mitochondria. This showed that the mitochondria at the concentration used (2.6% by weight) emitted fluorescence with an intensity 1/25,000 that of the incident light. Estimating the concentration of mitochondrial protein in the retina to be about 1% by weight (a similar order of magnitude to values known for other tissues) we can estimate that the eye's mitochondria would emit 0.4 times the intensity of the mitochondria in the cuvette. We further allow for the geometrical differences between the eye and the spectrofluorophotometer, assuming isotropic fluorescence radiation. For the eye with a 1mm diameter exit pupil and optical distances as in Figure 3.1, .00022 of the emission from the retina will pass through the pupil (if we increase the exit pupil to 8mm, the maximum dilated pupil size, this figure rises to .013); the spectrofluorophotometer

Figure 4.3. Geometry of the spectrofluorophotometer for mitochondria; a mirror may be placed in the cuvette location as shown by the dashed line.

with geometry as shown above collects .048 of emitted light.

Combining these factors, and including the transmission of the ocular media at 440 nm and 520 nm (.75 and .80 respectively, for excitation and emission), the first estimate for the light level expected from the eye would be

$$(65) \quad \frac{I_{out}}{I_{in}} = \frac{1}{25,000} \times 0.4 \times \frac{.00022}{.048} \times .75 \times .80 = 4.4 \times 10^{-8}$$

for a 1 mm exit pupil (for an 8 mm exit pupil the output is 61.5 times higher, and the light output goes nearly as the square of the exit pupil diameter for sizes in between).

For a 50 microwatt light input intensity to the eye this yields 2.2 picowatts. In photon-counting terms, we use the energy of one photon at 520 nm, and find

-72-

$$(66) \quad \frac{2.2 \times 10^{-12} \text{ W.}}{3.82 \times 10^{-19} \text{ J/photon}} = 5.7 \times 10^6 \text{ photons/sec.}$$

If this formulation is accurate, this figure can be used as an estimate of expected light level in our instrument design. The monkey experiments described below help substantiate this figure by allowing us to derive a second estimate of light output by separate means.

4.3. Monkey experiments

In order to better approximate the actual conditions of human retinal fluorescence measurement, and to actually test the idea that we could observe retinal fluorescence in vivo and its variation with oxygenation, a number of fundus-photography experiments were carried out in Rhesus and Cynomalogous monkeys, which have retinal structure similar to that of humans.

In these experiments, laser light at 457.9 nm was used to illuminate the monkey's retina, which was then photographed through a filter which would block the incident light wavelength but permit fluorescent light in the 520 nm region to pass through. The amount of light returning from the retina was measured under various conditions.

4.3.1. Materials and methods.
Rhesus and cynomalogous monkeys were screened for experiments, so that only eyes which were normal by history and by ophthalmoscopy were used. The monkey used for a day's experiment was anesthetized with intramuscular ketamine for both induction and maintenance. This agent was chosen over phenobarbital because it preserves respiratory function (thus eliminating the need for respiratory support) and because barbiturates have a direct blocking action on the electrontransport chain. The experimental eye was dilated with Mydriacyl; the other eye was allowed to close normally or was taped shut. A drop of natural tears solution was applied to the experimental eye at regular intervals to avoid dehydration.

For illumination and photography, an apparatus used for laser—contour angiography[60] was used. In this apparatus (Figure 4.4) an argon ion laser was fitted with a Littrow prism and tuned to the argon line at 457.9.

Figure 4.4. Apparatus used for monkey experiments. The animal lies supine
on the cotton chuck with its head between the pads of the head-
holder. Light leaves the laser to the left and is turned so that
it travels rightward through the beam expander, after which it is
focused on the small mirror which is suspended from the upper
fundus camera assembly, and directed from there into the eye.
Light returning from the eye is directed into the fundus camera.
The viewing position and the camera back are located at the top
of the fundus camera assembly.

This beam was expanded and spatially filtered to produce a uniform parallel
beam; in this beam could be placed a phase grating or a grating with dark
stripes which would produce a striped illumination pattern on the retina.
The beam was then focused on a small mirror which reflected light through
the nasal side of the eye and illuminated the retina in the region of the
optic disc, with a total illumination of 110 mw. Light leaving the pupil

-74-

on the temporal side was captured by the fundus camera and recorded on the film, which was focussed on the retina. Before reaching the film the light passed through one of two filters: an SB40 fluorescein barrier filter or a Kodak Wratten #58 filter, both of which pass wavelengths from 500 to 700 nm and have high opacity at 457.9 nm. The monkey is positioned in a head-holder which allows travel in three linear dimensions and one rotational dimension. Photographs of 1/4 to 1 second exposure were used.

For hypoxia studies, a breathing mask was placed over the monkey's nose and mouth. The breathing mixture could be selected to be room air, 100% oxygen or 100% nitrogen. Arterial blood gases were drawn via a femoral arterial catheter to measure blood oxygen pressure and saturation.

For photographs, Kodak Tri-X Pan 135 film was used and push-processed in Kodak D19 developer to an ASA rating of 1600. A standard-density step wedge was contact-printed on the same roll of film and developed along with the photographs to measure the film's exposure vs. density (D/log E) curve and compensate for differences in sensitivity between film rolls.

Optical density could be measured at any location on the negatives by a microcomputer-based microscope macrodensitometer[61] which could be programmed to look at the same sites on each of several different negatives.

4.3.2. Results. Figure 4.5 shows a typical monkey fluorescence photograph. (In some of the photographs small pinholes in the barrier filter gave rise to Airy-disk artifacts). The retinal tissue, optic disc and vessels are clearly seen. The collagen in the sclera directly behind the disc fluoresces brightly at these wavelengths; in some eyes the scleral collagen extends past the borders of the disc, forming a crescent-shaped high-fluorescence region. The blood vessels, located on the inner surface of the retina in the peripapillary region, are dark, indicating that the fluorescence arises from behind the vessels, i.e. in the retina, choroid or sclera.

Although the lens and cornea fluoresce in these wavelengths, their fluorescence is diffused over the entire image by the optical system, allowing the in-focus retinal fluorescence to become the major factor determining the intensity variation across the picture -- thus the image is sharp.

When the striped grating is placed in the laser beam, the fundus is illuminated in a striped pattern. Since the observing axis is titled 12°

Figure 4.5. Retinal fluorescence photograph from a Rhesus monkey.

from the illumination axis, we can observe locations in which only anterior fundus fluorescence could be detected, and other sites in which only the posterior fundus could be detected (Figure 4.6). The thickness of the observed section for each site is calculable. Our observations show that the fluorescing structure is thicker near the disc and thinner near the macula, corresponding to the relative retinal thickness in these regions.

Sample experimental measurements of retinal fluorescence taken with the animal breathing 100% oxygen for several minutes, followed by 100%

Figure 4.6. Use of a striped grating for observation of three-dimensional structure. Expanded laser light passing through the grating illuminates the fundus in a striped pattern. The observing axis is tilted from the illuminating axis, so that some returning light rays carry fluorescence from only the posterior fundus, others from only the anterior fundus.

nitrogen for 40 seconds, followed by restoration of oxygen breathing, are shown in Figure 4.7. In these experiments, In nearly all sites except those which lie over the optic disk (and are therefore subject to a strong collagen fluorescence signal, such as site 10), fluorescence is 15-30% less after nitrogen breathing compared to oxygen breathing; the average decrease

is 22.8% with a 5.1% standard error (p<.01), demonstrating the change in fluorescence that would be expected if FAD were becoming reduced under hypoxic conditions. After reinstitution of oxygen breathing, the fluorescence rises again (see below).

Figure 4.7. Comparative fluorescence of 10 sites on the monkey retina after breathing 100% oxygen for several minutes and 100% nitrogen for 40 seconds.

The blood oxygen saturation is 100% under oxygen breathing and 43% under nitrogen breathing, with a return to 100% after 30 seconds back on oxygen.

Curves of the change in fluorescence with a continuous series of

photographs at 2.1-second intervals, starting from the change to nitrogen
breathing, are shown in Figure 4.8.

Figure 4.8. Curve of fluorescence vs. time during (first line) nitrogen
breathing; (second line) restored oxygen breathing. The fluores-
cence values are normalized to a scleral-crescent site on each
film frame.

Note that there is a delay between the time oxygen breathing is restored
and the time the fluorescence curve starts to rise.

4.3.3. Pitfalls. One important question is whether the recorded light represents fluorescence at all, or whether it is just the much stronger incident light breaking through the blocking filters. The following experiment is used to check for this possibility: instead of one filter, two blocking filters are placed in series and the resulting "fluorescent" light level is compared to the usual single-filter light level (Figure 4.9).

Figure 4.9. Two-filter experiment.

If the light which passed through the single filter were incident light of wavelength λ_1, then the second filter would attenuate it with its

λ_1 response and the two-filter image would be far dimmer than the single-filter image. If, however, light passing the first filter were fluorescent light of wavelength $\lambda_2 = 520$ nm, the second filter (Wratten 58) would only cut its intensity by about 55%. The experiment shows the latter to be the case, indicating that the majority of light is at fluorescence wavelengths.

In the O_2-N_2 experiments, for some sites the rise in fluorescence when oxygen is restored is not as great as the fall was initially. It is possible that this is related to drying of the eye during the two-minute period of continuous photography or to local autoregulatory mechanisms which constrict blood vessels and thus cause blood flow to be slower at the time well-oxygenated blood begins to recirculate.

Another question to consider is the changing absorption of hemoglobin as it becomes more and less oxygenated. If the hemoglobin absorption were to increase when it became deoxygenated, this would also result in a decreased fluorescence signal, causing confusion with a signal decrease attributed to reduction of FAD. In fact, the absorption of deoxyhemoglobin is _less_ than oxyhemoglobin at 457.9 nm, and nearly equal at 520 nm, so this is not a problem; if the change in fluorescence were due to hemoglobin changes it should have increased with poorer oxygenation.

4.3.4. Discussion. These experiments do not use a second wavelength to obtain numerical values of FAD_{ox} and FAD_{red}, nor are there compensating wavelengths to accurately remove artifacts as discussed in the previous chapter. Nonetheless, this experiment provides strong evidence that we could observe and localize monkey retinal fluorescence in wavelengths attributable to FAD, and that a change in fluorescence could be demonstrated in the living retina with changes in oxygenation.

The amount of light input to the monkey's eye is much brighter than

the acceptable human safety standards (see Chap. 9), especially for illumination times longer than a few seconds. In fact, on post-experiment inspection one eye of one monkey showed some retinal changes attributable to laser overexposure. We would need to use considerably lower light levels in order to have a system safe for human use.

As we did with the mitochondrial experiment, we can use the data from the monkey experiments to provide an estimate of the amount of light we can expect when we are measuring human retinal fluorescence. Film negatives from the monkey photos show a range of optical density from 0.3 to 1.0 over the retina for 1/4 second exposures at ASA 1600. Photographic film sensitivity can be characterized by the relationship between the light energy striking the film and the resultant optical density of the negative. Some sample values for Kodak Tri-X 135 film at its usual ASA 400 rating at 520 nm are:

Energy $(ergs/cm^2)$	O.D.
.40	1.0
.08	0.6
.025	0.3

Now, 1 erg = 10^{-7} W-sec, and a 1/4-second exposure with film pushed to 4 times its usual speed is roughly equivalent to a 1-second exposure at normal speed. Thus, for an O.D. range from 0.3 to 1.0 under these conditions the light emerging from the eye ranges from 2.5×10^{-9} to 4.0×10^{-8} W/cm^2. As the retinal image takes up 1.4 cm^2 on the negative, and the total illumination is 110 mw, we can calculate

$$(67) \quad \frac{I_{out}}{I_{in}} = \frac{2.5 \times 10^{-9} \text{ to } 4.0 \times 10^{-8} \text{ W/cm}^2}{(110 \text{ mw}) (1.4 \text{ cm}^2)}$$

$$= 1.6 \times 10^{-8} \text{ to } 2.6 \times 10^{-7}$$

For 50 microwatts illumination (and assuming we use the same barrier filter for human work) this corresponds to 8.0×10^{-13} to 1.3×10^{-11} W/cm^2, or 2.1×10^6 to 3.4×10^7 photons/sec.

This figure is the same order of magnitude as the figure reached from the mitochondrial data, lending support to both calculations. Unfortunately, it also indicates that the reduced illumination required for human study makes recording data on Tri-X film (or nearly any commercial film) unlikely to be successful.

4.4. Human photography studies

An attempt at recording human retinal fluorescence was made, using a standard Zeiss fundus camera which provides broad-spectrum flash illumination up to 155 mJ/cm^2 (this light level, while found in clinical use, actually exceeds ANSI light level safety standards when used without a filter[62] -- see Chap. 9) with a Kodak Wratten #34 filter for illumination

and Wratten #58 as an emission filter. Various films and developers were used, including an ASA 3000 surveillance film (Kodak RE 3475) and a hydrazine-based developer[63] purported to push Tri-X and Ektachrome to ASA 20,000; however, no convincing image could be made to appear on the photographic film. It is clear that safe illumination levels require electronic light detection in order to produce a useful fluorescence image.

4.5. Light level requirements

We can readily derive an expression for the amount of light signal that is actually needed in order to get a reasonable signal-to-noise ratio.

In our case, as in the case of most instruments which try to distinguish between two states (such as CT scanners), the key parameter is not the overall signal level so much as the contrast between the two states we are distinguishing -- oxygenated and deoxygenated tissue, or oxidized and reduced FAD. The amount of common-mode fluorescence exhibited by both FAD_{ox} and FAD_{red} is actually a detriment to our ability to distinguish the states, because of its statistical noise. The other major source of noise, if we are using electronic photodetection, is the inherent dark current (dark counts, in the case of photon counting) of the detector. The effective contrast-to-noise ratio for each discrete image area (pixel) is given by

$$(68) \quad CNR = \frac{\text{change in fluorescence } (FAD_{ox} - FAD_{red})}{\text{shot noise + dark noise}}$$

This figure increases as overall light output increases. Let N_r be the total fluorescent light output (in photons) from a given spot when all FAD is reduced, and let $\underline{x}$ be the fractional increase in fluorescence when FAD

-84-

is oxidized, i.e. $\dfrac{FAD_{ox} \text{ fluorescence}}{FAD_{red} \text{ fluorescence}} = 1 + x$

then the CNR is on the order of

$$(69) \quad CNR = \dfrac{N_r \, x}{\sqrt{N_r\,(1+x) + \text{dark noise}}}$$

$$= \sqrt{N_r \dfrac{x^2}{1+x}}$$

if the dark counts are insignificant (true for most good photomultipliers, which feature dark counts of 10–100 counts/sec).

We can solve this for the light output N_r required for a given $\underline{x}$ and a desired CNR. This relationship is shown in Figure 4.10. Selected examples include:

x	CNR	N_r (counts)
.05	10	42000
.10	10	11000
.20	10	3000
.10	5	2750
.20	5	750
.10	2	440
.20	2	120

Figure 4.10. Contrast-to-noise ratio (CNR) as a function of the number of photons counted for a pixel. x is the fractional increase in fluorescence from reduced to oxidized FAD.

From our derivations above we estimate light output of 10^5 photons/sec/microwatt illumination. Using this figure, we find that for a 15% FAD_{ox}-FAD_{red} fluorescence variation and a 10:1 CNR we need illumination of .051 microwatt-seconds (5.1 x 10^{-8} J.) for each pixel. Changes in some parameters, particularly the exit pupil size, can reduce this requirement.

4.6. Instrument requirements

It is clear that we require an instrument with electronic detection to be able to detect the fluorescent signal at the light levels derived above. The use of electronic imaging also allows us to process images more easily,

and to combine the various wavelengths and introduce corrections readily.

The design of a noninvasive retinal fluorometric scanner would include the following requirements:

1. Sensitivity. As described above, the instrument must be able to provide good signal-to-noise ratio with inputs in the 10 to 50 microwatt range and with the inefficiency of the fluorescence process. A solid-state detector such as a photomultiplier or ISIT would be necessary.

2. Resolution. The resolution of an instrument which would be used to help direct photocoagulation targets must at least be better than the size of the photocoagulation spot itself (about 400 microns). More importantly, we would like resolution which allows us to see geographic landmarks in the eye, and a resolution which is finer than the minimum width of a hypoxic area caused by capillary blockage. (This minimum width is essentially a "maximum bandwidth" of hypoxia and tells us what our spatial sampling frequency must be.) 100 micron resolution would help us to see detail clearly. The maximum field of the instrument should be comparable to fundus-photo fields (30°)

3. Variable scanning. It is not initially known what the optimum parameters will be in practice for the retina scanner, and the parameters may differ for different uses. It should be possible to vary the scan as to image size (so that rapid surveys could be made but fine detail could be seen where necessary), resolution (fluorescence may require less resolution than reflectance in order to capture enough light), speed (for the same reason), and it should be easy to change the position of the scan, whether by

adjusting the patient's eye position or the scanner itself.

4. __Adjustment for ametropia__. It should be possible to adjust the instrument for persons with focusing corrections. This is especially necessary in the diabetic population, which even in the pre-retinopathic stage has more average visual correction than the population in general.

5. __Desensitivity (relative) to eye and head movement__. Small movements should not severely affect the focus and resolution of the instrument.

6. __Multi-wavelength operation__ , since as we described above more than one image must be recorded to find the FAD signal in the presence of the other emission determinants. The images at different wavelengths should be taken at the same time or as close together as possible.

7. __Interactive work__. The images of the eye should be immediately available for review; first, to ensure that the positioning and data acquisition is acceptable, and second, so that the study of a subject can be guided by the initial images while the subject is still on hand: regions of interest can be studied more closely, etc. A direct or TV-based viewfinder is also essential.

8. __Rapid image acquisition__. Although long exposures would enhance the signal-to-noise ratio, they also lead to blurring due to fine movements, and in the extreme case to variations in the physiologic state during the capture of an image. Other than ocular microtremor (which occurs at 80-90 Hz and has an amplitude less than our projected pixel size), significant eye movement can be avoided with "exposures" less than 1 second.

9. _Image processing_. The corrections needed to the raw fluorescence image require a number of operations performed on one image or several images together. The images must be carefully held in registration, and the operations described should be available either on the system directly or on an outside computer.

The scanner itself must be equipped with sufficient algorithms to clearly display, mark and partially analyze the pictures acquired. Simple operations such as averaging, 2-dimensional convolution (to find edges, for example), brightness histograms and thresholding are very often necessary in initial evaluation of raw images before processing occurs.

10. _Subject comfort_ should include comfortable sitting position and headholder.

11. _Ease of use_. The changeable parameters and mechanics described above should be as easily accessible to the experimenter as possible -- most of the system should be under control from a central terminal, and focusing adjustments, disks, controls, etc., should be in easy reach. Programming for new experimental situations, and for image processing and analysis, should be as quick and easy as possible.

12. _Safety_. The dangers of intense light exposure to the eye have been well studied[64], at least for laser exposure, although not as well for scanning or imaging instruments. The light level passing into the eye must be kept at a safe level, and suitable failsafe means developed to protect against accidental exposure to unsafe levels.

13. _Range of subjects_. Humans, monkeys (for experiments involving

induced hypoxia), and possibly rabbits are being considered as objects for study with the instrument. It should be capable of recording from all of these.

In the next chapters design concepts of a retinal scanning instrument are explored.

OPTICAL DESIGN

5.1. The Optics of the Eye

Let us examine again the diagrammatic representation of the Gullstrand model eye, in Figure 5.1.

Figure 5.1. The Gullstrand model of the eye's optics, when focused at infinity.

All subjects for experimentation, whether human or animal, are pretreated with a drop of 0.5% tropicamide (Mydriacyl, Tropicacyl) to dilate the pupil and paralyze accommodation; thus we may assume that the eye is at

its farthest focus -- infinity for a emmetropic eye, somewhat less for the nearsighted.

Shown in the figure, along with the physical positions of the cornea and lens, are the first-order planes of the effective optical system formed by them. Because of the difference between refractive indices of the media anterior and posterior to the cornea-lens combination, the principal planes (where the cone angle formed by parallel rays entering the eye is defined) are more anteriorly located than the nodal planes (light heading for the center of the first nodal plane leaves the eye, parallel to its initial course, heading from the center of the second nodal plane). As mentioned before (chapter 3), the pupil can dilate to 9 mm maximum diameter; thus $\tan^{-1}(9/16.9) = 12.5°$ is the maximum total half-angle of any input or output cone. The nodal distance, 16.9 mm, will define the size of the retinal spot (and total scan). Our optical system performs imaging of both the retina and the pupil, and we will make use of both of these planes in the discussion.

The optic disc is located 15° to the nasal side of the visual axis. The disc is rich in structure, and is also a site of major changes in DR. We often will be interested in observing the disc; so one parameter of our system's optics is that we can scan at least as far out as the disc.

Many persons are not emmetropic, but are nearsighted (myopic) or far-sighted (hyperopic) with corrections measured in diopters (inverse meters). A deviation from emmetropia in either direction is called ametropia. For small corrections, each diopter of ametropia implies that the eye's optics are focusing distant objects 0.3 mm in front of or behind the retina. In a system designed to image sharply on the retina, we must take account of this variation so that we can observe a wide range of people, particularly

since diabetics have a larger average correction even before retinopathy develops[4].

5.2. Criteria for the Optical System

The following criteria are among the most important to be satisfied in designing the optics of the retinal imaging device:

1. <u>Brightness</u>. The system must be capable of supplying enough illumination to the retina so that the very dim fluorescence (see Chapter 4) is still detectable. A corollary to this is that a high percentage of the light returning from the retina should pass through the eye and the emission optics to be detected. From our theoretical discussion we know that this increases monotonically with φ; thus it is in our interests to have as much of the eye's pupil area available as possible for passing emission light from the eye.

2. <u>Resolution</u>. The system should be capable of resolving sufficient detail to be useful in its application. A copy of a fundus photograph was digitized and sampled at varying resolution; it was determined that for resolution of 100 μm or better, the structure of the fundus was readily recognizable, though blurred at the coarsest resolutions. Laser photocoagulation spots are 400 μm in diameter; if the images from our instrument were to be used in an examination of or adjunct to laser treatment this would be the limiting resolution. For detection of early changes of DR the limit is the "spatial bandwidth of hypoxia" — we want to sample finely enough to detect all local changes due to hypoxia. This

limit can be estimated from noting extent of damage on series of retinal photographs in diabetics, and from the mean vessel-vessel distance on micrographs -- these show the area of tissue around a vessel over which no further perfusion is needed. The distance is on the order of 80-100 μm.

3. Safety. The total light introduced to the subject's eye must be well within the ANSI Standard Values for illumination safety. (see Chapter 9).

4. Field. Standard fundus photographs cover an area of 30° in each of several views of the retina. While other constraints may cause us to limit our field for the sake of higher speed or resolution, we would like to have a field close to this size.

5. Reflections. Specular reflections from the lens and cornea, which will be much brighter than the retinal diffuse reflectance, must be eliminated from the output.

6. Effect of lens and cornea. The fluorescence from the lens and cornea should be minimized. Absorbance from these structures should be minimized, and more importantly, should not change over the retinal image field, lest this difference be taken for a retinal metabolic variation.

7. Wavelengths. The wavelengths should be chosen as in Chapter 3 to excite FAD vigorously, transmit well through the eye and be little absorbed by hemoglobin and others in the retina. Use of more than one wavelength would help in optimizing the response and in adding special illumination, e.g. isobestics, to compensate for absorbing components.

8. Ametropia correction. Sufficient adjustment should be available to

rapidly correct the system optics for persons who are nearsighted or farsighted. One or more lenses in the system should be moveable to achieve this end.

The design of the optical system, described below, has been chosen to maximize, or at least make reasonable compromises among, these criteria.

5.3. Optical Design

We begin the discussion of the optical system by looking at a few significant design choices which were made early: the choice of illumination source and detection method, the dual-focus scheme by which both retina and pupil are imaged, and the choice of whether the excitation or emission pathway, or both, should be precise imaging pathways or nonspecific illuminators/collectors.

5.3.1. Light source.
The product (light source brightness) x (detector sensitivity) is the minimum fluorescence or reflection power that can be detected with our system. Our estimates of expected light level in Chapter 4 are only order-of-magnitude estimates. We must decide how much light intensity we will need, and how much can be tolerated safely. The first-order rule would be that more light (up to safety limits) increases performance. This would suggest use of a laser, particularly an argon laser whose 458 nm line is nearly ideal for FAD excitation.

Other than providing the highest brightness by far, the laser suffers from a number of problems, including space, the requirement for cooling and for large power supplies, limitation to one wavelength, cost, and the greater possibility of accidental overexposure of the eye. The HeCd laser (standard line 441 nm) is less expensive and smaller than the Ar laser but

has rapid intensity variations up to 10%. These could be compensated for if necessary.

The mercury or xenon arc lamp is probably the next brightest source of continuous illumination. A mercury arc has a strong line at 436 nm, which is not ideal but is still capable of exciting more fluorescence than the broadband xenon arc. The mercury arc also has lines at 514 nm which is isobestic for hemoglobin. Certainly the light level of the arc will be more than adequate for reflectance imaging. The mercury arc was thus chosen as the light source for our basic system. Consideration has been given to the use of an argon laser for the FAD fluorescence channel only; an imminent change in the physical layout of the laboratory may make it possible to do this.

In choosing the particular arc lamp for the application, it is important to remember that we will be putting light into a very narrow cone and imaging it into a very small spot on the retina. As Figure 5.2 shows, there is an invariant (known as the optical invariant) at every focused plane in a given optical system:

(70) $h \tan \theta = \text{constant}$

Or, in three dimensions,

(71) $A\Omega = \text{constant}$

where A is the area of the focused image and Ω is the solid angle of the cone.

Thus, if we set the size of the retinal spot and the width of the eye's pupil that will be used for excitation, then we have already set the system geometry and we will be limited not by arc power, but by brightness (power per steradian per unit area of arc). The Oriel HBO-100W 100-watt

Equations shown in figure:

1) $\dfrac{h_0}{h_i} = \dfrac{d_0}{d_i}$

2) $\tan\theta = D / d_0$

$\tan\phi = D / d_i$

3) $h_i \ D/\tan\theta = h_0 \ D/\tan\phi$

$\longrightarrow h_i \tan\phi = h_0 \tan\theta$

Figure 5.2. Origin of the optical invariant.

mercury arc features an exceptionally small arc (0.25 mm gap) and provides the highest brightness of any lamp investigated. The brightness multiplied by the optical invariant is the total illumination incident on the retinal spot.

5.3.2. Detector. Unlike light source brightness, we can push detector sensitivity indefinitely without incurring any safety problems. Photographic film has been shown to be much too insensitive, even when we are

pushing the safety limits on excitation[62]. Electronic detection is necessary. Of the major types of electronic photodetectors, the photomultiplier has a clear advantage in low-light-level S/N ratio over the photodiode, avalanche diode, and intensified silicon imaging tube (ISIT)[65].

Cathode sensitivity for S/N = 1
(typical minimum detectable illumination)

Photodiode	10^{-7} W.
Avalanche diode (gain=100)	10^{-11} W.
ISIT	10^{-12} W.
Photomultiplier	10^{-16} W.

Although the ISIT has the desirable property of area imaging, the greater absolute sensitivity of the PMT was deemed advantageous. Imaging by means of moving scanners is necessary on input or output.

Photomultipliers operate at a still lower effective noise level if operated in the photon counting mode. Rather than measuring or integrating the total current as a PMT-based detector would do in standard operation, photon counting relies on (1)the fact that a PMT's response to a photon of light is a packet of charge (produced by release of photocathode electrons when the photon strikes the cathode, followed by multiplication of the number of electrons at each successive dynode) and (2)the fact that the probability distribution of the magnitude of charge produced at the anode from a single photon at the cathode is Gaussian. The "dark current" generated by thermal disturbances within the PMT itself has a distribution as shown in Figure 5.3. If we have an amplifier which is fast enough to amplify each charge packet, followed by a discriminator which establishes a threshold as shown in the figure and produces output pulses only when the charge exceeds the threshold, then Bayesian analysis demonstrates that we will be detecting as many true photons with as little noise from "dark counts" as possible. A side benefit of photon counting is that the light emission output is in digital pulse form, simplifying the data-acquisition

Figure 5.3. The distribution of charge packets produced by incident photons and by noise sources within the photomultiplier. Whereas ordinarily all of these currents would be recorded as due to light, selection of a threshold as shown allows us to detect as many true photons as possible with a minimum of "dark counts".

problem.

Photomultipliers designed for photon counting feature higher voltage between the cathode and first dynode than between the dynodes, so that the current due to photons will have an extra amplification advantage over the current generated internally. The EMI 9814A photomultiplier has this feature and has a bialkali photocathode which is relatively selective for the 400-550 nm detection range.

5.3.3. Dual-focus. It is clear that there must be one pathway — either excitation or emission or both — which provides sharp focus on the retina, if we are to create a retinal image as the final output. In a scanning system, or any system wherein the light from a fixed source is projected on variable parts of the retina, there will be beams approaching the eye from a variety of incident angles to reach the different parts of the retina (See Figure 5.4). These beams could be designed so that they all pass through different portions of the cornea and lens; there will be a crossover plane somewhere in the system (there may be two — a crossover plane for X and one for Y) at which the beam will not change as the scan proceeds.

If we adjust our system so that this crossover point occurs near the center of the lens (in a front-to-back sense), we gain a number of advantages. First, the incident beam passes through a nearly equal volume of lens for all positions in the scan. This will equalize the effect of lens fluorescence over the entire image field, so that variations across the field are less likely to be due to lens-fluorescence artifact. Second, the eye's lens as well as the retina is assigned a fixed place in the overall optical system, which simplifies adjustments. Third, the area of lens and cornea used by the excitation path over the entire scan is minimized. This means that stops which must be placed in the emission path to block lens and corneal fluorescence (and specular reflections, for reflectance work) are as small as possible for a given intensity of illumination, i.e. for a given optical invariant. As mentioned above, the greater area available for the emission pathway increases the total emitted light.

Webb[66] utilized the idea of maintaining a fixed beam on the eye's lens with some success. Because he centered the excitation pupil within

Figure 5.4. Incident beams, from a scanned system, shown imaged on two
            points on the retina. The rays cross through the same area of
            the lens in this configuration.

the emission pupil, however, there was still significant range over which

specular reflections from the lens and cornea could be transmitted through

the emission pupil into the detector; thus the stops must be quite large.

We tried this scheme and found intolerable corneal reflections plaguing our

image quality. The arrangement finally developed for our retina scanner
displaces the excitation and emission pupils

Figure 5.5. Fixed, distinct excitation and emission pupils. A: the excita-
tion pupil is centered within the emission pupil. B: the excita-
tion and emission pupils are side by side. The size of the stop
needed to block specular reflection and lens fluorescence are
smaller in this scheme.

side by side (Figure 5.5). Ray tracing demonstrates that in this way the

locus of points in the pupil plane where specular reflections can appear is

entirely on one side of the midline. A larger emission pupil is permitted

in this arrangement.

In order to place a stationary spot on the lens, we must design so that the x and y scanning mirrors are conjugate to (imaged upon) each other and conjugate to the lens of the subject's eye. There are then two focus adjustments which must be made for a given subject: a retinal focus which allows a focused scan pattern to reach the subject, and a pupil focus which forces the stationary spot to match up to the subject's lens. This spot will not be sharp, for it is not an image of the arc; rather it will be an unfocused disc which is nonetheless unchanging as the scan changes.

5.3.4. Choosing the imaging pathway. Fundus cameras and other low-level eye photographic devices which use film operate on the principle of the Maxwellian view[67].

Figure 5.6. Maxwellian viewing system, and ordinary retinal focus system.

In this system a small source of light is focused on a spot on the pupil;

the light continues on to illuminate the retina. The film is conjugate to the retina and records the retinal image. Because there is a large lens used for the imaging and because the cone is not predefined by the pupil and retinal spot size, the optical invariant can be larger for this system and thus more light can get to the retina; more exactly, the same amount of light reaches the retina per unit angle, but reaches it over a wider angle.

Other fundus cameras use the Maxwellian system in the opposite sense: the retina is flooded with light from a large diffuse annular source, and the sharp imaging is done, again in the emission path, by imaging the retina on the film via a small area of the pupil (typically 1.5 mm diameter) and a large imaging lens as in Figure 5.9a.

For flash photgraphs, the Maxwellian system works quite well. Its disadvantages include: inability to get a sharp retinal spot, placing on the emission path the requirement for sharp imaging; and potentially, limitation of overall scan angle.

If we choose to let the emission path do the sharp imaging of the retina, as the above systems do, we would need a large excitation cone to provide wide-angle illumination, and the Maxwellian system could be used.

The primary disadvantages to imaging in the emission path are light level, image quality, and focusing. In our system, because of the inefficiency of native fluorescence in the retina, we must ensure that the emission path transmits as much light as possible. While it is desirable, we do not have to be so efficient on excitation; we can excite the eye with as much light as safety will allow -- but even with this excitation brightness the fluorescence emission will be very small. It has already been stated that, for any given point on the retina, the angular width φ of the emission cone leaving that point should be large for high emitted light output;

thus we want to make sure that a large area of pupil is available for emission. Using the full pupil diameter introduces aberrations from the eye itself which impair image quality. We also need large-diameter lenses in the emission path, which would have a small f/# and cause further deterioration in imaging quality and field depth. The optical parameters that pass the most light in this circumstance are in conflict with those which give the best image quality. We have the excitation system designed to use large f/# lenses and a small pupil of the eye (and thus does sharp imaging), while our emission system is arranged for maximum light output.

Additionally, in a system which images emission light, the light is on all parts of the retina at once; thus, to avoid wasted retinal exposure to light we need a detector which is capable of simultaneous recording of light from a large area of retina. Although this theoretically allows the entire scan to be completed at once, we would not be able to use a single photomultiplier, and would sacrifice sensitivity by using film or run great cost by using multichannel plates. Furthermore, in such a system the subject would not see a sharp raster but rather a uniform field of light; thus we could not set the focus to be that setting where the subject sees the raster, but must adjust it based on the quality of the recorded image — a more tedious process.

Another peculiarity of our system — the fluorescence of lens and cornea at similar wavelengths to retinal fluorescence — imposes a further restriction. If we used a large annulus of light for excitation, even if it all became incident on a single retinal spot so that we could use a PMT, we would stimulate fluorescence of these anterior components which would be very difficult to exclude from the collected light.

The problems of scattered light are found in both excitation and emis-

sion imaging systems. In the former, excitation light that is scattered by the lens illuminates the "wrong" area of the retina, and is collected with that light which was directed to the proper retinal spot. All light reaching the PMT in a certain time interval is taken to be from the correct spot. In emission imaging, light emitted from one spot on the retina may be partly scattered by the lens on the way out so that it is directed toward the "wrong" spot on the film, which would be taken to be emission by the corresponding wrong part of the retina.

It is possible that we could do imaging on both excitation and emission paths, by synchronizing two sets of scanners. The only advantage to this over the system chosen is the better performance with respect to scatter. In exchange for this advantage, we would have to deal with considerably more complexity in control and synchronization. We have chosen at this point to use the simpler method of one imaging path, and for the reasons above the excitation path is chosen to do sharp imaging while the emission path is assigned to collect as much light as possible.

5.4. Specific design

Once the general principles of the design were established, the exact positioning and selection of optical elements were optimized by means of a design program, where the degrees of freedom of the system were selected by the program's user and the results of the choices were displayed for scrutiny and adjustment. A sample output of the program is shown in Figure 5.7 (next page). A graphical ray-tracing program was written to assist in this design. Figure 5.8 is a diagram of the system's optics. The arc lamp is imaged through a condensing lens onto a pinhole -- 200 μm and 50 μm

Settable parameters:
retinal half-field = 12.00 degrees
input pupil = 2.00 mm
L1 to Mx = 160.00 mm
L1 F.L. = 33.33 mm
L2 F.L. = 80.60 mm
L2MAG (mx/my) = 0.86
pinhole diameter = 200.00 um
L3 F.L. = 51.70 mm
L3MAG (r2/ph) = 0.95
LF F.L. = 80.00 mm
Full pupil = 8.00 mm

	FL	diam.req.	DI	DO	RI	RO	MAG	f/#	angle
L1(ex.)	33.33	19.90	42.10	160.00	59.10	33.33	0.51	1.68	12.00
L1	33.33	25.90	42.10	160.00	59.10	33.33	0.51	1.29	12.00
L2(x)	80.60	33.37	149.92	174.32	276.59	113.75	2.43	2.42	6.90
L2(y)		16.55						4.87	4.16
L3	51.70	14.71	40.24		100.82	106.12	0.95	3.51	4.17
LF	80.00	30.40	160.00		126.67	217.13	1.71	2.63	9.98

	Tilt	Size
Mx	1.60	10.75
My	2.33	12.50

System characteristics:
 retinal spot diameter = 235.64 microns
 Lengths:
  pupil to Mx =   202.10 mm
  Mx to My =   324.24
  My to ph =   146.36
  Total =   672.70
Relative retinal spot power = 110.88

Figure 5.7. Optical design program output, and graphical drawing of the
           optical system.

-108-

Figure 5.8. Three-dimensional layout of the optical system.

pinholes have been most commonly used — which is conjugate to the subject's retina and thus determines retinal spot size. The arc image overfills the pinhole rather than exactly matching its size; as noted above, brightness is conserved (a consequence of the existence of the optical invariant), so that we may overfill the pinhole to any extent without losing illumination power as long as the cone coming out of the pinhole is sufficiently wide (the requirement is: pinhole diameter • tan (half-angle of cone leaving pinhole) $\geq$ retinal spot diameter • tan (half-angle of cone [in water] from excitation pupil to retina)).

The light path is turned vertical by mirror M_T so that the scan mirror orientations will be correct. L_3 establishes a beam of the proper size at the proper height for M_y, the Y-scanning mirror. A stop placed on L_3 establishes the size of this beam and thus the size of the cone at the subject's pupil.

After the beam passes M_y, it is being scanned in a vertical direction. L_2 images M_y onto M_x so that there will be one plane which is stationary in X and Y. The full raster is now created and turns from M_x toward the subject.

L_1 is a 33 mm F.L. Volk aspheric lens used for indirect ophthalmoscopy. It is placed so that a retinal conjugate plane (conjugate also to the arc lamp) is located at its focus. Light leaving L_1 is therefore in a parallel bundle, with the angle of the bundle to the centerline being determined by the current position in the scan raster. All of the bundles converge on the pupil of the eye (which is our stationary plane). For an emmetropic eye focused at infinity, the parallel beam is focused onto a sharp spot on the retina, which moves as the scan moves. The spot diameter is set by the nodal distance of the eye and the optical distances in the

system, and is equal to the pinhole diameter times 1.18.

Light fluorescing or reflecting from the retina heads back through the entire pupil, passes through L_1 and is directed through the emission field lens L_f, which is located just behind M_x. M_x plus an additional stop surrounding it block light which could come from lens fluorescence or specular reflection; the aperture in this stop defines the exit pupil of the emission path. Light passing the field lens is imaged by output lens L_o onto the photomultiplier surface, which is also in a plane conjugate to the pupil so that all parts of the scan are incident on the same area of the photocathode.

L_1 is located in both the excitation and emission paths, just as it is in indirect ophthalmoscopy. The two surfaces of L_1 form two small reflected images of the light source (effectively, M_x) near the photomultiplier. Small stops in the emission path block these reflections.

A filter wheel may be placed between L_2 and M_x on excitation, or L_f and L_o on emission, to provide the fluorescence excitation and emission filtering.

All lenses in the system (except L_1) are achromats; all are coated, and all except L_1 operate at f/#'s of 2.4 or greater. An earlier version of the scanner optics using simple lenses and smaller f/#'s proved to be disadvantageous not only because of poor image quality but also because that poor quality necessitated the placement of additional limiting apertures in the system which cut down the light level by a factor of 80 compared to the current system.

5.4.1. Focusing. Focusing adjustment is necessary for the pupil and retinal focus. The retinal focus is a function of how far L_1 is from the retinal plane -- 33 mm for emmetropic eyes. L_1 is moved until the subject sees a focused raster. Once this adjustment is made, it does not need resetting, and the pupil is focused by motion of the headholder.

5.4.2. Viewing. It is necessary to see where the scanner is targeted as the subject enters the headholder and positions his eye in the field. A qualitative measure of this is gained by knowing the angle at which the subject's eye is rotated and the position in space at which he sees the raster. For precise viewfinding, it is possible to observe directly by splitting the beam in the emission path and observing through an eyepiece. This does not work under fluorescence conditions, where the light is much too dim for the operator to see. It also requires diversion of some of the light output which otherwise would be available as signal, and forces the operator to be near the viewfinder eyepiece. We have chosen to use the photomultiplier output itself as a viewfinder. From the operator's position, a television display provides a representation of the scene seen by the PMT. Details of this display are found in the next chapter.

The current system provides satisfactory resolution, freedom from reflections from human or glass optical elements, safety, and correction for moderate ametropias. Brightness is more than adequate for reflectance pictures; still increased brightness could be useful for fluorescence work. The field half-angle of 15° is just sufficient to include the disc and macula in a single image if necessary.

The light presented to the photomultiplier during a scan is converted into charge pulses which must be accumulated, time-sliced and processed to

form a usable image. The electronics which perform these functions are described in the next chapter.

CHAPTER 6

SYSTEM ELECTRONICS

Once the light emitted from the eye strikes the photomultiplier cathode, it is converted into an electric charge pulse, and all further processing is done in the system electronics. Electronic subsystems are also involved in control of the scan itself, data display, storage and processing. Because of the combined demands for low light detection, high speed, and image processing, a number of special subsystems have been developed for the retina scanner. This chapter reviews some of the major special hardware elements of the system and explains how particular demands were addressed.

6.1. Tasks

The hardware and software of the scanner have several tasks:

1. The parameters of the scan raster are definable and adjustable so that the scan is controllable with respect to speed, position, size, and scan-line separation in space. This is necessary in order to allow movement to areas of interest in a larger field, to include or screen out particular areas, or to gain higher resolution for some small-size areas.

2. In order to get fluorescence pictures and compensating reflectance pictures without worrying about subject movement or metabolic change, rapid change of excitation and emission wavelengths by means of a filter wheel is made available.

3. The photomultiplier is switched on and off, and a shutter in front of it closes to protect the PMT when the room lights are on, as during patient repositioning.

4. The current pulses from the photomultiplier are amplified, passed through a thresholding discriminator which separates signal from dark current, and converted to digital pulses for further processing.

5. The pulses from the photomultiplier are collected, counted, partitioned into time slices which represent geographical divisions on the retina, and the resulting <u>brightness array</u> is passed to the main computer. The pulses are synchronized with the position of the scan so that the display corresponds precisely with eye geometry.

6. The central processor operates on the brightness array, processes it appropriately for display and for detection of FAD signals and other absorption/fluorescence as outlined in Chapters 2-3. Handling of multiple pictures, image processing, and other mathematical functions are performed by the CPU.

7. Display facilities are provided, both for displaying of the images (brightness arrays) before and after processing, and to serve as a viewfinder for the optical system as mentioned in Chapter 5. Graphical display for image processing purposes is also provided.

8. Storage of programs and images are provided, along with image parameters so that stored images may be recalled directly at any future time.

6.2. System Description

The overall system block diagram is shown in Figure 6.1. In the diagram can be seen the various subsystems such as the flying-spot scanner

Figure 6.1. System block diagram.

with its hybrid control circuitry which provides the variable raster patterns, and the filter wheel, also centrally controlled, to provide the variable wavelength control. The shutter and the "gate" prevent light from damaging the photomultiplier when necessary. The current pulses from the PMT reach the amplifier and discriminator and the resultant TTL-level pulses are collected by the photon-counter board, which does the time-

slicing described above. The brightness values are collected from this board by the computer, and used in subsequent image processing work.

Display of the PMT image in real-time is provided by the X-Y display, which monitors the PMT amplifier, while the processed digital image information is displayed on the graphics display in halftone form. The X-Y display and the photon counter board both receive synchronization pulses from the flying-spot scanner controllers.

6.2.1. Division of Labor. It was imperative in the overall system design to decide which operations could be controlled by the central processing unit, and which could be operated semi-independently as satellite controllers. If the processor had to continuously generate the scan pattern, count the data pulses, or poll safety status lines — each of which is a reasonable job for a microprocessor -- the overall workload would exceed the processor's capability, or at least leave it no room for user interaction. The philosophy with the individual subsystems in the retina scanner has been that those jobs which require high-rate operation (photon counting, scan pattern generation) or which demand continuous operation (safety monitoring, scan pattern generation again) should be capable of independent operation once their basic parameters are preset by the central processor. They can communicate with the CPU through polled lines and interrupts as necessary. Slower tasks, tasks requiring user interaction, and the like remain under central control.

6.3. Subsystems Description

6.3.1. Computer. The computer is an 8085-based "Matrix" development system developed at the MIT Biomedical Engineering Center for their line of microprocessor-based medical instruments. It supports 56K of memory and runs at a 2 MHz clock rate, with an assist from an AMD9511a companion mathematics processor chip to speed up computation. The most useful feature of the system is its expandability: 21 slots are provided in the backplane for regular (e.g. memory) and special purpose hardware, which can communicate through the usual synchronous data bus, or by interrupts or direct memory access. New designs might use a 68000 or 80286-based system for the same purpose, or one of the faster 4 or 6 MHz 8085 systems.

6.3.2. Flying-spot scanner interface. The scanner interface card (Figure 6.2) provides a means for the computer to control the motion of up to three scanning mirrors. The interface is designed to provide raster scanning as well as random point-to-point scanning of the mirrors.

The basic philosophy of the raster scan is that a triangle-wave (or sawtooth) drives the X-motor back and forth, while a Y-position driver is stepped through a preset sequence of line positions, changing lines each time the triangle wave returns to its leftmost position. There is one triangle-generator section and two selectable-point generator sections on the scanner interface card.

The triangle generator circuit [68] uses an Intersil 8038 function generator to generate a triangle wave or a sawtooth wave (depending on the setting of the balance controls) at frequencies ranging from 10 Hz to 3000 Hz. There are three D/A converters, which respectively set the rate, amplitude (gain) and offset of the triangle pattern. DAC1 provides an input voltage level to the 8038's voltage-controlled oscillator; the 8038 creates a triangle-wave output which is multiplied by DAC2 to provide gain control,

-118-

SCAN PATTERN GENERATOR
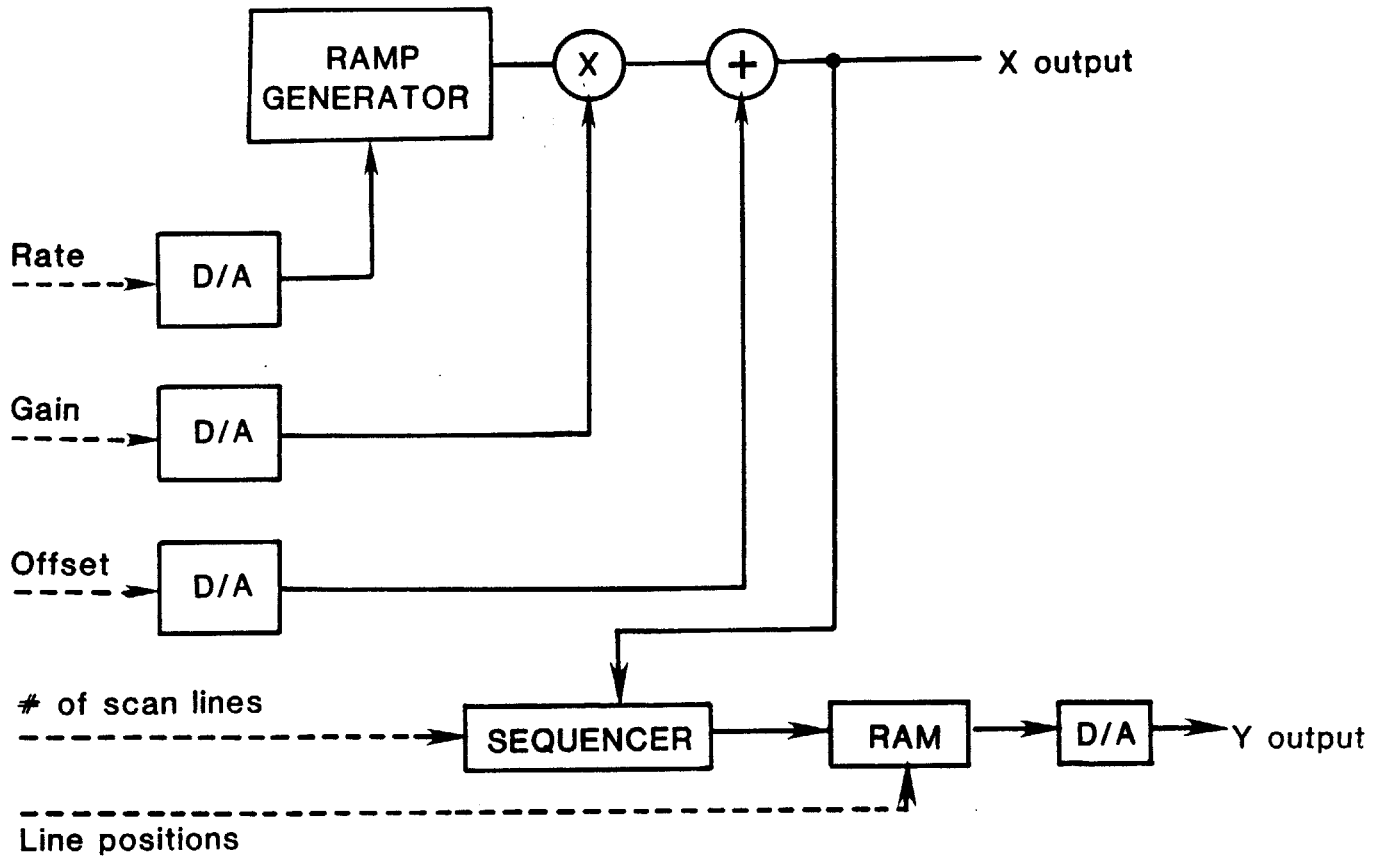
Figure 6.2. Block diagram of the scan interface card.

and the current output of this DAC is summed with the output of DAC3 to provide left-right offset adjustment.

The gain control is not symmetric: since the 8038 puts out a triangle wave between -5 and -10 V, the gain multiplication produces a signal whose range is at most [-5,-10] and at least [0,0]; thus the offset must be adjusted with each gain change if the scan is to stay centered in the same x-position. The software which presets these DACs automatically performs

this adjustment (see Chapter 7).

The position-sequence generator uses 2101A-2 static RAMs to store a sequence of up to 256 numbers corresponding to a sequence of angles (from a choice of 256 evenly-spaced angles) at which M_y is to be placed. For example, it may store a sequence of evenly spaced numbers to produce an evenly spaced vertical raster sweep, or a single line scan could be produced. A separate register marks the number of memory locations which are to be sequenced through before starting over again. Another register allows the central computer to set the sequencer to the last line in the sequence, so that the next change of line will cause the scanner motors to move to the top line. This feature is used to ensure that all data recording starts on the top line.

Only 256 bytes of address space are allocated by the computer for each backplane slot. In order to allow up to a 256-line scan, and leave room for the control registers of the scan interface card, all writes to the 2101A-2's are preceded by a special prefix instruction which sets two latches on the card to disable the control registers and enable the memory write line for the next cycle.

The change to the next position in the line sequence can be triggered directly from the computer -- thus the scan could be moved to selected points at will -- or it can be set to trigger whenever the triangle-wave generator reaches its lowest (leftmost) point. This is implemented by edge-triggering on the square-wave output of the 8038, which is the derivative of the triangle wave.

The output of the memory is applied to DAC4 which produces the appropriate voltage to drive the Y-mirror.

The role of the central processor in the scan interface is to set the

four DACs and write the appropriate line position sequence to the onboard memories. Once these parameters are set, the triangle wave and sequencer work independently to continuously generate the desired scan pattern.

6.3.3. Scanner amplifiers. The voltage signals provided by the scan interface card are used as input to the scan amplifiers (General Scanning A-601). These are voltage-to-current converters to drive the galvanometer scanner mirrors (General Scanning G-100PD). The galvanometers provide position feedback signals which are used by the amplifiers to improve their response through use of a closed-loop amplifier. (Without this improvement, the natural pole of the scanner mirrors at 300-400 Hz would prevent our using them at rates up to 1 kHz triangle wave); they are also used to provide synchronization for the X-Y display and photon counter board (see below).

The final scan line of any standard raster is occurring while the Y-mirror is moving from its bottom-most to topmost line. This scan line is not horizontal but appears as a retrace line coming across the rest of the raster. The operator may choose to ignore this line in the data analysis.

6.3.4. Filter wheel. The filter wheel is a disc which can hold up to four 2" square filters at 90° intervals. It is driven by a stepper motor with sufficient torque to change filters in 0.3 second.

The control of the filter wheel is directly from the computer. Two cross-coupled TTL flip-flops provide the appropriate pulse sequence to the output amplifiers which provide current to the stepper motor. The computer provides one pulse at a time by individual command. A separate register on the filter wheel control board changes the sense of the flip-flop inputs and thus provides choice of forward and backward motion.

6.3.5. Shutter and gate. A 25mm-aperture shutter (Vincent Associates, Inc.) is located in front of the PMT to keep room light from damaging the cathode. It requires a +5V. input to keep the shutter open, 0V. to keep it closed. Control of the shutter is available in two ways: a main switch on the optical table can close the shutter regardless of the computer; when that switch is set to AUTO, the input to the shutter comes from a simple latch which is computer-set as the operator desires.

The PMT is also protected by a gating circuit (supplied with the PMT). The gate is an optocoupler which controls the power supply between the photocathode and first dynode. When 5V. is applied to the gate input, power is applied normally; when no current is applied, the cathode-dynode voltage is reverse-biased approximately 100V, which reduces the PMT sensitivity by a factor of 100. The gate alone is not enough to protect the PMT from room light at full power, however. The shutter and power supply together provide adequate PMT protection.

6.3.6. Amplifier/Discriminator. The amplifier/discriminator (Pacific Precision Model AD6) has two stages. The first stage is a high-gain high-speed current-to-voltage amplifier which converts a 10^7-electron pulse (the typical gain of the 9814A PMT) into a 150 mv pulse lasting 6 nsec. An internal control sets the pulse width. The amplifier has a direct output, which we have coupled to the X-Y display to provide the intensity signal, and which is also fed into the discriminator.

The discriminator uses an ultraspeed comparator (Am615) to check the incoming voltage peaks against the threshold voltage specified by the user. Those peaks which pass this stage are conditioned by ECL logic to provide an ECL output and a TTL output pulse train at a maximum rate of approximately 125 MHz.

6.3.7. Photon Counter Board. (Fig. 6.3.)

COUNTER/TIMER
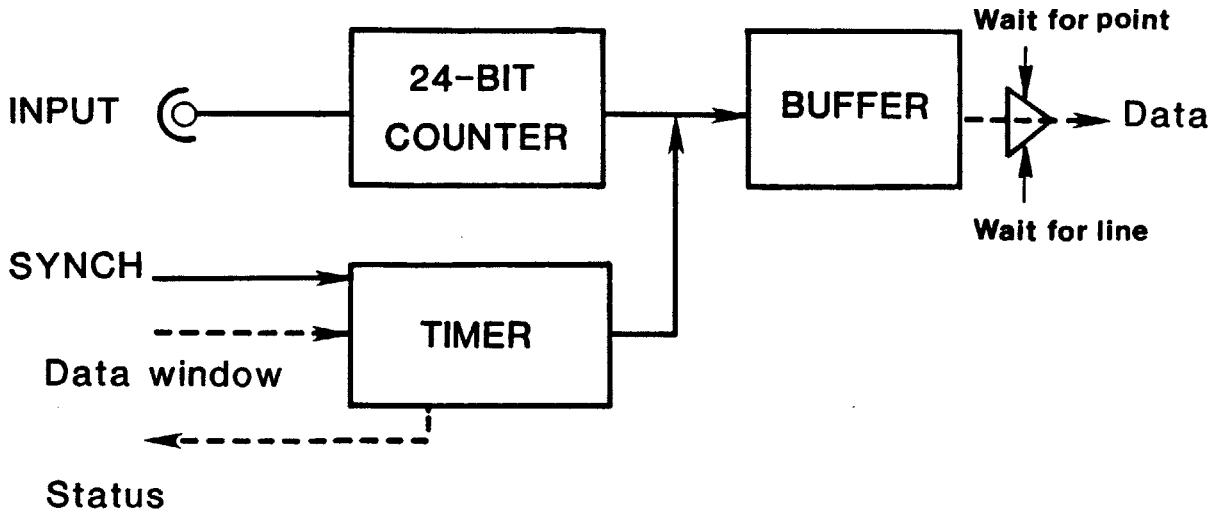
Figure 6.3. Block diagram of the photon counter board.

The photon counter board is probably the most complex single electronic subsystem in terms of the number of different functions it must perform. It is based around the Intel 8253-5 counter-timer chip, which can be programmed to produce a train of 500 ns pulses with any desired period from 1 μs to 32.768 ms in 500 ns increments. A gate input provided on the chip allows the pulse train to begin at a specified time.

Between pulses, the board is accumulating photon counts into its counter chain. The counter chain is a 24-bit counter in parallel with a 24-bit latch. The first (least significant) four bits of counting are implemented with 74S112 flipflops, which are guaranteed for 80 MHz count

rate (125 MHz typical); the rest of the chain is made up of 74LS161 and 74LS691 synchronous counters which are rated to handle 20 MHz. (80 million photons/sec corresponds to 30 nanowatts at 520 nm wavelength.)

When a timer pulse occurs, the contents of all counters are transferred to the corresponding latches immediately. Following a short (100 ns) delay the counters are cleared and counting continues from zero. The saved count is the number of photons which were recorded during the previous datapoint's time-slice. Counting thus continues without any long pause for the computer to read the data, which it does by accessing three 8-bit memory-mapped registers.

In order for the program loop which reads in the brightness array to be as short as possible, the programmer may add a special bit to the memory address when he wants to read from the latches. When this bit is set the photon counter board holds down the READY ($\overline{WAIT}$) line of the microprocessor to stop execution until the next pulse from the 8253-5 timer, at which time (after a delay of at least 250 ns, by which the counter data is safely stored in the latches) the instruction which had the special bit set will finally complete, with the correct data returned. Another address bit is set when the program is reading from the last 8-bit register it wishes to read for a given datapoint; this address bit clears the wait-for-datapoint-done flipflop so that the cycle may repeat, with the program waiting on a read instruction until the next datapoint is complete. (Figure 6.4). Two status lamps on the photon counter board edge can be set up so that they light when there is a data overrun (two datapoints were transferred to the latch before the first one was read out) or data overflow (magnitude of number of photons exceeds 24 bits in one point-time).

Counting of the first point in each scan line must begin at precisely
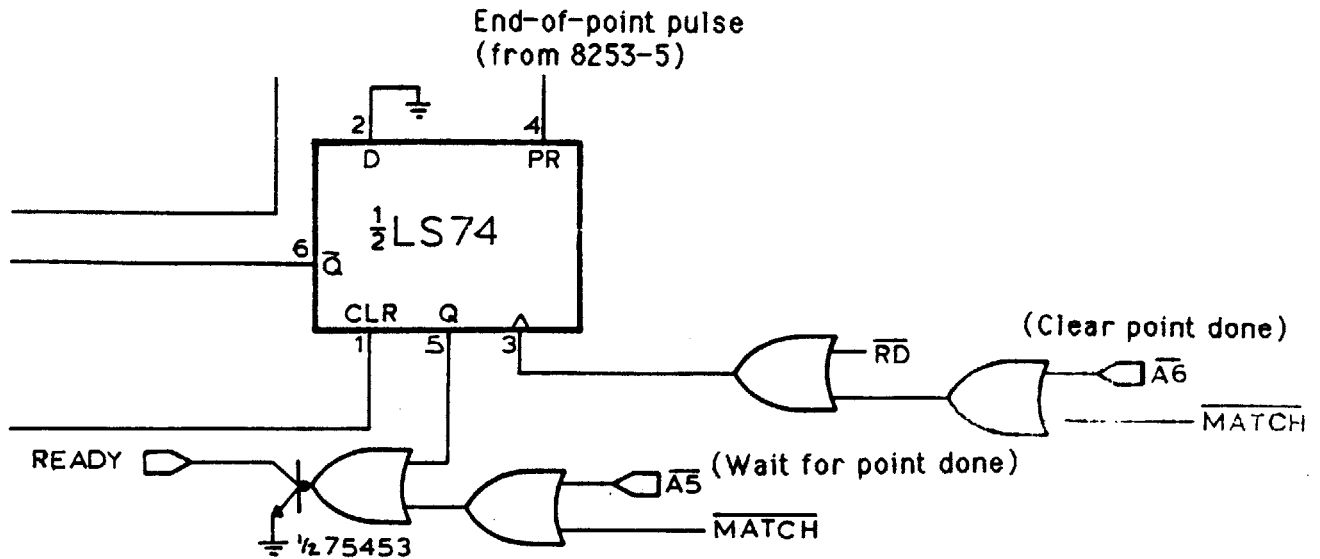
Figure 6.4. The wait-for-point-done and clear-point-done section of the photon counter board.

the same time after the scan line begins; otherwise the different rows of the image will be shifted one from another in X. In order to achieve this synchronization, a velocity-feedback signal from the X-scan amplifiers is brought in to the photon counter board. The signal is filtered to remove high-frequency noise not due to actual mirror motion, and a zero-crossing detector (with hysteresis to prevent multiple sync pulses due to noise) provides a pulse each time the scanner motor begins its rightward motion. The pulse (labeled VFN) then performs three jobs:

1. By using another of the spare address bits, the processor can tell the board to hold the READY line low until the beginning of the

next scan line. The VFN signal presets a flipflop which releases
this line. Thus when the data-recording time on a scan line is
less than the time for the scan line itself, the processor can
wait until the next line begins before reading more data (use of
the READY line instead of a program loop allows much less latency
between the actual start of the line and the processor's readiness
to take data -- important because we would like to be able to have
very short data-point times without overruns).

2. The VFN signal is applied to the gate of the 8253-5 timer, syn-
chronizing the start of the first timing interval.

3. To gain more precise synchronization, the pulse is also used to
clear a two-bit divider (74LS112) which generates the 500-ns clock
pulses for the 8253-5 from a base 8 MHz (125-ns) clock. Thus all
scan line data-collection begins within 125 nsec of each other,
relative to the physical start of the scan line.

In the description of the key algorithms in the next chapter the
sequence by which a program would be written to obtain an image of desired
size and resolution from the photon counter board is explained.

6.3.8. X-Y display/Viewfinder. The need for a viewfinder became clear
early on, due to the difficulties of precisely positioning a subject in the
darkened room with low illumination light. Since the scanner does not work
with a standard raster, a normal television could not be used for instan-
taneous display. The best available solution is a display which can be
driven in X and Y directly so that the intensity of a point seen by the PMT
is referenced to its position in the raster.

We acquired an X-Y display (QSI Industries) originally designed for
use as an arcade-game display monitor; this display is considerably less

Figure 6.5. The synchronization of the data collection.

expensive than commercial laboratory X-Y displays. The display accepts a ±10 volt X or Y input for full-screen spread, and, while it is designed for TTL levels, the Z-axis (intensity) can actually be driven in roughly linear fashion between 0.8 and 4.0 V.

Three amplifier trains, one for each input, were built to adapt the scanner to the X-Y display. The X and Y inputs come directly from the X and Y position feedback signals of the scan mirrors. In order to keep the display centered on the screen when we shift our scan offset in X and Y,
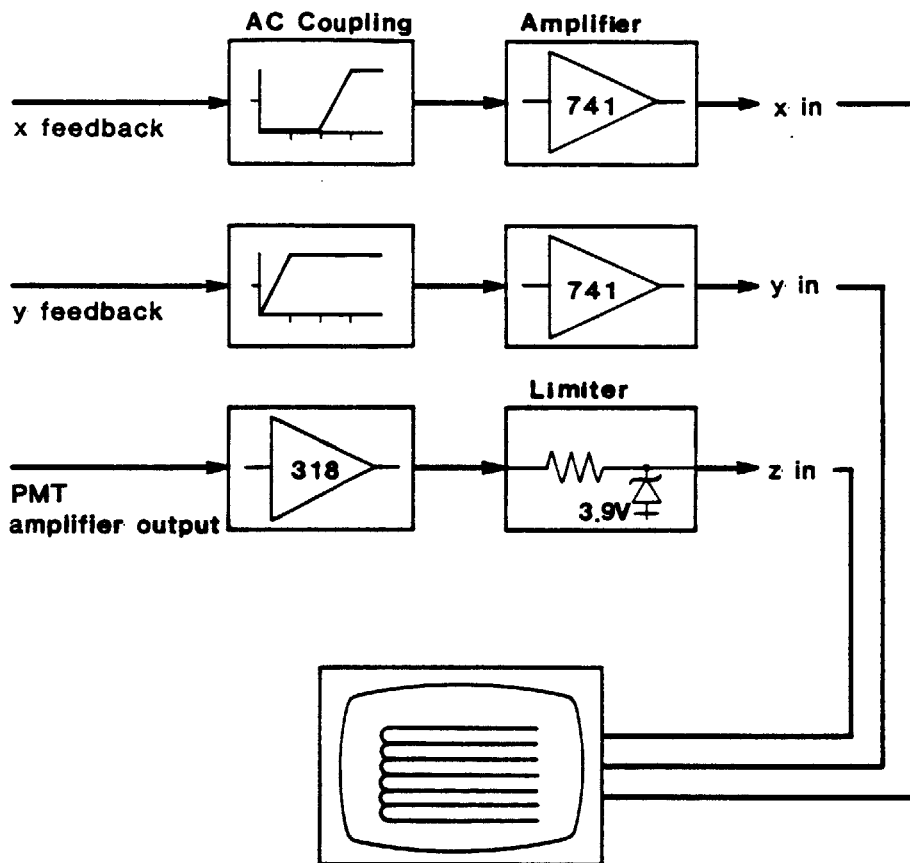
Figure 6.6. X–Y display operation.

both channels are AC coupled with the pole in the response low enough in frequency to account for the actual slow variations of the scanners during a scan. Gain controls allow the display to stay the same size when the size of the scan changes.

The display contains "spot killer" circuitry which protects it against phosphor burnout by blanking out the display when either X or Y motion is too small. Because of this, slow or narrow scans may blank out for part of their duration. The value of the spot killer threshold has been adjusted down to remove this annoyance while still retaining protection for the display.

The Z input comes directly from the PMT amplifier. Gain and offset

controls keep it in the desired range; the output is limited by a Zener diode in any event to remove a potential problem whereby if the display momentarily became excessively bright, some of its light could couple back into the PMT which would make the display brighter still.

The display is coarse but is more than adequate for its task.

6.3.9. Graphics/Digital Display. The digital display comes from a master's thesis project by Fredrikson[69]. It is a 1024 x 1024 graphics display, with bit memory. The display is written by the processor, 1 or 8 bits at a time. Because the display does not provide grey levels, a half-toning algorithm (see Chap. 8) has been developed which provides "continuous-tone" display.

6.3.10. Disc Controller. Because no disc controller was available for the Matrix system, one was developed based on the Western Digital FD1771 controller. Like the photon counter board, it uses READY circuitry to speed up tight program loops to rates adequate for the data transfer. Two double-sided Shugart SA851 floppy disc drives are managed by the con-troller.

6.3.11. DMA Card. Direct memory access has the potential for further increasing transfer speed and in particular to allow even shorter data-point times than the current polling-loop scheme allows. A general DMA controller based on the Zilog Z80A-DMA has been designed and is available for future implementation.

CHAPTER 7

ALGORITHMS FOR SYSTEM CONTROL AND IMAGE ACQUISITION

This chapter describes the basic algorithms and programs which are used for control of the hardware and optics of the scanner, including a very brief description of the Stoic operating system. The individual programs are primarily important as they relate to the whole image acquisition process, which is summarized at the end of the chapter.

7.1. General capabilities

The use of application software of this type is to allow the experimenter full control over the system parameters necessary to perform and vary the experiments -- in general, to record eye images under different conditions. Control is available from the console of the retina scanner of:

1. The size, rate, number of lines, and X-Y position of the scan;

2. The shutter;

3. The length of time (space) given to an individual pixel [picture element], and the dimension of the image array in X and Y;

4. Initialization of acquisition of an image or a sequence of images;

5. Changes in the filters used;

6. Storage on disc of any picture, raw or processed;

7. Image processing which prepares the raw data for display, storage, enhancement, registration with and combination with other images (See Chapter 8).

8. Utilities which allow easy debugging and correcting of programs, comments and annotations.

7.2. STOIC

The operating system and standard language used for programs running on the Matrix computer is STOIC (Stack-Oriented Interactive Compiler), developed at the Biomedical Engineering Center as an outgrowth of FORTH. STOIC is a _threaded-code_ system; that is, programs are written as collections of macro names ("words") which are found elsewhere in core and invoked in sequence. These macros may themselves be collections of other names, and so on recursively; to execute a word either singly (by typing its name) or as part of a group of words which comprise a larger word is to invoke all of its component words, which may themselves have other components or which may run short assembly-language sequences. The prime advantages of threaded code are small memory requirements for new programs (words) and the ability to build up large words by using smaller constructions. A penalty in speed is paid because of the necessity of managing all of the pointers to the different levels of invocation whenever a new word is called — thus a STOIC program would take as long as an assembly program plus 100 μs times the number of STOIC words in total that are invoked; but the space used to define the program will be much less.

STOIC uses a stack for storing arguments to any word; the word may take them from the stack and may push others on the stack to be left for the next word. Thus the + operator takes the top two items off of the stack, adds them, and returns the result to the top of the stack. Because of the stack, STOIC constructions are worded in reverse-Polish form. In order to add 2 and 3 and print the result one would type

2 3 + =

Arguments may be kept on the stack, or they may be saved as variables or constants for later use. Conditional and loop constructions are part of the language.

New words are defined by typing the proper sequence of words and assigning them to a new name. For example,

'ADD23 : 2 3 + = ;

will always add 2+3 and print the result whenever ADD23 is typed. More practically,

'AVERAGE : + 2/ ;

takes the average of the top two numbers on the stack by adding them and dividing by two. Typing

5 17 AVERAGE

would find the average (i.e., 11) and store it on top of the stack, where the next word could use it.

STOIC provides its own disk operating system to create, modify, and read files from the disk. Along with text and data files, two special types of files are used to store a large set of words for particular applications. Once a set of words has been defined and is in core memory, a core-image or a bootstrap file could be created, respectively, by:

'.FILE1 WRCI

'%FILE1 WRITE-BOOT

The file created by WRITE-BOOT can be invoked whenever the system has just been reset (as on power-up); the core-image file created by WRCI can be invoked anytime from STOIC (by the command '.FILE1 RDCI). Both files save all word definitions that were known in core at the time they were written, and when they are invoked those definitions will become the currently-known set. Thus, we have three core image files .ED, .FILT5, and .IP which con-

tain commands for text editing, data acquisition, and image processing routines. Invocation of a core-image file is similar to calling up a new menu of available commands.

7.3. Memory organization

The memory map of the system when running .FILT5 (image acquisition routines) is diagramed in the table below.

Address (decimal)	Contents ''
0	ROM containing bootstrap, disk routines
2048	Low-level assembly-language debugger
4096	Lowest available user location
	Basic STOIC words
18084	Retina scanner data acquisition words, scanner control words, tables, frequently-used IP programs
29000	Space for pixel buffers, normalizer buffers, editor and image workspace
57344	Top of memory (0E000 hex)
	I/O page -- addresses of I/O registers
61440	Unused
65536	Top of address space

The system memory available to the user runs from 1000 to 0E000 hex (4096 to 57344 decimal). In this space will be found all of the defined words (called the dictionary), both those known to basic STOIC and those particular to our research work. All of the buffers which store and process images are found in this space, after the defined words; their sizes are fixed when the core image file is invoked each time (see ALLOCATE, below). In the range 0E000-0F000 hex (57344 to 61440) are found all of the addresses which correspond to the data registers of the photon counter board, digital display, and other hardware. 4K at the top of the address space remains unused primarily for historical reasons; the I/O addresses could be shifted to that space if necessary, giving a continuous stretch of memory to 60K.

7.4. Algorithms of the Retina Scanner

In the descriptions below, a name followed by two numbers in parentheses is the name of the word and the number of arguments pulled off and left on the stack. For example, RR(1/0) denotes that the word RR takes

one argument off the stack and leaves none on -- so the stack pointer is two bytes (one argument) higher on completion than it was on invocation. Also, TOS denotes the last argument to the command (which is on top of the stack when the command is invoked) and NOS refers to the next-to-last argument.

7.4.1. Scan pattern control.

The software controlling the scanner mirrors needs only to set the various DACs in the scanner interface hardware, as well as the Y-position memory elements. A small amount of work must be done to make the commands comfortable for the experimenter to use. As mentioned before, the position of the center of the raster (in X) will change when just the Gain DAC is adjusted. Ideally, however, the rate, gain, and offset commands should affect only one parameter of the scan apiece. Thus for both the gain and offset commands, two DACs are modified to compensate for the hardware and effectively create independent control.

Since the hardware registers are write-only, shadow registers in core are maintained so the experimenter knows where the current settings are. The commands RR(1/0), GG(1/0) and OO(1/0) adjust the rate, horizontal gain, and horizontal offset of the (center of the) scan, respectively. All arguments range from 0 to 256 although mechanical limitations actually cut these ranges down some -- for example, the maximum practical rate with our particular scanner motors and mirror loads is 1440 Hz, which is brought on with 128 RR. For loading the Y-position memories, the LINES(1/0) command sets the sequence-count register with the appropriate number of lines. The basic command MEMLD(2/0) loads the Y-position RAM address TOS with the value NOS -- a loop of such commands would be required to set all the lines of an entire scan. For ease of use, the HH(1/0) and SPACING(1/0) commands

calculate and perform such loops automatically, and allow the user to set the center height and inter-line spacing of the raster, respectively.

The file 'DFAULT, which is a file of STOIC commands as one would type them in, is automatically called when the '.FILT5 core image is invoked. Among the default commands included in this file are commands setting the initial raster to a pattern found to be useful for most work.

7.4.2. Filter wheel and shutter. These are controlled by writing directly to their registers. The commands ONSH(0/0) and OFFSH(0/0) simply write 1 or 0 into the shutter's control register to open and close it (as long as its master switch is set to AUTO). For the filter wheel, the words FORWARD(0/0) and BACK(0/0) change the hardware direction register to set the next step in one direction or the other. The actual stepping is done by the word PULSE(0/0), which writes into the stepping register of the driver board. The word STEPS(1/0) is the most commonly used version and generates the number of PULSEs requested by its argument. 50 steps is one quarter turn, or the distance between two adjacent filters. The variable MRATE is set to the length of time to pause between successive PULSEs, for fast or slow filter changes.

7.4.3. Photon Counter Board. Particular attention is to be devoted to the basic routine ADDD1, which forms the core of the two major data acquisition routines GETDATA and ADDDATA. The main demands of this program are that the computer be ready for a new scan line, receive only as many data points as are requested per scan line, and read successive datapoints as fast as possible, for the minimum data-point collection time is limited by the tight data-reading loop.

The relationship between the number of data points to be collected per

line, XPTS, the datapoint time PTTIME, and the scan rate as set by RR are demonstrated in Figure 7.1.

GG

spacing

PT TIME

XPTS=15

HH

OO

X=O

Point size ∝ GG×RR×PT TIME

Y=O

Figure 7.1. Relationship between various scan parameters.

If we are taking data only on one half of each scan line (i.e. while it is traveling in one direction only), then we want to make sure that

(72) $2 \bullet \text{PTTIME} \bullet \text{XPTS} \leq \frac{1}{2} \bullet$ scan-line time(μs)

(the factor of 2 on the left is because the variable PTTIME is set in 0.5 μs increments.). The word CHK checks this computation and also makes sure that the value of XPTS • YPTS is not greater than the buffer size.

The system user sets the parameters XPTS, YPTS, PTTIME and calls GET-DATA, which clears the working buffer, or ADDDATA, which does not. The working buffer is pointed to by the variable PCPTR. By default, for single-buffer, single-picture chores, PCPTR points to the start of the photon-counter-buffer space (PC-buffer space) which was allocated when the

core image was invoked.

The _ADDD1_ routine is invoked next. This algorithm, written in assembly language for speed, involves four main steps:

1. It writes to the SCAN RESET register on the scan interface board, forcing the scan to the last line so the next scan line will be the correct place to start taking data.

2. It then executes an instruction which will not return until the VFN pulse arrives to indicate that the mirrors have reached the start of the first scan line.

3. Once on the scan line, a tight program loop is entered which waits for the next data point to be ready, reads in the low and middle 8 bits of the data, clears the data-point ready flag, and loops if we have not already taken XPTS datapoints on this line. The data read in is added to the two words pointed to by the 8085's HL register pair. At a 2 MHz clock rate for the 8085, this loop runs in 38 μs; if PTTIME is any less than (twice) this figure, data will eventually be lost (overrun) because the computer is not ready for subsequent data points quickly enough.

4. The program then returns to step 2 to wait for the next line, until all of the lines have written their data into the computer memory.

As currently configured, this maximum-speed routine leaves the array in correct order but with the high byte and low byte of each 16-bit word reversed. _SWAPBYTES_ is run after the GETDATA's and ADDDATA's are completed to restore the correct orientation.

The companion program _SAVEDATA_ will start from the address pointed to by PCPTR and write XPTS • YPTS 16-bit quantities to the disc. It assigns

file names sequentially; all names begin with the current date.

7.4.4. Status. Status of the system parameters is obtained at any time from the words CS and SS, which give information about the photon counter board parameters and the scan interface card parameters, respectively. Their output is shown in Figure 7.2.

```
0/0) CS
DATA ARRAY: 75 (X) BY 40 (Y)
          (NOTE: 5500 POINTS MAXIMUM!)
TIME : 80 X 0.5 US

0/0) SS
RATE = 64
GAIN = 130
OFFSET = 64
# LINES = 40
HH = 128
SPACING = 2

0/0)
```

Figure 7.2. Output from CS and SS.

7.4.5. Buffer allocation. The program ALLOCATE is run whenever the .FILT5 package is invoked, in order to establish buffer space for images and display memory. For a picture of NPTS points, 2*NPTS bytes are required for the PC-buffer and NPTS bytes are required for the N-buffer ("normalized" buffer, see Chapter 8) which stores values for display and enhancement. In a given experiment, several PC-buffers may be required; usually only one N-buffer is necessary. The experimenter is asked to specify how many PC-buffers are needed and the maximum value of NPTS

planned for that day's activities. This value cannot be so high that when the buffers are allocated there will be less than 2K bytes of memory remaining (this value is chosen to save space for small STOIC programs which often are written after the buffers are in place). ALLOCATE will create the buffers as specified and set the value of MAXPTS so that the user does not set XPTS • YPTS to be larger than this value.

7.5. Running an Experiment

The sequence used in running an experiment is as follows:

1. The arc lamp and computer are turned on. The .FILT5 package is invoked (thus reading in the starting parameters from DFAULT) and the scanner motors and other equipment is then turned on.

2. The subject places himself in the headholder and the retinal focus is adjusted until a sharp raster pattern is seen (it is often useful to spread out the raster lines to 3 SPACING to facilitate this).

3. The desired XPTS, YPTS, scan parameters (if different than the default) are set.

4. The shutter is opened and voltage applied to the photomultiplier. Now an image should appear on the X-Y display; the subject's head is positioned so that the stationary plane appears to line up with the subject's pupil; further fine adjustments of the scan position or the headholder are made until a sharp image without vignetting appears on the X-Y display.

5. The system and subject are now ready for image recording. Usually a single word is used which includes commands to blank out the

console screen and the digital TV (ERASE and TVERASE), adjust the rate if necessary, run GETDATA/ ADDDATA as many times as desired, perhaps including some filter changes, then restore the TV display for further processing.

7.6. Utilities

A number of general utility programs for use in debugging, program generation and record keeping are briefly described here:

7.6.1. Editor. A displaying editor similar in style to EMACS has been written for the Matrix computer. Macro-commands in STOIC as well as single-character EMACS-like commands perform a range of editing functions.

7.6.2. DISCO. It is easy in STOIC to keep writing newer and newer words to fit the current situation until a word which works well has been written, but it is difficult to go back and find out what that word actually consists of, if it was typed directly into the dictionary the way AVERAGE was in the example above. DISCO is a STOIC discompiler, i.e. given the name of a STOIC word it will search through the dictionary looking up all of the component words and print their names. DISCO is also useful for debugging an errant STOIC word.

7.6.3. TRACE. The TRACE utility allows one to step through a STOIC program one component word at a time. Any information desired may be printed out after each word is invoked; the standard version prints the name of the next word and the current contents of the stack. This is usually enough information to reveal where an error is located. TRACE is made to turn off whenever the system reset is pressed, so that long loops being

traced will not go on indefinitely.

7.6.4. BIGDIR. BIGDIR provides a facility for maintaining a system master directory -- a directory of all files on all floppy disks. With the BIGDIR package loaded, the user inserts a disk whose contents have been recently changed and types !UPDATE; the contents of that disk directory are then updated in the master, which may be searched at any time to find on which disk a file is located.

7.6.5. ANNOTATE. ANNOTATE stores the current system parameters (the output of CS and SS) in a file. This permits the experimenter to automatically record the size, number of pixels, etc., pertaining to a series of pictures, for later retrieval.

7.6.6. %TERM. %TERM is a bootstrap file (actually written in assembly language) which serves as a terminal-emulator program so that the Matrix may upload and download files directly from larger systems.

7.6.7. Unix-stoic routines. The main computer at the Biomedical Engineering Center is a PDP 11/44 running Unix. Several programs (in the c language) have been developed to allow the 11/44 to read and write floppy disks in STOIC filesystem format. This is a faster method of uploading/downloading files between the Matrix and this particular computer.

CHAPTER 8

IMAGE PROCESSING

The image is now stored, in its raw form, in the computer memory; this data represents the actual fluorescence or reflectance intensity returning from the eye. We already know that a number of operations must be performed on the data to extract information relevant to oxygenation, including combination with other pictures according to the formulas expressed in Chapter 3. Pictures that are to be combined will also need preprocessing, mainly to ensure that two data points being combined refer to the same point on the retina, even when patient movement and other factors change the geography of the images.

In fact, there are a number of areas in which we have discovered the importance of image processing algorithms to extract special information from the raw data, remove noise sources, and present the current data in a manner useful to the experimenter. This chapter deals with these algorithms, in the general sense of how to solve a particular image-processing problem, and in the specific sense of the actual methods chosen for the retina scanner.

The image processing procedures described herein can be grouped into three categories:

1. Algorithms which allow direct display of the image recorded by the photomultiplier, and which display other information important to evaluation of the image quality and suitability for further processing.

2. Algorithms which highlight various features of the image and which attempt to remove some types of noise. This includes the set of 3x3 local mask operations for image transformation.

3. Algorithms which are used for bringing separate images in register with each other and operations which combine points from two or more images.

8.1. Histogramming and Display Processing

The first thing the experimenter generally wants to do with the image he has recorded is to see it directly on his screen, so that he knows if image quality is satisfactory (e.g. no motion or blinking) and also to look for structures in very dim light conditions which may not be visible on the X-Y display.

The contrast provided by raw data may be too small to allow excellent viewing of the detail and structures of the eye, usually because of a constant baseline level of illumination due to FMN, scatter, etc. Furthermore, the digital display will only produce graphics, i.e., black dots and white dots. In order to get good gray-level image display on this monitor, we need to resort to halftones. And for the best visualization of the structures, we need contrast enhancement methods to produce adjusted brightness values for the halftoner.

8.1.1. Tone Scale Adjustment Using Histograms. The contrast of a displayed image is adjusted by changing the brightness values, or tones, of the individual points which make up the image. This is most easily seen if we consider the histogram of the brightness of all pixels in the image (Figure 8.1). If we alter the brightness levels of all the points, in a smooth manner so that the sense of darker and lighter is maintained, we change the overall contrast of the final image. If we spread the tone scale out, then each image point becomes farther removed in brightness from
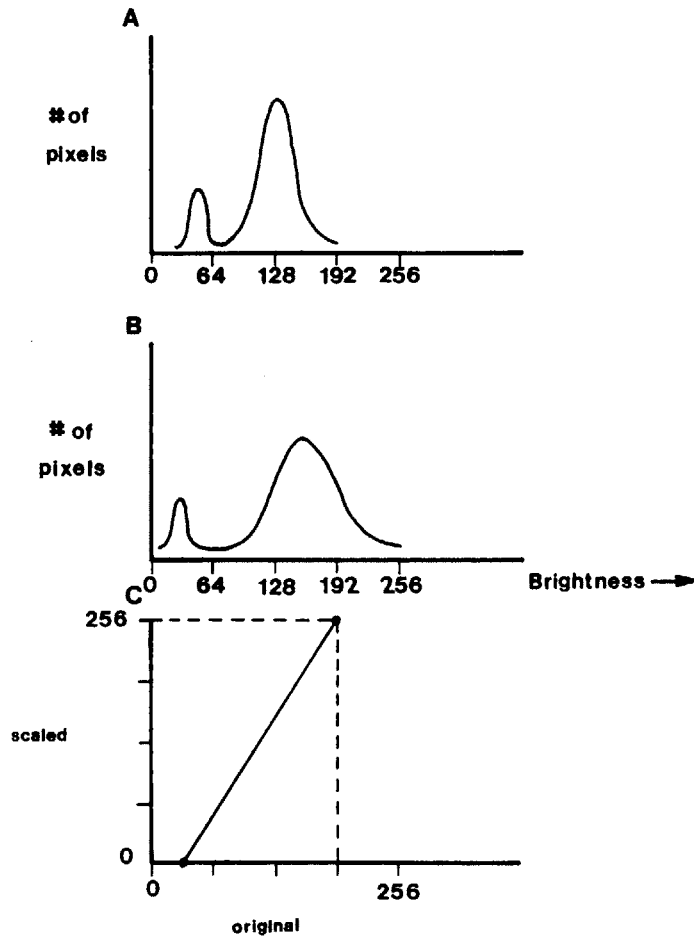
Figure 8.1. A. Brightness histogram of a hypothetical image. B. The same histogram with the brightnesses smoothly changed (tone scale transform) provide a higher contrast image. C. The tone scale transformation.

the overall average brightness, and the picture appears to have higher contrast. Compressing the tone scale yields less contrast.

Our raw data images usually have a range of brightness variation superimposed on a constant average level of brightness. To get the highest possible range of contrast, and to normalize the brightness values to the range 0-255 (the range of brightnesses known to the halftone process), we perform a linear contrast enhancement[70]. This involves selecting levels of brightness which represent the darkest (NMIN) and brightest (NMAX) tones of interest, and linearly adjusting all of the pixels so that pixels whose

brightness was NMIN become 0 (black), pixels whose brightness was NMAX become 255(white), and all pixels falling between these two levels in brightness map into tones between 0 and 255. The formula is

(73) $B_N = (B_P - NMIN) \cdot \dfrac{255}{NMAX - NMIN}$

where B_P is the brightness of the point in the raw photon-counter ("PC") buffer and B_N is the contrast-enhanced, normalized tone which is placed in the normalized ("N") buffer. Pixels in the PC-buffer which are brighter than NMAX also become white; pixels darker than NMIN become 0 or black. The transformation is shown in Figure 8.2.

Figure 8.2. Linear contrast enhancement and normalization.

An automated means of finding the values of NMIN and NMAX speeds processing and removes the necessity of manually examining all raw data to select these values. With linear contrast enhancement, the most straightforward method is an automated ends-in search. It is assumed that the

-146-

interesting region of brightness is the region in which the large majority of (raw) pixels is found, and that there is a small percentage of pixels which for statistical or other reasons may be quite far in brightness from the desired range. An ends-in search operates by removing a fixed fraction of the total number of pixels from each end of the brightness histogram (1% has worked best for us) and setting NMIN and NMAX to be the brightness of the points at the bottom of the 1st and 99th percentile of brightness, respectively. This selection process is shown in Figure 8.2.

One problem that we face with this scheme is that there is sometimes a significant fraction of pixels with raw brightness of 0 — due to the points being so bright that they produce a steady stream of charge from the photomultiplier, thus saturating the amplifier/discriminator which puts out only a constant level instead of a stream of pulses. The significant number of pixels at 0 brightness forces NMIN to 0 when the rest of the data may be far away in brightness. Modification of the histogram process so as not to consider points of brightness zero helps make the automated search hit the desired value more reliably.

There are other means of processing histograms for contrast enlargement. Histogram equalization[71] works on the principle that the more points that are found in a given range of raw brightness, the more interesting is that range and the wider the contrast should be within that range. If our allowable N-buffer brightness range were 100 tones, i.e. 0 to 99, then the raw picture would be histogrammed (or just sorted by brightness) and those pixels in the 0th percentile would be assigned brightness 0; in the 50th percentile, brightness 50; in the 99th percentile, 99. Similarly for 256 allowable tones in the N-buffer, we convert the brightest 1/256th of the raw pixels to 255, and so on down. Histogram

of
pixels

0   256   B ➡

of
pixels

0   256   B ➡

Figure 8.3. Histogram equalization. A. The original brightness histogram,
with lines drawn to show how the transformed brightnesses are
created so as to give maximum contrast to the brightness regions
with the most pixels. Areas under the curve between adjacent
vertical lines are equal. B. Brightness histogram of the
transformed data.

equalization gives better overall contrast than the linear enhancement
technique, but also is a nonlinear process that tends to create contrast
where none should exist -- particularly among pixels away from the
interesting region which just happen to be at brightnesses also found in
large proportion in the region of interest. We have used linear enhance-
ment steadily in our work. The STOIC routine HISTO displays a histogram of
incoming data; CREST displays a histogram of the PC-buffer and is a faster
routine. FNN silently performs a histogram and calculates the values of

NMIN and NRANGE (which is NMAX-NMIN). NORM3 performs the actual linear transformation of the PC-buffer pixels, based on NMIN and NRANGE, and stores the result in the N-buffer.

8.1.2. Halftoning. Several methods of halftoning are explained by Stoffel and Moreland[72]. Our choice is based on speed of operation, compatibility with our display, and sensitivity to image detail. That choice is a global fixed-threshold halftone algorithm[73]. In it, a halftone screen is established over the entire display area -- the larger the halftone screen, the more distinguishable brightness levels will be available, but the larger the screen size in pixels must be. The picture, contrast-enhanced and normalized to the allowable brightness range (found in the N-buffer following execution of NORM3) is overlaid onto the halftone screen; those points which are brighter than the corresponding halftone screen level are made into white dots on the display; those which are darker become black dots. The 8x8-point screen represents a standard pattern size. In a large area of brightness 0, all of the dots will be black; in a large area of brightness 128, half of the dots will be black and half will be white, giving the impression of gray from a moderate distance.

In order for a halftone to effectively show small detail, the minimum size of detail must be comparable to or larger than the halftone screen size. Detail which is much smaller than the halftone screen falls randomly over different areas of the screen and may be dark in some locations, light in others. Thus, there is an inverse relationship between highest detectable spatial frequency (improved with small halftone screens) and distinct number of brightness levels discernible (improved with large screens). To avoid this problem, when our images are halftoned for display, they are first enlarged by reproducing each point HMAG times horizontally and VMAG

Brightness Array								Halftone Screen							
193	193	193	193	193	193	193	193	128	0	32	160	8	136	40	168
193	193	193	193	193	193	193	193	192	64	224	96	200	72	232	104
193	193	193	193	193	193	193	193	48	176	16	144	56	184	24	152
193	193	193	193	193	193	193	193	240	112	208	80	248	120	216	88
63	63	63	63	63	63	63	63	12	140	44	172	4	132	36	164
63	63	63	63	63	63	63	63	204	76	236	108	196	68	228	100
63	63	63	63	63	63	63	63	60	188	28	156	52	180	20	148
63	63	63	63	63	63	63	63	252	124	220	92	244	116	212	84

Figure 8.4. Halftone screen and the halftone image resulting from the brightness array shown.

times vertically (HMAG and VMAG are automatically set to fill a large portion of the television display with the whole image). This not only helps fill the screen but also ensures that the detail will be reproduced faithfully.

Rather than do the point-by-point comparisons, which would take intolerably long for a typical image (which may be 200,000 dots on the television) we have stored, in core, templates for the 64 distinct patterns available with the 8x8 screen. The halftoning task, SCTONE, takes each N-buffer point in turn, finds the appropriate halftone template, accounts for the enlargement of the point as described in the last paragraph and the

position of the point on the television display, and "prints" a section of that halftone template on the display. The time to display an image ranges from 1 to 2.5 seconds for the 8085 microprocessor, compared to 35 seconds for a point-by-point comparison.

This technique provides good contrast and detail in the midtones. Edges in the image are readily noticeable. The highlights and shadows tend to lack detail as compared with other techniques, and a Moire effect is occasionally noticeable when large areas of common brightness appear in an image.

A more complex scheme which we have tried is the error diffusion algorithm. In this technique the first pixel in the N-buffer is compared to a level of 128, and if the pixel is brighter a white dot is written, otherwise a black dot. Then the error between the dot that was written (255 or 0) and the actual brightness of the pixel is added or subtracted from the surrounding pixels in the N-buffer; this will bias their comparison when the algorithm reaches their location. In a typical scheme, 7/16 of the error is passed to the pixel next over to the right; 5/16 to the pixel immediately below; and 2/16 to the pixels below left and below right (it is assumed that pixels above and to the left have already been processed). This means that a dot on the display is less likely to be white if its neighbors are white.

The error diffusing algorithm provides a full tone scale, no Moire effect, and good detail detection. However, edges tend to be blurred, the texture may be noisy, and the process is much slower since no templates can be stored.

8.1.3. _Display of other image information_.   Two   routines   provide

graphical display of brightness information in the image, allowing the

experimenter to scan through the image and receive interactive feedback on

the data values and image quality.

ROWPLOT displays a graph of the (normalized) brightness across one

scan line.   The scan line is chosen as an argument to the command, e.g.

10 ROWPLOT to see the 10th line (a coordinate map of an image may be super-

imposed on the image on the digital display by the command AXES).   It has

proven especially useful for observing contrast variation and edge sharp-

ness along a line.   For example, a scan line which crosses a blood vessel

will have a ROWPLOT displaying an abrupt drop in brightness (see Figure

8.5).

O

scanned
row

YPTS

O   XPTS

disc

B

vessel

O   XPTS

Figure 8.5. ROWPLOT; the scan crosses a blood vessel and the bright optic
            disc.

CROSSHAIR displays a crosshair over the image, which can be moved

around the image by single-letter commands from the keyboard. The row and column position and the N-buffer and PC-buffer brightness values are displayed for the pixel under the crosshair. CROSSHAIR is useful, for example, when it is desired to choose a different NMIN and NRANGE than the FMM program selects (in order to enhance a particular region of the image); moving the crosshair about the region of interest will reveal the range of brightness in this region and allow intelligent selection of NMIN and NRANGE.

8.1.4. Image enlargement. One problem with the "enlargement" of images for display by simply reproducing their values in a block HMAG wide and VMAG high is that the image becomes a series of blocks, rather than a real enlargement. Especially viewed closely or in great enlargement, the image appears to have false edges and a coarse, digitized effect. Photographic enlargers perform a different transformation when they enlarge an image; a perfect enlarger of a continuous image would compress the frequency spectrum by the amount of the enlargement, i.e. if an image in one dimension $f(x)$ were doubled in size, creating $g(x) = f(x/2)$, the resulting Fourier transform would be

$$(74) \quad G(\omega) = \int_{-\infty}^{\infty} f(\frac{x}{2}) \, e^{-j\omega x} dx = \int_{-\infty}^{\infty} f(u) e^{-j\omega 2u} d2u$$

$$= 2 \int_{-\infty}^{\infty} f(u) e^{-j(2\omega)u} du$$

$$= 2 \, F(2\omega)$$

Enlargement in the digital world will create new pixels whose brightness values must be interpolated from the nearest known brightness values in the original picture. As Figure 8.6 shows, this block (sample-and-hold) enlargement scheme provides poor interpolation between the four nearest

Figure 8.6. Three different methods of interpolation for image enlargement. (a)sample-and-hold (pixel reproduction),(b)bilinear, (c)raised cosine. On the right side are shown the brightness levels of an inerpolation between a bright point (at the back corner of the plane) and a dark point (at the front).

known points of the original N-buffer which we are trying to enlarge. We could enlarge by linear interpolation of the new points among the four nearest known neighbors (part B of the figure) which gives a smoother image; however, the region around a single bright point in this scheme will have square-shaped lines of equal brightness and will not look smooth. From a signal processing standpoint, the "triangle" shape of this interpolation has frequencies higher than the sampling limit implied by the known data points ($f_{max} = 1/(2 \text{ points})$); the contours we see are the result of aliasing.

The raised-cosine interpolation (the interpolation between one known point and the next rolls off as a cosine curve between the two) has the advantages of smoothed, circular lines of equal brightness, much like a photographic enlargement, and a resultant effect on the Fourier transform closest to the photographic process (see appendix for complete derivation). Because the enlargement follows a cosine curve of frequency $1/(2 \text{ points})$, there is no aliasing. The picture tends to be mildly blurred, as a photo enlargement would be.

8.2. Algorithms for Feature Enhancement and Noise Elimination

The displayed image allows the experimenter to see an image which is unaltered except for contrast enhancement. It is occasionally desirable to enhance significant structures, such as blood vessels, notably to see if such structure exists in a low-contrast image and to improve the auto-registration process (described below). The effects of certain forms of noise may also be lessened with the use of simple transformations described here.

8.2.1. Local Masks. While certain filtering operations on an image can be done by means of Fourier transformation, spectrum modification, and inverse transformation, the sequence is time-consuming for a microcomputer (although not unbearably so, particularly with the Am9511 arithmetic processor which is installed); furthermore, it is not necessarily the best way to show local features such as a single edge. A series of local operations have been developed[74] which use the concept of local area masks (or 3-by-3 masks) to perform local image transformations (Figure 8.7).

$$P_{new}(x,y) = AP(x-1,y-1) + BP(x,y-1) + CP(x+1,y-1)$$
$$+ DP(x-1,y) + EP(x,y) + FP(x+1,y)$$
$$+ GP(x-1,y+1) + HP(x,y+1) + IP(x+1,y+1)$$

Figure 8.7. The concept of local area masks.

Such a mask contains numbers in a 3-by-3 array which is placed over the image with the center over a particular pixel on the image. The mask thus overlays the center point and its nearest neighbors in all eight

-156-

directions. To perform a mask operation, each pixel brightness value under the mask is multiplied by that number in the mask, then all nine numbers are added together and the result is the brightness value of the center pixel in the transformed image. Then the mask is moved so as to be centered on the next pixel, and the process is repeated for all pixels in the image. Put another way, the 3-by-3 mask is convolved in two dimensions with the image.

For example, the mask

1/8	1/8	1/8
1/8	0	1/8
1/8	1/8	1/8

creates a new image where each pixel in the old image is replaced with the average value of its neighbors (such a mask might be useful for display, and for smoothing over single points that are much brighter or darker than their neighbors, i.e. "salt-and-pepper" noise). It should be clear that

0	0	0
0	1	0
0	0	0

is the unity transformation.

A mask of the form

-1	1	0
-1	1	0
-1	1	0

produces a transformed image wherein the brightest points (the points with the highest value) will be those which in the original image were much the highest value) will be those which in the original image were much

brighter than their left-hand neighbor, and which are below and above points brighter than their left-hand neighbors. In other words, this transformation will isolate vertical edges in the original image with brightness increasing to the right. Other similar masks will find horizontal and 45° diagonal edges. The negative of the Laplacian

$$(75) \quad -\nabla^2 f = -\frac{\partial^2 f}{\partial x^2} - \frac{\partial^2 f}{\partial y^2}$$

can be represented by

0	-1	0
-1	4	-1
0	-1	0

(because $\partial^2 f/\partial x^2$ is $(f(i+1,j) + f(i-1,j) - 2f(i,j))$). The negative Laplacian measures the difference between a point and its neighbors' averages; this can produce subjective sharpening of a blurred picture in many cases, if it is needed for display or if edge enhancement in all directions is desired. Adding the negative Laplacian to the original picture (i.e. replacing the 4 in the middle box with 5) produces a copy of the original image with sharper edges. (We have not made use of this particular operation for our images to date).

The program MASKSET creates a 3-by-3 mask from its 9 arguments; the program MASK performs the convolution of the mask with the N-buffer (the result goes in another N-buffer-sized buffer named MBUF). The edge-enhancing and salt-and-pepper-noise-reducing masks have been used to improve pictures for display; more importantly, the edge-enhancing routines have been used to improve the registration process.

8.3. Registration and Combination

Matching up two images for eventual combination is a process that has been studied at considerable length;[75,76] the images can conceivably differ by translation, rotation, and magnification simultaneously. It is necessary to exactly line up at least one point on both images to account for each of these distortions. Some registration algorithms stress the location of the most prominent small feature on an image in order to match it up to the corresponding feature on the other image.

In general, our eye images are identical in magnification, because of the fixed position of the subject's head, and they are generally free from rotational distortion as well. We have significant x- and y- translational movements from picture to picture, however.

The use of a straight two-dimensional cross-correlation has been successful for us in registering two pictures of reasonable image quality:

$$(76) \quad R(i,j) = \frac{1}{(XPTS-i)(YPTS-j)} \sum_{x=0}^{XPTS-i} \sum_{y=0}^{YPTS-j} B(x,y) \bullet C(x+i,y+i)$$

where $\underline{B}$ and $\underline{C}$ are the brightness arrays (raw images) being brought into register. (we use the sample correlation with a divisor as shown rather than 1/(XPTS•YPTS) because of our finite sample size and because we are comparing samples rather than estimating a process; use of the fixed denominator creates too much bias toward the (0,0) point. This scheme will give erroneous results, however, if the registration distance is too far).

Correlation of ordinary images from the scanner works reasonably well; when the images are edge-enhanced, which tends to emphasize blood vessels, the correlation is improved, in effect because we have given it narrow areas of detail to bring into register.

<u>ACORXY</u> does a two-dimensional correlation; <u>ACOR</u> is a one-dimensional

-159-

version. OFFSET performs an correlation and returns the expected offset between the two pictures. SHIFT takes two arguments (the x and y offset) and creates a new image which is the translation of the original image by that offset.

With the pictures in register, any mathematical function may be executed by use of COMBINE, which performs a user-specifiable operation on each pair of points from the two buffers given as its arguments.

CHAPTER 9

SAFETY CONSIDERATIONS

In any device which passes light into the eye, consideration must be given to the danger of causing damage to the eye from the incident radiation. This issue is explored briefly here.

9.1. Light Damage to the Eye

Numerous studies[64,77] have been performed on rats, monkeys. amd rabbits to assess the amount of light irradiance that can be tolerated under varying conditions before damage occurs to the eye. A review of twenty-three of these studies is provided by Lanum[77]. In this review the illumination level for each study is converted as accurately as possible to the same standard unit, i.e. watts/cm^2 at the entrance pupil of the eye. When this conversion is made all of the studies reveal thresholds on the order of 10^{-5} to 10^{-2} watts/cm^2 over a wide range of wavelengths and exposure times. Above the threshold, damage occurred as detected by reduced electroretinogram signal, histologic evidence of destruction of photoreceptors, degeneration of the outer segments and depigmentation of the retinal pigment epithelium (RPE). These studies form the basis of the standards adopted by the American National Standards Institute (ANSI) for human light exposure, as described below.

The nature of damage to the eye caused by light has been elaborated somewhat by Ham[78], Sliney[79], and others. In the spectral range from 400-1400 nm, their studies suggest at least three types of damage due to radiation:

1. Mechanical disruption of the retinal structure due to shock waves

-161-

from extremely high power pulses which are absorbed by the RPE and choroid. This effect does not concern us at our power levels.

2. Thermal damage due to absorption of energy by these tissue layers causing a local increase in ambient temperature.

3. Actinic (photochemical) damage which is a particularly bothersome problem at the shorter wavelengths of the visible spectrum.

Thermal injury to the retina results when local temperature rises more than 10 ° above normal. At these and higher temperatures irreversible denaturation and coagulation occurs in proteins. The duration of overheating is important, as well; for the same radiant energy (radiant power times exposure duration) applied to the eye, thermal effects predominate over photochemical injury at shorter exposure times (See Fig. 9.1) due to the inability of the local circulation to carry away heat rapidly.

Thermal damage is assumed to be due to energy absorption by melanin in the RPE and in the choroid. Thus the damage threshold is lower for those wavelengths where RPE melanin absorbs more strongly -- specifically, toward the shorter wavelengths. This trend is mitigated as the wavelength drops to 400 nm and below because much of the incident light is absorbed by the crystalline lens of the eye.

Photochemical effects also have a lower absolute threshold at the shorter wavelengths, and predominate when the actual incident power is too low for thermal damage but overall incident energy is high because of long exposure time. For exposure times of ten seconds or longer, damage threshold may be considered to be a function only of total energy up to exposures of approximately twelve hours. The eye does not practically integrate damaging exposures for longer than this time.

The actual mechanism of photochemical damage is not completely known;

Figure 9.1. Plot of retinal damage thresholds for various visible
wavelengths applied to intact eyes. For short exposures, thermal
effects predominate because heat cannot be rapidly carried away.
For photochemical damage, the product of power times duration is
roughly constant for a given wavelength[62].

photo-oxidation in the retina and the RPE (which usually are well oxy-

genated) or oxidative processes involving the membranes or the visual pig-

ments have been proposed. Melanin produces free radicals when exposed to

light, and histology shows that the first change to occur in photochemical

damage is clumping of melanosomes (melanin-containing bodies) in the RPE.

Possibly, degradation of the RPE melanosomes leads to impaired metabolic

support from the underlying choroid, followed by deterioration of the pho-

toreceptors.

9.2. Safety limits

Whatever the exact mechanism of photochemical and thermal damage, their effects are clear, and concern for retinal damage from optical instruments has prompted regulation and standards for light intensity, of which the best known are the ANSI Standards for Safe Use of Lasers, ANSI-Z136.1[80,81]. In the Class I laser safety standard, the standard used for lasers which will be directly observed by the unprotected eye, the following maximum permissible exposures (MPE) in the 400-700 nm range are set by ANSI:

Emission duration(sec)	energy limit	Radiance limit on 7mm pupil
1.0×10^{-9} - 2.0×10^{-5}	2.0×10^{-7} J	5.2×10^{-7} J/cm^2
2.0×10^{-5} - 1.0×10^{1}	7.0×10^{-4} $t^{3/4}$ J	1.8×10^{-3} $t^{3/4}$ J/cm^2
1.0×10^{1} - 1.0×10^{4}	3.9×10^{-3} J	1.0×10^{-2} J/cm^2

These standards are designed to give a 10x margin of safety.

Delori[62] points out that a number of commercially available fundus cameras actually exceed these limits, especially when used with longer exposures, e.g. several minutes of fluorescein angiography. The flash used in our own original human experiments (Chap. 4) provides illumination up to 155 mJ/cm^2 in 4.3 ms into a 1.75 mm entrance pupil, which is in excess of the standard.

In our system, there are several factors which make these figures overly conservative. Because we have a raster which scans a large area of retina, the energy incident to a particular local area of retinal tissue is smaller than would be for laser light which is stationary on one spot. As long as the scanning motor continues to operate at a minimum rate, this large raster provides an extra margin of safety. Furthermore, our light

spot has a finite size on the retina, whereas the standards were established for a focused spot whose size is due to diffraction, and optical distortions, only; the focused spot applies much more energy to a local area. Despite these protections, it is useful to use the ANSI standards as a first-cut limit for light exposure.

Using the middle time range in the table above, we find that for a 50 µs exposure, the allowed incident light is 8 mW; for an entire scan of 0.1 s the MPE is 1.2 mW; and for prolonged illumination of 10 s the power at the pupil may not exceed 400 µw if it is to conform to the ANSI standards. (There are also standards for maximum source radiance in units of $J/cm^2/sr$; for our system these standards are not as stringent as the radiant-energy standards).

If we are to have constant illumination of the patient's eye for viewfinding, focusing, etc., then the power to the eye during these non-data-taking times should be considerably reduced from that level used for data collection; the highest light intensities should be reserved for those times when they are absolutely needed, i.e. when fluorescence is being recorded.

9.3. Providing Instrument Safety

In order to provide patient safety with our particular scanner, all possible conditions must be taken into account which would allow dangerous light levels to reach the eye. Included among these are excessively high light source output, excessively long exposure, excessively close patient proximity to the light source, optical system abnormalities which increase light levels at the eye, and failure or slowdown of the scanning motors.

The arc lamp light source in current use in the retina scanner provides light intensity at the eye's pupil that is an order of magnitude lower than the requirements specify. With the higher light sources proposed, the instrument would exceed the limits unless filters and other measures are used to cut down the spot light level. In order to maintain patient safety, special hardware has been designed to sense abnormal conditions and shut down key elements, such as the arc lamp, in case of risk to the patient. The hardware also is designed to act to protect the photomultiplier from light damage.

9.3.1. Design philosophy. The safety board hardware is based on a series of sensors which can monitor the light levels and motor velocities mentioned above, check them against pre-specified limits, and, if an abnormal condition is found, act to cut off further exposure to the light and/or protect the photomultiplier.

In designing safety monitors, reliability is at a premium and extra considerations are necessary to assure safe system operation:

1. Independence of computer operation. Once the limits on each measurement are set, the safety monitor should run independently of the 8085 microprocessor and be resistant to many computer failures and slowdowns.

2. Failsafe sensors and actuators. These elements should be designed so that single component failure can never cause safety to be violated. In the event of component failure, either steps must be taken to ensure that the monitoring function is still intact, or an alarm must be generated and the light exposure stopped immediately.

3. Fault annunciators. The user must know which safety test/condition

-166-

has caused an alarm. Some display is necessary to reveal the unusual condition.

Circuitry to implement safety monitoring and shutoff has been designed for the scanner. Photodiodes are used to monitor light output, which may be tapped off the optical system with beamsplitters. The scanner motor amplifiers provide a feedback signal which is differentiated and rectified to monitor velocity. An analog switch, sample-and-hold and A/D converter allow any sensor voltage to be read directly by computer programs.

For each sensor there is an upper and lower limit register. Each register is set by the computer to the desired bound by loading and latching a value into its D/A converter; the analog output is then fixed until later changed from the computer. The upper and lower bounds are constantly compared to the incoming sensor value, and an alarm is generated if either is exceeded. The presence of upper and lower bounds aids in the failsafe effort; e.g. if a photodiode fails, the light output would fall below the lower bound (Fig. 9.2).

Two levels of alarm actions are provided: one closes the optical shutter and cuts off voltage to the gate of the photomultiplier; the other performs these functions and also shuts off power to the lamp. Each comparator may be tied to either alarm line. As with other parts of the circuit, these actuators are so constructed that opening or shorting a single component will cause failure in a safe mode.

When a fault condition occurs, the operation of the actuators will generally cause the fault condition to vanish; in fact, closing the shutter or extinguishing the arc lamp will cause lower-bound faults on the photodiode channels. In order to maintain an accurate record of what caused the fault, an $\overline{S}$-$\overline{R}$ latch on each channel latches up when that channel's compara-

SAFETY

Upper limit

Monitoring

Sensors

Lower limit

LATCHING
D/A

LATCHING
D/A

Shutter

Lamp
cutoff

PMT
cutoff

Figure 9.2. Schematic of the safety monitor operation. When the sensor voltage exceeds the upper limit or drops below the lower limit, current flows out of the comparator to trigger the alarm cutoffs, turning on the LED in the process. Not shown are the latches which record any fault conditions which have occurred since they were last reset.

tor switches for a fault condition, causing an LED to indicate a fault on that channel.

The safety monitor provides an extra measure of security against thermal and photochemical damage to the patient's eye, allowing the use of light sources with power exceeding that of the arc lamp to achieve improved

signal-to-noise while still maintaining safety.

CHAPTER 10

SYSTEM TESTS

In this chapter I will discuss some of the characteristics and capabilities of the retina scanner as designed and built. This discussion begins with some quantitative assessments of system performance measures, and goes on to demonstrate the images acquired under various conditions, in reflectance and fluorescence. Throughout the chapter, the methods and procedures used to obtain different results will be mentioned, along with various problems which were encountered and managed along the way, so that the procedures can be repeated and refined by other workers. The use of the model eye in testing the system and in experimentation to evaluate the mathematical model of Chapters 1-3 is also discussed.

The next and final chapter briefly summarizes the progress made in developing a means of measuring retinal oxygenation and evaluating the model, and outlines the next steps to be taken in developing a noninvasive means of measuring retinal oxygenation and establishing the role of hypoxia in metabolic retinopathies.

10.1. Performance Measurements

10.1.1. Light levels. As mentioned before (Chapter 5), the present system uses a mercury arc lamp filtered to transmit 436 nm for excitation. The original optical system designed for the scanner faced problems due to use of uncoated singlet lenses, which (a) transmit less light because of increased reflections, and (b) provide a poorer quality image which may (and did) require extra stops to maintain a good image over a smaller area.

Additionally, the optical design was optimized in the second system (see Fig. 5.7) to maximize light transmission to the eye along with optimization of magnifications and reduction of required lens speed. Light levels (in microwatts) recorded at different sites in the optical systems are displayed in the table, using one of three excitation filters: a yellow broadband long-pass (designed to pass from 500 nm on up), an Oriel 436 nm interference filter, and a Melles-Griot 436 nm interference filter:

Site	No Filter	Yellow LP	Oriel 436	M-G 436
After filter	735	489	51.7	33.4
After L3 aperture stop	438	277	28.3	19.0
At pupil	340	230	21.0	11.4
Old system after filter	660	460	50.0	
Old system at pupil	4.8	3.3	0.26	

We find that (a) the use of better lenses and optimization of dimensions in the second optical system yields 80-fold increased illumination at the subject's pupil due to the lessened requirement for stops, decreased reflections, and lessened aberrations; (b) the light level in the blue 436 nm region is still considerably less than the official safe illumination levels, even for exposures of considerable length of time. If we had a higher-brightness source we could increase illumination and hence signal-to-noise ratio.

When the yellow long-pass filter is placed in series with one of the 436 nm filters, one can measure:

Oriel + yellow long-pass $0.011 \ \mu W$

M-G + yellow long-pass $< 0.0001 \ \mu W$

Thus the Melles-Griot filter has greater protection against breakthrough illumination which would produce false (reflectance) results on

fluorescence images.

10.1.2. PMT Cathode Voltage; dark current. The sensitivity of the photomultiplier (in amps per watt, or size of anode charge pulse per incident photon) is dependent upon the cathode-to-anode voltage, which is settable from the front panel of its power supply. Dark current is also increased by increasing the voltage, in a roughly exponential fashion (Ideally, the discriminator threshold should be adjusted for each different PMT voltage level, but this is impractial in the current arrangement; the threshold set, 100 millivolts, represents the best compromise). In practice, PMT voltage is optimized by observing the image on the viewfinder (real-time) monitor and selecting voltage for appropriate sensitivity. For most fluorescence work, the PMT is run at 1900-2000V; for reflectance work in the yellow, it is operated at 1000-1100V. Dark counts are typically 0 to 10 counts for a 30-microsecond pixel.

10.1.3. Spot Size and Resolution. There are two parameters affecting the spatial resolution attainable with the system. The first is the actual size of the light spot on the retina (i.e., when the scanners are not moving); the second is the actual movement of the scanner, and consequently of the excitation spot, during the time of collection of data for a single pixel.

By slowing down the scanner motor rate and shortening the single-pixel collection time (PTTIME), each pixel represents as small a physical scan distance as possible. Thus the spatial resolution is maximized under these conditions. The retinal spot size itself is finite (in the current configuration it is nearly equal to the pinhole size; it can be reduced arbitrarily by making the pinhole smaller, at the cost of excitation intensity)

and limits resolution under these conditions; however, with a 50 micron pinhole pictures such as those shown in Figure 10.1 can be achieved (remember for this and all pictures that a bilinear normalization was done during the processing -- the small salt-and-pepper variations seen are smaller numerically than they are on the picture, because image brightness values have been expanded to fit the full 256-gray-level range of the display).


Figure 10.1. High-resolution scanner image of a human retina (all photos are _in vivo_).

In this picture the scanner moves 50 μm across the retina during one pixel time; the retinal spot size (not including scanning) is also 50 μm. The disk is 30 pixels wide; the smallest blood vessels visible are 25-50 μm in diameter. (Although two structures separated by this distance could not be resolved under these conditions, detection of one object of this size is easily accomplished.) This distance is considerably smaller than the size of a photocoagulation spot.

As mentioned earlier, the critical size is the characteristic width of hypoxia, i.e. the smallest size of tissue which can become infarcted with a single significant lesion. This width is unknown at present although it is bounded by the size of small capillaries (8 μm) and the size of obvious cotton-wool infarcts (100-300 μm). Further means for determining this critical size are discussed in the next chapter.

Using scan rates still slower than the above, both direct measurements of point spread function (using a razor edge as object) and observation of a ruled grating reveal that lines 20 μm apart can be separated. The scan rate used for this observation (16 msec/scan line), if extended over a 30-to-50-line scan, would cause the raster to appear to move in front of a subject's eye, causing movement of the eye itself during the scan. If shorter PTTIMEs (pixel times) can be made available, higher resolution can be achieved at the same scan rates.

10.1.4. Image size. The size of the image is limited by the maximum range of the scanners. Figure 10.2 shows an image including the optic disk at one edge and the macula at the other, a distance of 4-5 mm (22 degrees) on the retinal surface. The scan size can be up to 50% larger than this if desired.

Figure 10.2. Retinal (reflectance) scanner image encompassing optic disk and macula.

10.2. Model Eye Studies

For performing experiments on fixed objects, such as chemical preparations of FAD, a model eye was constructed (Fig. 10.3) consisting of a plastic chamber into which a methylmethacrylate contact lens (serving as a "cornea") and a plastic intraocular lens (normally used as a surgical implant) were assembled. A glass spacer marked the distance to the retina; any slide, capillary tube, cuvette or other object could be placed in the

Figure 10.3. Diagram of the model eye.

model eye at this distance. A viewing lens on the back allowed proper alignment of such objects. When filled with water, the optical properties and dimensions of the model closely approximate the human eye.

The retina scanner with the model eye in place has been used to test some elements of the mathematical model of reflectance and fluorescence from the eye presented in Chapters 2-3. According to the model, the fluorescence from a given concentration of FAD in the eye is given by

(1) $\quad F = \eta I_0 \frac{\varepsilon}{\varepsilon + a} (1 - e^{-B_1 L})$ (in the forward direction,)

(2) $\quad = \eta I_0 \frac{\varepsilon}{\varepsilon + a} \cdot \frac{1}{2} (1 - \cos\phi) \frac{B_1}{B_1 + B_2} (1 - e^{-(B_1 + B_2)L})$ (in and back)

where, among other things, I_0 is the incident intensity and B_n (the subscripts stand for the excitation and emission wavelength) are proportional

-176-

Figure 10.4. The model eye placed in the headholder at the subject position of the retina scanner.

to absorption coefficients and to concentration of the species.

If we measure the fluorescence of tubes of higher and higher concentration of FAD we approach saturation (the mean free path of an incident photon becomes much smaller than the thickness of the cell), and the rise in counts with concentration slows down. Thus the curve of observed fluorescence vs. concentration at a given wavelength for a given species should have the form

(3) $F = \alpha (1 - e^{-KC}) + D$

where all variables except C are independent of FAD concentration. α is the

asymptotic value. The extra term $\underline{D}$ accounts for dark counts.

An experiment testing this formula with increasing concentrations of FAD in capillary tubes has been performed (Fig. 10.5) using the longpass filter to block excitation-wavelength light.

Fluorescence vs. Concentration
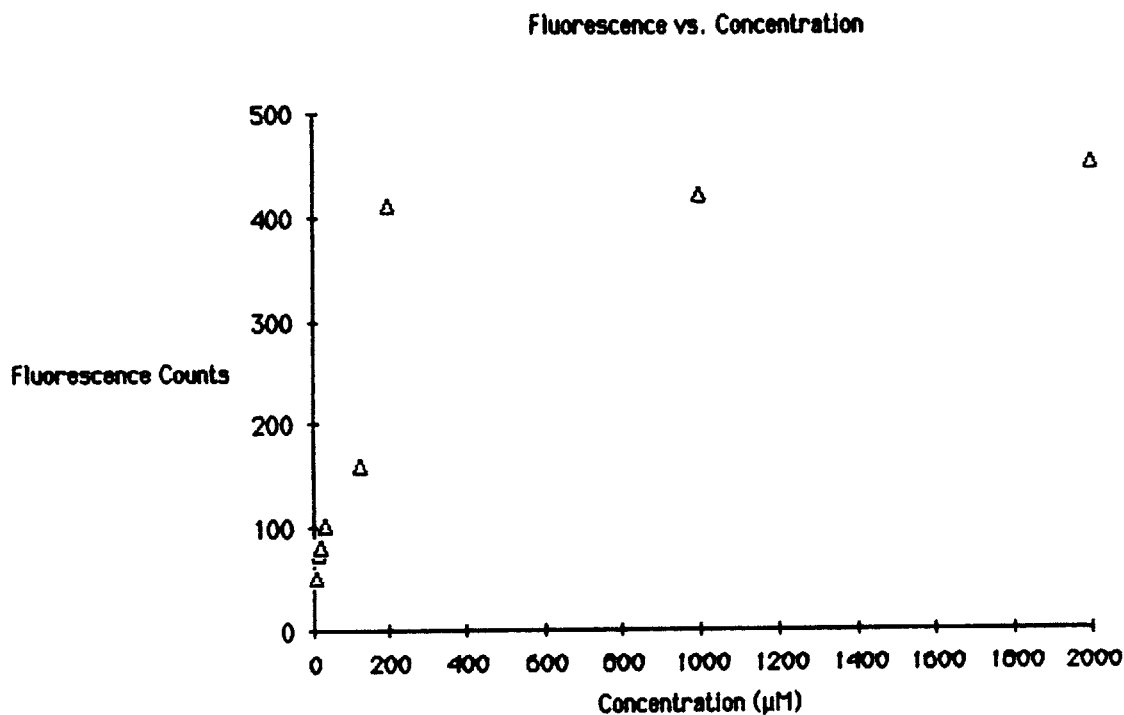
Figure 10.5. Curve of fluorescence counts (436 nm excitation, >500 nm emission) vs. concentration of FAD in capillary tubes.

The data definitely display asymptotic behavior, with the asymptote near 450 counts. The points appear to be an excellent fit to an exponential of the form shown above except for the point at 200 µM, which pulls the curve up somewhat (without this point the best fit is $F = 415 (1 - e^{-2.94\ C(mM)}) + 36$ with correlation $r = 0.998$). Suggestions about possible deviation from exponential behavior due to scatter are discussed in the next chapter.

10.3. Imaging the Eye

10.3.1. Setting up the Subject. Figure 10.6 shows the experimental scanner layout as seen from the operator's position.

Figure 10.6. The retina scanner layout as seen from the operator's position. The X-Y viewing scope, image-display TV and floppy disk drives are closest to the operator, along with the power-supply controls for the PMT and lamp (the operator's computer console is just in front of and to the left of the photo).

The following figure shows the subject's end of the table. The subject is seated at one end of the table with the head placed against the headholder as shown. A fixation lamp is available (the small rod suspended from the top of the headholder) but the raster pattern itself often makes a more useful target. The scan parameters -- scan rate, size, and line position;

Figure 10.7. The subject in position to be scanned in the experimental scanner layout. The headholder may be moved in three transla-tional axes and one rotational axis. Also visible in this pic-ture are the fixation lamp (suspended from the top of the head-holder), several of the system's lenses, the x-scan mirror (rear center) and the filter wheel. The arc lamp is at left fore-ground; the hooded path to the photomultiplier is at rear.

X and Y dimensions, pixel time -- are set up based on the image to be recorded (subject to the restriction that the amount of time spent record-ing data on a line must be less than half the time the X-mirror takes to traverse the line forward and backward, otherwise a duplicate "mirror-image" forms in the right half of the image memory; in other words the res-triction is XPTS x PTTIME < (1/2) SCANTIME). The proper raster is now being presented to the subject.

As mentioned in Chapter 5, there is a retinal and a pupil focus

adjustment to be made. The retinal focus corrects for ametropia and is a matter of setting the distance between L_1 and the retinal conjugate preceding it in the excitation path. This is done by moving L_1 with an adjusting knob until the sharpest raster pattern is seen. To adjust the pupil focus, the position of the patient's eye is moved by moving the headholder. When the patient approaches the proper position from a distance (with the photomultiplier voltage applied and the gate and shutter open), the operator sees in the real-time viewfinding display first an image of the subject's face, then the outer eye with lid and brow (for the illumination raster scans this area) with a small bright corneal reflection (reflex). As the eye gets closer to the proper pupil focus, the corneal reflex enlarges, taking on a dark center as the system's optical stops block the reflection. At the correct point, the bright reflection disappears; in the dark center, now filling the image space, the retinal structures are seen.

Horizontal translation and angular movement (rotation of the eye in a left/right direction) are used in balanced fashion to position the desired part of the retina in the image area; the fixation target is used as a guide. The same is true in the up/down direction. For fluorescence experiments, the filters can be changed to the fluorescence configuration (in which reflections are much less a problem) and the eye position further adjusted for maximum signal on the XY display.

10.3.2. Recording an image. With setup complete, the operator issues a single command which usually does the following:

1. Black out the computer console and image display, so stray light from them doesn't contribute to the data;

2. Check that the time formula above is not exceeded and that there are not too many datapoints for the memory size;

-181-

3. Record and save one or more image frames, perhaps with filter wheel motions in between to use different wavelengths;

4. Create a histogram of the brightness values;

5. Use the histogram to normalize the pixel array for display; and

6. Display the image.

After the image is examined, the operator may use COMMENT and ANNOTATE to record text and parameters related to the image, may combine the image with others, or may use it as a guide for further images to be recorded.

10.3.3. Image characteristics and artifacts. Examples of reflectance images taken with the retina scanner are shown in Figures 10.1 and 10.2. We will examine some of the characteristics and artifacts of images produced, to demonstrate problems that can occur and how they are managed.

In Figure 10.8 we see the image of the retina of a rhesus monkey in reflectance. The dark spot to the right of the optic disk (right center of image) is a Purkinje reflection from the monkey's cornea. Moving the monkey's head under guidance from the viewfinder TV can cause the spot to disappear; moving the head the wrong way makes it larger. The spot is dark, not light, because the spot is so bright that it causes a steady stream of photons to the amplifier/discriminator such that the input to the discriminator section is always above the threshold — the individual photons are not separated as they are for fainter areas. The discriminator never sees a crossing of the threshold, and so never produces pulses (just a steady high output) — thus a count of zero is recorded. The same is true for the central area of the optic disk, where the light is bright enough to saturate the discriminator section in the same fashion. This condition arises from having too much voltage on the photomultiplier

Figure 10.8. Image of a rhesus monkey eye.

cathode. Turning down the voltage produces a normal image.

(Photon counting is modeled as a Poisson process with the unit time interval being the minimum distinguishable interval between pulses -- about 8 ns in our case. If an image area is so bright that the number of counts is on the near one per interval, then some counts will be missed because more than one photon is coming in in some intervals. There are formulas for compensating for this error; however, adjusting the voltage proves to be an effective way of regulating the image data. Adjusting the

discriminator threshold is another).

The source of reflections was also explored by the examination in the retina scanner of black rabbits in which the cornea and lens were surgically removed. A glass cover slip was glued to the free scleral border and the eye stabilized by fastening the slide to a metal suspending ring (Fleiringa ring) so as to put only minimal traction on the eye. Reflectance images of the rabbit retina were thus free of Purkinje reflections; although the images are sparse in detail (the rabbit retina has a structure somewhat different from ours) the bright spots did not appear in the image. The optical stops made no difference in the image, adding evidence that the Purkinje images were indeed the source of the spots in the human and monkey images. The particular placement of the stops necessary to remove the reflections in those images was final confirmation of this.

On some early images, a pattern of grey lines, nearly the length of a scan line and appearing regularly every 3-4 lines depending on scan rate, was superimposed on the image. Measurement showed these lines to be repeating at 60 Hz; the source was the arc-lamp housing and fan, which were vibrating with the line rate and causing the arc itself to wander, imperceptibly to the naked eye but very clearly in the recorded images. Improvements to the arc-lamp mount solved this problem. Such pure-frequency noise could also be eliminated through digital methods by means of a variation on a 60-Hz notch filter which accounts for the gaps in time between successive rows of data (by phase-matching the artifact to a template which is itself a 60-Hz wave with gaps of the same length).

The beginning-of-line-detect circuitry (Chapter 6) is of clear importance in creating a steady sync pulse so that image rows line up properly. That circuit is a comparator with hysteresis -- necessary to avoid glitches

and multiple row-counts at the end of a line -- using the scan velocity as input. Because of this, very low-gain, very slow-rate scans which have small velocities may give signals too low to reliably trigger the line detect. Figure 10.9 shows the artifact produced when a line detect is triggered by a scan-motor signal glitch instead of a valid beginning-of-line turnaround.


Figure 10.9. Image showing one line in which the beginning-of-line detect triggered on an erroneous signal.

Addition of a lowpass filter to the input of the hysteresis comparator, and

amplification of the x-velocity signal, have greatly improved this matter.

10.4. Fluorescence images of the eye

The mercury arc lamp used with the Melles Griot 436 nm line narrow interference filter produces 11.4 microwatts illumination at the eye, quite a bit less than the levels permitted in the derivation of Chapter 9 (1.2 mw for 0.1 second exposure; 400 μw for 10 second exposure). Using the shutter and filter wheel effectively (so that light input is reduced when not recording an image) we could increase our incident light an order of magnitude or more for fluorescence imaging. With the arc lamp and standard image parameters we record counts of 40-200 photons per pixel. The next page shows reflectance and fluorescence versions of a human eye. The fluorescence image is unprocessed (see below).

The low light levels of fluorescence in the eye produce more problems than simply the low signal-to-noise ratio. Changing from viewfinding to fluorescence imaging involves changing either the excitation or the emission filter so a reflectance path becomes fluorescence. If the excitation is changed (from yellow to blue), the focus point is changed due to chromatic aberration of the lenses; thus a larger area of lens and retina are illuminated. Use of achromatic lenses corrects this problem partially. Additionally, a change in color of the raster is distracting and causes the eye to move. For these reasons rotating the emission filter is more practical. In addition to changing the filter, the PMT voltage must be increased between reflectance and fluorescence measurements (see Chapter 11).

In order to increase the signal-to-noise, we need to get higher photon

Figure 10.10. (a)Reflectance image of a human eye. (b)Fluorescence image of the same eye, summed from 49 individual frames but otherwise un-processed (compare to Fig. 10.13, page 192).

-187-

counts in each pixel. This can be done by increasing the pixel time, slowing down the scanner motors enough to give comparable resolution. When this is done, our pixel counts increase proportionately to the time increase. However, a raster slowed to much more than 100 msec period can be perceived by most human eyes as a moving line, which the eye naturally follows instead of staying in place as it does for a faster raster. Because of this eye movement, we cannot readily use this technique to increase the number of photon counts.

We can instead take several rapid image frames and add them together. Although this takes as much time as the other, the eye is not distracted by a moving line -- most subjects can hold the eye fixed for a few seconds on what they perceive to be a stationary raster. Images have been recorded which are the sum of 50 to 100 rasters, which have proportionately increased counts. The extra signal is traded for defocusing due to those slight movements of the eye which do occur and cause misregistration of the images (in the current system there is not enough memory to hold the images separately, nor is the disk fast enough to collect the images in real time; thus we cannot use the registration algorithms in this case and the images are simply added on top of each other). The image shown in the figure above is a sum of 49 images. Although this method did succeed in bringing out shadowy structure outlines, the blurring from eye position changes takes a toll.

10.5. Image enhancement

The effect of applying some of the image transformations of Chapter 8 on the retina scanner images is shown in the followng photographs. Figure

10.11(a) is our original image, showing the disc and the macula. The following photographs show the same photo processed in various ways. In Figure 10.11(b) a smoothing filter has been applied to the image. The result is less precise definition of areas in exchange for a more continuous and more curved-looking appearance, free of "salt-and-pepper" noise. The appearance is smooth and somewhat blurry at close range; yet this technique has a role in improving noisy images.

Conversely, we can edge-enhance to provide images with sharp borders and contours. This is useful for registration of multiple images; in most eyes there are distinctive contours of vessels or other markings to serve as fiducial marks. Figure 10.12 is the result of a Laplacian operator applied to the original image. Even with this simple 3x3 convolution operator, the vessels are greatly enhanced, along with a few false contours. We can set the Laplacian representation of two images in register easily, then use the same shift to register the originals.

A smoothing operation is used to improve the quality of the noisy fluorescence image of Figure 10.13(b) on the following page. The top image is the reflectance photo, the bottom is the fluorescence photo passed through a 5x5 smoothing filter. The structural contours are much more evident in this picture compared to the unprocessed image.

The chief caveat in using these methods to improve picture quality is that we are doing just that — improving the picture as we see it. The true data is still that of the original unprocessed image; information may be lost in image transformations such as the above. Also, to be able to fairly compare images, the same general rules of transformation should be applied to all. Radiologists and pathologists who are looking for qualititative information may process images to their liking; when we look to see

Figure 10.11. (a) Reflectance image of a human eye. (b) The same image, pro-
cessed with a 3x3 smoothing filter.

Figure 10.12. The same image as Fig. 10.11(a) with a Laplacian edge-enhancing operator applied. Vessel-edge and disk contours stand out.

Figure 10.13. (a)Reflectance image of a human eye (same as Fig.10.10(a)).
(b)Fluorescence image of the same eye (same as Fig.10.10(b)) with
a smoothing operator applied. Although the detail is still quite
shadowy, structures are much more readily apparent than on the
unprocessed image.

if a fluorescence image is present, as in the above examples, the processing can be helpful to us. When quantitative processing is to be done, one must be extremely wary of prior modifications to the data.

CHAPTER 11

FUTURE WORK

11.1. Where we are

This thesis outlines steps which have been taken toward the eventual goals of understanding the role of hypoxia in metabolic eye diseases such as DR, and developing a means to measure and follow such hypoxia.

A model has been proposed and elaborated based on the general properties of reflectance, fluorescence and absorption in a multicomponent medium in an eye which has focusing power and a limited aperture. A number of biochemical, animal, and model-eye experiments have been performed which address parts of this model, and which to date are encouraging as far as the use of the model to measure hypoxia. However, the model remains unproven experimentally in the general case. The experiments in monkeys also point encouragingly to the possibility of observing hypoxia through fluorescence measurments if a device of sufficient sensitivity can be made.

We have made strides in the operating philosophy, design, and development of an instrument which has high enough sensitivity and image quality to permit us to view and map the extremely dim native FAD fluorescence of the retina and observe the fluorescence changes as FAD converts to FADH with increasing hypoxia. In its present state, the instrument is a computer-controlled high-resolution video ophthalmoscope with image processing and multiple wavelengths. With multiple summing of images, fluorescence can be seen in the retina with structure similar to that expected by the model and by animal experiments. The low efficiency of fluorescence in the eye, the low light level (compared to the safety

limits) used for illumination, and the suboptimal properties of 436 nm excitation in exciting FAD fluorescence compared to 458 nm take away from the image signal; structures can be seen on the image but they are shadowy and require processing because of the low signal-to-noise level.

To complete the task -- to have a fully functioning system capable of doing the planned future experiments -- requires that some improvements be made to the existing instrument.

11.2. Improvements

11.2.1. Scatter. As mentioned in Chapter 2, light in the eye undergoes absorption, reflection, refraction, and scattering. Recent experiments have suggested that the effect of scatter on image quality is significant. In the current scanner all imaging is done in the excitation path; the photomultiplier's field of view includes the entire pupil, except that area which is blocked by Mx for excitation; thus any photon leaving the pupil (at other than the excitation window) will count as part of the brightness of the particular pixel being scanned. Any excitation photon which scatters in the eye will not strike the desired small spot on the retina but will still count, if it fluoresces where it DOES strike, as a photon emanating from the nominal spot. In fact, a fraction of the excitation beam is scattered and illuminates a large area of the retina, causing a diffuse whole-retina signal to be mixed in with the signal from specific points on the surface. This leads to defocusing of the image. Experiments with the model eye show that scattered fluorescence increases from 10% to 50% that of the nominal-spot fluorescence as a test solution of FAD becomes

more concentrated from the 20 μM to the 2 mM range.

Although this phenomenon is not yet fully understood and calculated for the optics of the eye, it appears that it may explain some of the blurring seen in fluorescence images. It is also possible that scatter helps to explain why the curve of fluorescence vs. FAD concentration in Chapter 10 rose faster than expected with concentration, although this remains to be proven experimentally.

Since scatter blurring arises because the PMT receives input from the entire retina, the answer would be to provide a detector which itself has spatial selectivity much like the excitation path. Using a scanning detector such as an ISIT would not solve the problem, as explained in Chapter 5; scatter in the emission path after fluorescence would produce photon images on the ISIT cathode in areas other than the desired spot; these would then add together with the brightness recorded when the scanner was supposed to be on that spot, to give the same defocused result. (The ISIT's recording time is much longer than a single pixel-time, so we do not have time to isolate the desired retinal spot area).

If we use the same PMT we have now, with the same time-slicing scheme, and provide an optical scanning system much like the excitation path and synchronized with it, then only the light which struck the desired retinal spot at the time that the excitation was supposed to be on that spot will be recorded for that spot. Plans for such a second scanning path are in progress.

11.2.2. Other optical improvements. Among the other developments which could improve the overall functionality of the scanner optics are:
   1. All focusing should be done by moving the system optics, not the patient. Doing the latter causes problems because movements are

not steady, especially not in animals. Addition of a second focusing lens assembly in the common path could make this possible.

2. Certainly, more light in excitation is needed. The switch to an argon laser with appropriate safety precautions, or to another adequate light source offering more intensity and better wavelength, will be the biggest factor providing improved signal-to-noise ratio in the image.

3. Lower-noise photomultipliers exist than ours, as do PMT's with higher sensitivity in the 520 nm range. Although dark counts are not a problem in reflectance, they are not that far away from the signal level in fluorescence mode; if the fluorescence levels stay low, reduction of the noise by acquiring a lower-noise and/or broader-spectrum PMT (made by RCA and others), cooling it with dry ice or electrical means, or using magnetic lenses (which reduce random excitations by allowing only electrons from the central region of the PMT surface to reach the dynode chain) will all increase S/NR.

4. When proper light levels are achieved, the introduction of a scanning monochromator in the optical system will allow us to measure the fluorescence spectra very carefully. This will permit us to see if the curve matches the FAD curve (suitably adjusted for ocular media) and will give us several datapoints instead of just two for separating the FADH from the FAD and the other components.

5. The variation of light level which sometimes appears across an image -- dark at one edge, light at the other -- relates to difficulty in tilting the subject's eye in just the right position with

the current headholder, and possibly to shadowing by the excitation-path stop (Mx) in the emission path. Controlled experiments should be able to elucidate the mechanisms which produce this variation, and to outline steps for removing it as a problem.

11.2.3. Electrical and software improvements. Further gains in image quality can be attained by addressing some of the following ideas concerning the electronic hardware and computer software:

1. Although the halftone scheme has served us well, a continuous-tone display for images would still be superior in image quality and flicker tolerability. Many such displays are now available with complete memory and interfaces.

2. I mentioned in Chapter 10 that there is not enough memory on the 8085 to allow multiple rapid-sequence images to all be stored, from which they could be better registered and processed before being combined. Even keeping the 8085 bus with its 16-bit address space, it would still be possible to add unlimited fast virtual memory by establishing a bus window which could open a small area of address space to any segment of the memory.

3. Along the same line, faster disks and a DMA facility for disk transfers would reduce the overhead in time spent waiting for these transfers when doing multiple rapid-sequence images. The DMA has been designed and a prototype wired up.

4. All in all, it might be better if the processor itself were faster, since the minimum PTTIME is currently limited by a polling loop to read the photon counter, which requires 32 µs to run. In the next generation of the scanner, it may be wise to think in terms of 68000-based systems which have more memory, faster

processing speed, continuous-tone displays as standard peripherals, and other conveniences.

5. The problems with discriminator saturation (chapter 10) and with the need to switch PMT cathode-anode voltage rapidly when going from fluorescence to reflectance and back, would all be obviated if the PMT power supply were operated under computer-based remote control. Many available PMT supplies have remote voltage control options. Voltages may be programmed at fixed levels for different procedures, or an autoscale could be built which keeps the sensitivity at just under the saturation range. The value of voltage used must be passed to the computer so that the image brightness can be calibrated, if necessary, for the different sensitivity present at different voltages. The above could apply to the discriminator threshold voltage as well as the PMT supply voltage.

6. The PMT amplifier has two outputs at present -- one to the real-time display, one back into the discriminator. Superfast buffers (most likely from discrete transistors) should be used to further isolate these signals, to make sure that adjustment of one does not affect the quality of the other.

7. An ECL prescaler stage in front of the Schottky TTL flipflops in the low-order photon-counter chain would raise the effective count rate, making sure no photons were ever lost by missing a count strobe. The chief disadvantages of this are power requirements and heat.

8. Use of graphic devices such as a joystick, or their simulation in software, would be an easier method of moving the scan raster area on the retina, compared to the keyboard commands now used.

11.3. Future experiments

As a guide to further experimentation, I will suggest the following experiments and studies as important steps to the full understanding of hypoxia in DR and its measurement with noninvasive technologies such as the scanner.

11.3.1. FAD/FADH mixtures; mitochondrial preparations.

Some work has already been done in observing FAD and FADH preparations in the scanner with the model eye. As a means of confirming the scanner's performance, a monochromator can be used to measure spectra of FAD and FADH at different concentrations and of mixtures of the two components. Care must be taken in making and preserving FADH as it rapidly converts back to FAD in the presence of oxygen (for method of conversion, see [82]). Plain wide or narrowband filters may be used instead of a monochromator to see the effect of component variation on measurements with such filters, which provide more light level than a monochromator.

By the same token, measurements can be made of blood or hemoglobin and other absorbers and fluorochromes in similar fashion. Curves such as the fluorescence-concentration relationship in Chapter 10 can be plotted for all such combinations, fully testing the optical/chemical theory presented in this thesis.

11.3.2. Multiple-wavelength images and analysis.

The determination of redox state from fluorescence and reflectance measurements depends on the ability to make measurements at multiple wavelengths and to combine them (Figure 2.9 is the simplest case). With the improvements described above

the system would be fully ready to make such measurements in rapid succession and to register and combine them properly. Experimentation with these experiments will yield the best wavelengths to use to optimize contrast-to-noise for the redox calculation.

In previous discussions the usual assumption was that multi-wavelength experiments would be done by using single-wavelength excitation and varying the emission band through several closely-spaced narrowband areas. I will now suggest that the implementation is far more likely to succeed if the wavelength variation is done on the excitation side and a broadband emission filter is used. The reason, of course, is that the excitation beam can be made as narrowband as possible and still have reasonable incident light levels to the eye; it is just a matter of getting the appropriate power out of the light source in the first place, so that the filter's attenuation brings it down to a bright but safe level. If narrowband emission is used, we are still limited on the excitation side by safety requirements, and the resulting dim light output must be reduced even more by the emission filters. It is a straightforward extension of the theory presented here to account for wideband emission. Tests on mitochondria with close observation of the excitation spectra of oxidized and reduced environments can provide the best choices of wavelength to use — presumably the excitations would be in the 450-490 nm band, using fixed or tunable lasers.

11.3.3. Separation of choroidal and retinal circulations. One of the main biological obstacles to the use of the hypoxia measurement technique is the dual blood supply to the retina. The choroidal circulation, as mentioned in Chapter 1, supplies the posterior segments of the retina, while the retinal circulation supplies the anterior. Most mitochondria are

located in the photoreceptors themselves, which may receive blood supply from both circulations, although it is likely they receive majority supply from the choroid. Yet it appears to be the retinal circulation which is impaired most in diabetes. If hypoxia plays a role in DR, is it only tissue fed by the retinal circulation that is affected, or are both circulations changed in their ability to oxygenate tissue?

It is surgically possible to obliterate the retinal circulation, at least locally, by ligating major arteries or veins, leaving only the choroid to supply tissue. Then the change in local oxygenation and in fluorescence measurement could be measured in an animal using the scanner. In humans, branch vein occlusion and central retinal arterial occlusion are pathological entities; patients with these abnormalities could be studied serially with the scanner, ideally beginning when the lesion is new.

The role of light bleaching can be investigated in similar fashion, by shining bright light into the eye just before an image is taken, and watching for a difference in fluorescence properties because of the metabolic changes brought on by bleaching. In this as in all experiments care must be taken to avoid bleaching by the fluorescence illumination source itself whenever possible (i.e., keep the subject in the dark or under dim light until just before the image is to be recorded).

11.3.4. Retinal isolates. The preparation of an isolated retina in a nutrient bath has been accomplished and maintained by a number of investigators [83]. Establishing such a preparation in the model eye would allow study of the fluorescence properties of the retina without worrying about abnormalities in human lenses or the limitations of the closed eyeball and narrow pupil. Traded off against these advantages must be the fact that such a preparation will never completely mimic normal in vivo conditions.

11.3.5. Bandwidth of hypoxia . As noted, the minimum spatial resolution we require is that which allows us to see the smallest practical area of hypoxia created in the pathogenesis of DR (or any other diseases studied). We have upper and lower bounds for this area (Chapter 10) but no full studies to determine the true minimum.

Once we are able to reliably identify areas of hypoxia, a study should be undertaken in which histograms of hypoxic-region size are made to determine how many pixels wide such a region might be (obviously, the criterion for determining a hypoxic area would have to be insensitive to the size of such an area for this technique to be fair). If the histogram shows a bell curve of probability vs. region size, and the small-size end of the curve (mean minus an appropriate number of S.D.'s) is much larger than one pixel wide, then we can be confident that we have sufficient resolution to identify most practical hypoxic areas. If the curve is truncated or stacked, or the minimum approaches one pixel in width, then the resolution needs to be increased, for small hypoxic areas are smaller than one pixel size and are probably being missed.

11.3.6. Diabetic patient prospective study. The final test of the ability of retinal FAD fluorescence measurement to predict future eruptions of DR in a local area of the retina must come from a several-year, prospective patient study, following diabetic patients in semiannual or other intervals with retinal fluorescence scans and with other techniques currently used to assess DR -- fluorescein angiography and fundus photography. Statistics would be done to correlate patients' eyes and local retinal areas deemed hypoxic in retina scanner studies with areas found (concurrently or later on) to show local retinopathy. Such a study has already been designed and planned at Boston University Medical Center.

11.3.7. Pre/post photocoagulation study. As an investigation of the role of laser photocoagulation in treating diabetes and of the role of hypoxia in this treatment, scans should be made before, immediately after, and at regular intervals after photocoagulation treatment in a group of patients. Careful analysis will be required to understand the change in fluorescence that is attributable to scar and to burned tissue, compared to the change in the local area surrounding a photocoagulation spot, compared to the change at a more distant point. Photocoagulation of local and generalized areas can also be performed in animals, to be followed with FAD scans to measure the changes.

Eventually, in the best of worlds, the experimentation could reach the point where the predictive value and sensitivity of FAD retinal scans for future local retinopathic outbreak would be high enough that local photocoagulation in humans might be possible. When that takes place, evaluation will include serial retinal scans and clinical evaluation of the progression of disease including fundus photos and fluorescein angiograms. Comparison would be made with panretinal photocoagulation in patients deemed to have similar numbers and sizes of lesions. A crossover from the local photocoagulation group to the panretinal group would be possible.

When planning for and thinking about such impressive eventualities as this last, it is good to remember that the major goal to be realized from the scanner development is actually in the understanding of the mechanism of diabetic retinopathy and other metabolic retinal diseases — in evaluating the existence and role of hypoxia in their progression. Even if the power of the technique presented in this thesis for identifying future areas of retinopathy — and for facilitating safer, more localized treatment — is not fully realized, a more complete understanding of the early

development of the major retinal diseases, which may be its major contribution, is likely to be part of the chain of new knowledge which eventually will result in improved therapies for metabolic blindness.

APPENDIX A

DIFFUSE REFLECTANCE WITH A NONCONCENTRIC EXIT PUPIL

In Chapter 2 we present the problem of calculating the expected dif-
fuse reflectance (Lambertian source) when the exit pupil and entrance pupil
meet only at their apex, at the reflecting surface, and go off in cones of
different widths at different angles. I present here the derivation for
the reflectance in such a situation. We will see that the adjustment for
these factors is a multiplier dependent only on the angles themselves, and
thus that processing of the results is no more complicated than with
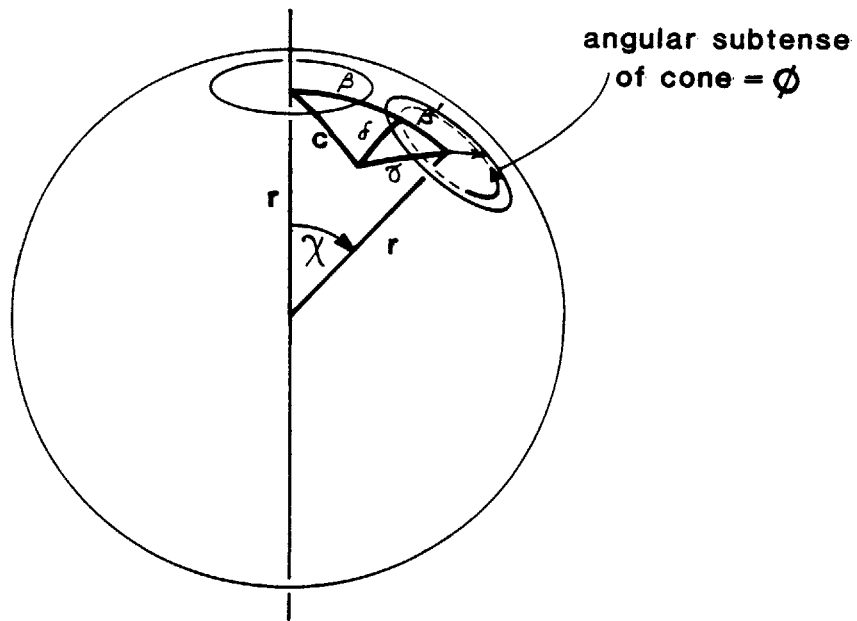transmission measurements.

Figure A.1. Geometry of diffuse reflectance

The input is from a single ray I as shown. The diffuse reflectance is
captured only in a cone of half-angle φ, whose center forms an angle χ with
the incident ray. β and β' are distances on a great circle joining the

incident ray and the emission cone center at radius r. δ is a distance perpendicular to β.

Our method is to create an integral which treats each area dA of surface inside the exit zone (the intersection between a sphere of radius r and the exit cone) equally, then introduce as the integrand the expression for magnitude of diffuse reflectance over each dA.

To integrate over a single circle (annulus of width $d(r\gamma)$ in the exit zone, and radius given by $r \sin(\gamma)$) we just use the contour form

$$(1) \quad \int_0^{(2\pi r \sin \gamma)} ds \; d(r\gamma) = 2\pi r \sin \gamma \; d(r\gamma)$$

(We could break up the contour more explicitly in terms of δ and β' but it turns out to not be necessary). then the whole area of the exit zone, as a fraction of the entire sphere, is

$$(2) \quad A = \frac{1}{4\pi r^2} \int_0^{\phi r} \int_0^{(2\pi r \sin \gamma)} ds \; d(r\gamma)$$

This does indeed integrate to 1 for a whole sphere ($\varphi = \pi$).

Now, the intensity is a constant $4R_0$ times the cosine of the angle between I and the current point being studied. As the diagram shows, this angle is $\underline{c}$ and is described by the relation

$$(3) \quad \cos c = (\cos \delta)(\cos \beta)$$

$$= \frac{\cos \gamma}{\cos \beta'} \cos (\chi - \beta')$$

Then the reflectance in the cone is given by

$$(4) \quad \frac{R}{4R_0} = \frac{1}{4\pi r^2} \int_0^{\phi r} \int_0^{(2\pi r \sin \gamma)} \frac{\cos \gamma}{\cos \beta'} \cos (\chi - \beta') \; ds \; d(r\gamma)$$

The inner integral may be broken down:

$$(5) \quad \int_0^{(2\pi r \sin \gamma)} \frac{\cos \gamma}{\cos \beta'} \cos (\chi - \beta') \, ds$$

$$= (\cos \gamma) \left[\int_0^{(2\pi r \sin \gamma)} \frac{(\cos \chi)(\cos \beta')}{\cos \beta'} \, ds \right.$$

$$+ \left. \int_0^{(2\pi r \sin \gamma)} \frac{(\sin \chi)(\sin \beta')}{\cos \beta'} \, ds \right]$$

Now χ is a constant and γ is constant over any one circle so

$$(6) \quad \frac{R}{4I_0} = 2\pi r \sin \gamma \, (\cos \chi)(\cos \gamma)$$

$$+ \, (\sin \chi)(\cos \gamma) \int_0^{(2\pi r \sin \gamma)} \tan \beta' \, ds$$

The integral in this term is zero because the contour symmetrically ranges over positive and negative values of β' and the tangent is an odd function. Now the full expression is

$$(7) \quad \frac{R}{4R_0} = \frac{1}{4\pi r^2} \int_0^{\phi r} 2\pi r \, (\sin \gamma)(\cos \chi)(\cos \gamma) \, d(r\gamma)$$

$$= \frac{2\pi r \cos \chi}{4\pi r^2} \left[\frac{r}{2} \sin^2 \phi \right]$$

$$(8) \quad R = R_0 \sin^2 \phi \cos \chi$$

for $\phi < \frac{\pi}{2}$.

This is the expression for diffuse reflectance in a cone of half-angle ϕ at an angle χ from a particular incident ray. If we want to range over all of the incident rays (which are in a cone of half-angle θ, but rather than being Lambertian are of even distribution), the geometry is reversed (the excitation "ray" from before is now the center of the emission cone, and we have proved that the emitted light depends only on the position of that line and the cone width) then we will be setting up exactly the same integral as (4) (with a different normalizing constant $2\pi r^2(1-\cos \theta)$ instead of $4\pi r^2$ because the full intensity R_0 is all contained in the exci-

tation cone) and find

(9) $$\frac{R}{R_0} = \frac{\sin^2\phi}{2\pi r^2(1-\cos\theta)} \int_0^{\theta r} (2\pi r \sin\gamma) \int_0^{\sin\gamma} \frac{\cos\gamma\,\cos(\gamma - \beta')}{\cos\beta'}\,ds\,d(r\gamma)$$

and when we simplify this we get our final answer:

(10) $$R = R_0 \frac{\sin^2\phi\,\sin^2\theta\,\cos\gamma}{2\,(1 - \cos\theta)}$$

This is the total effect of limited, nonconcentric excitation and emission cones and diffuse reflectance on the reflected light. R_0 in this case is the total reflected light as stated in Equation 11 of Chapter 2.

There are three important features of this result:

1. The correction is simply a constant once the optical parameters of the system are set up (including ϕ, θ, and χ); thus these optical alterations will not complicate the processing.

2. The decrease in reflectance as the excitation and emission pupils are separated goes as a cosine, and for the angles we have been considering (typically 11 degrees' separation between pupils — in fact with an eye with a 17mm nodal distance and 9 mm maximum pupil, 30 degrees is the absolute limit and pupil size must be very small) this is only a 2% decrease in signal. The advantages (including suppression of specular reflections) are tempting in this case.

3. The reflectance _is_ strongly dependent on the widths of the cones themselves — a change in one cone from 10 to 12 degrees can make a 45% increase in light level. (The best light throughput comes when, other constraints being unimportant, the excitation and emission pupils are the same size.

References

1. Leibowitz H., Krueger D., Maunder L., and others, "The Framingham Eye Study Monograph. An ophthalmological and epidemiological study of cataract, glaucoma, diabetic retinopathy, macular degeneration and visual acuity in a general population of 2631 adults, 1973-1975.," Surv. Ophthalmol..

2. Klein R., Klein B., Moss S., and others, "Prevalence of Diabetes Mellitus in Southern Wisconsin," Am. J. Epidemiol. Vol. **119**, pp.54-61 (1984).

3. Klein R., Klein B., and Moss S., "Visual Impairment in Diabetes," Ophthalmology Vol. **91**(1), pp.1-9 (January 1984).

4. Klein R., Klein B., Moss S., Davis M., and DeMets D., "The Wisconsin Epidemiologic Study of Diabetic Retinopathy. II. Prevalence and Risk of Diabetic Retinopathy When Age at Diagnosis is Less Than 30 Years," Arch Ophthalmol Vol. **102**, pp.520-527 (April 1984).

5. Klein R., Klein B., Moss S., Davis M., and DeMets D., "The Wisconsin Epidemiologic Study of Diabetic Retinopathy. III. Prevalence and Risk of Diabetic Retinopathy When Age at Diagnosis is 30 or More Years," Arch Ophthalmol Vol. **102**, pp.527-532 (April 1984).

6. Bresnick G., Engerman R., Davis M., De Venecia G., and Myers F., "Patterns of Ischemia in Diabetic Retinopathy," Proc. 80th Annual Meeting Am. Acad. Oph. Otol., pp.OP-694 to OP-709 (1975).

7. Young R., "Visual Cells," Scientific American Vol. **223**(4), pp.80-91 (October 1970).

8. Kohner E., "Dynamic Changes in the Microcirculation of Diabetics as Related to Diabetic Microangiopathy," pp. 41-47 in Diabetic Micro-Angiopathy: The August Krogh Memorial Symposium, ed. Ditzel J. and Poulsen J., A. Lindgren + Soner AB, Molndal (1975).

9. Ditzel J. and Standl E., "The Problem of Tissue Oxygenation in Diabetes Mellitus," Acta Medica Scandinavica Vol. **587**(Suppl.), p.49 (1975).

10. Read J., Ernest J., Goldstick T., Raichand M., and Petersen A., "Hyperglycemia and the Retinal Circulation in Man," Inv. Oph. Vis. Sci. Vol. **(Suppl.)**, p.168 (1980).

11. Kohner E.M., "The problems of retinal blood flow in diabetes," Diabetes Vol. **25**(2(suppl)), pp.839-44 (1976).

12. Yanoff M., "Ocular Pathology of Diabetes Mellitus," Am. J. Ophthal. Vol. 67, p.21 (1969).

13. Zuckerman and Weiter J., "Oxygen Transport in the Bullfrog Retina," Exp. Eye. Res. Vol. 30(2), pp.117-127 (1980).

14. Diabetic Retinopathy Study Group, "Photocoagulation treatment of proliferative diabetic retinopathy: The second report of diabetic retinopathy study findings.," Trans. Am. Acad. Oph. Oto. Vol. 85, p.82 (1978).

15. Raskin P., Pietri A.O., Unger R., and Shannon W.A., "The effect of diabetic control on the width of skeletal-muscle capillary basement membrane in patients with Type I diabetes mellitus," New England J Med Vol. 309(25), pp.1546-50 (1983 Dec 22).

16. Siperstein M.D., "Diabetic microangiopathy and the control of blood glucose [editorial]," New England J Med Vol. 309(25), pp.1577-79 (Dec 22 1983).

17. Silver I.A., "Changes in Tissue Micro-environment during Hypoxia," in Oxygen Measurements in Biology and Medicine, ed. Payne, Butterworths (1975).

18. Jobsis F., Mitnick M., and Snow T., "Pyrene Butyric Acid: A non-invasive probe for in situ intracellular oxygen concentration," Adv Exp Med Biol Vol. 75, pp.47-54 (1976).

19. Lubbers D., "Spectrophotometric Examination of Tissue Oxygenation," Advances in Experimental Medicine and Biology Vol. 37A, pp.45-54 (1974).

20. Benson R., Meyer R., Zaruba M., and McKhann G., "Cellular Autofluorescence -- Is It Due to Flavins?," J. Histochem. Cytochem. Vol. 27(1), pp.44-48 (1979).

21. Wood W.B., Wilson J.H., Benbow R.M., and Hood L.E., in Biochemistry: A Problems Approach, W.A. Benjamin, Inc., Menlo Park, CA (1974).

22. Chance B., Oshino N., Sugano T., and Mayevsky A. , "Basic Principles of Tissue Oxygen Determination from Mitochondrial Signals," in Physiology of Oxygen Transport to Tissue.

23. Chance B. and Schoener B. , "Fluorometric Studies of Flavin Component of the Respiratory Chain," in Flavins and Flavoproteins, ed. Slater E.C., Elsevier Press, New York (1966).

24. Barlow C. and Chance B. , "Ischemic Areas in Perfused Rat Hearts: Measurement by NADH Fluorescence Photography," Science, pp.909-10 (Sep 1976).

25. Barlow C. H., Harden W. R., Harken A. H., Simson M. B., Haselgrove J. C., Chance B., O'Connor M., and Austin G., "Fluorescence Mapping of Mitochondrial Redox Changes in Heart and Brain," Critical Care Medicine Vol. 7(9), pp.402-406 (1979).

26. Wexler P., "The Television Ophthalmoscope Pyridine Nucleotide Fluorometer: Measurement methodology for the detection of pyridine nucleotide fluorescence in the rabbit retina," in M.S. Thesis, University of Illinois (1981).

27. Wolbarsht M., "Ophthalmic Uses of Lasers," , pp.401-420.

28. Bakker N., "Fundus Reflectometry -- An Experimental Study," Doc. Ophthalmol. Vol. 38(2), pp.271-360.

29. Janssen F., "A Study of the Absorption and Scattering Factors of Light in Whole Blood," Med. and Biol. Eng. Vol. 10, pp.231-240 (1972).

30. Levi L., pp. 136ff in Applied Optics, John Wiley + Sons, New York (1980).

31. Quistorff B. and Chance B., "3-D Mapping of Metabolic State of the Rat Liver," Fed. Proc. Vol. 36(1328), p.528 (1977).

32. Wodick R. and Lubbers D., "Quantitative Evaluation of Reflexion Spectra of Living Tissues," Hoppe-Seyler's Zeitschrift Physiologische Chemie Vol. 355, pp.583-594 (May 1974).

33. Anderson N. and Sekelj P., "Light Absorbing and Scattering Properties of Non-Haemolysed Blood," Phys. Med. Biol. Vol. 12, pp.173-184 (1967).

34. Boettner E. and Wolter J. , "Transmission of the Ocular Media," Invest. Ophthal. Vol. 1(6), pp.776-83 (Dec 1962).

35. Mellerio J., "Light Absorption and Scatter in the Human Lens," Vision Res. Vol. 11, pp.129-141 (1971).

36. Laing R.A. and Masters B.R., "A Non-Invasive Time-Sharing Spectrofluorometer Study of Corneal Metabolism: The Effect of Barbiturates on Rabbit Corneal Endothelium and Epithelium," (1981).

37. Bennett A.G. and Francis J.L., "The Eye as an Optical System," pp. 101-131 in The Eye, Vol. 4, ed. Davson H., Academic Press, New York (1962).

38. Last R.J., in Wolff's Anatomy of the Eye and Orbit, W.B. Saunders Co., Philadelphia (1968).

39. Christensen S., _Acta Pharm. Tox._ Vol. **32**(Suppl. 2), p.27ff (1973).

40. Waddell W.J. and Bates R.G., "Intracellular pH," _Physiol. Rev._ Vol. **49**, pp.285-329 (1969).

41. Bessey O., Lowry O., and Love R., "The Fluorometric Measurement of the Nucleotides of Riboflavin and Their Concentration in Tissues," _J. Biol. Chem._ Vol. **180**, pp.755-769 (1949).

42. Ghisla S., Massey V., Lhoste J., and Mayhew S., "Fluorescence Characteristics of Reduced Flavins and Flavoproteins," pp. 15-24 in _Reactivity of Flavins: Proceedings of the Symposium,_ ed. Yagi, Univ. of Tokyo Press, Tokyo (1973).

43. Zak, Steczko, and Ostrowski, "Reduced Forms of the Flavoprotein from Egg Yolk," _Bull. Soc. Chim. Biol._ Vol. **51**, p.1065 (1969).

44. Long C.(ed.), in _Biochemists' Handbook,_ D. Van Nostrand Co., Princeton NJ (1968).

45. Geeraets W.J., Williams R.C., Chan G., Ham W.T., Guerry D., and Schmidt, F.H., "The Loss of Light Energy in Retina and Choroid," _Arch. Oph._ Vol. **64**, pp.606-615 (1960).

46. Chance B. and Lieberman M. , "Epithelial and Endothelial Sides," _Exp Eye Res_ Vol. **26**, pp.111-17 (1978).

47. Kuck J. and Yu N., "Raman and Fluorescent Emission of the Human Lens. A New Fluorophor," _Exp. Eye Res._ Vol. **27**, pp.737-741.

48. Young L. and Sheena D., "Survey of Eye Movement Recording Methods," _Behavior Res. Meth. and Instr._ Vol. **7**(5), pp.397-429 (1975).

49. Coakley D. and Thomas J.G., "The Ocular Microtremor Record and the Prognosis of the Unconscious Patient," _Lancet,_ pp.12-15 (March 5, 1977).

50. Lubbers D., Kessler M., Scholz R., and Bucher T., "Cytochrome Reflection Spectra and Fluorescence of the Isolated, Perfused, Hemoglobin-free Rat Liver during a Cycle of Anoxia," _Biochemische Zeitschrift_ Vol. **341**, pp.346-350 (1965).

51. Harosi F., "Absorption Spectra and Linear Dichroism of Some Amphibian Photorecptors," _Journal of General Physiology_ Vol. **66**, pp.357-382 (1975).

52. Ripps H. and Weale R.A., "Flash Bleaching of Rhodopsin in the Human Retina," _J. Physiol._ Vol. **200**, pp.151-159 (1969).

53. Fram I., Read J., McCormick B., and Fishman G. , "In Vivo Study of the Photolabile Visual Pigment Utilizing the Television Ophthalmoscope Image Processor," Proc Conf in Computers in Ophthal (5/Apr/1979).

54. Franceschini N., Kirschfeld K. , and Minke B., "Fluorescence of Photoreceptor Cells Observed in vivo," Science Vol. 213, pp.1264-1267 (Sept. 1981).

55. Lubbers D. and Wodick R., "The examination of multicomponent systems in biological materials by means of a rapid scanning photometer.," Applied Optics Vol. 8(5), pp.1055-1062 (May 1969).

56. Sigelman J., Trokel S., and Spector A., "Quantitative Biomicroscopy of Lens Light Back Scatter: Changes in Aging and Opacification," Arch. Ophthal. Vol. 92, pp.437-442 (1974).

57. Chance B., Schoener B., Oshino R., Itshak F., and Nakase Y., "Oxidation-Reduction Ratio Studies of Mitochondria in Freeze-Trapped Samples (subtitled) NADH and Flavoprotein Fluorescence Signals," J Biol Chem Vol. 254(11), pp.4764-71 (10/June/1979).

58. Sanadi D.R and Fluharty A., "On the Mechanism of Oxidative Phosphorylation: VII.," Biochemistry Vol. 2(3), p.523ff. (May-June 1963).

59. Chance B. , "Spectrophotometric and Kinetic Studies of Flavoproteins in Tissues, Cell Suspensions, Mitochondria and their Fragments," Flavins and Flavoproteins, Elsevier Press (1966).

60. Shapiro J.M., "Laser contour angiography: a new technique for quantifying the topography of the optic disc," Ph.D. dissertation, University of Michigan, Xerox-University Microfilms (1976).

61. Shapiro J.M., Kini M.M., and Philpott D.E., "Early diagnosis of glaucoma: A microprocessor-based system.," Proc. 31st Ann. Conf. Eng. Med. Biol. Vol. 20, p.161 (1978).

62. Delori F. C., Parker J. S., and Mainster M. A., "Light Levels in Fundus Photography and Fluorescein Angiography," Vision Research Vol. 20(12), pp.1099-1104 (1980).

63. Wyckoff, C., Personal Communication, 1980.

64. Lawwill T., Crockett S., and Currier G., "Retinal Damage Secondary to Chronic Light Exposure," Documenta Ophthal Vol. 44(2), pp.379-402, (1977).

65. Yariv A., in <u>Introduction</u> <u>to</u> <u>Optical</u> <u>Electronics</u>, Holt, Rinehart, and Winston, New York (1976).

66. Webb R. H. and Hughes G. W., "Flying Spot Television Ophthal-moscope ," <u>IEEE</u> <u>Frontiers</u> <u>of</u> <u>Engineering</u> <u>in</u> <u>Health</u> <u>Care</u> (<u>proceedings</u>), pp.209-212 (1980).

67. Westheimer G. , "The Maxwellian View," <u>Vision</u> <u>Res</u> Vol. **6**, pp.669-82 (1966).

68. <u>Electronics</u> (October 30, 1975).

69. Fredrikson G.M., "Design of an Interactive High-Resolution Graphics Display," in <u>S</u>.<u>M</u>. <u>Thesis</u>, MIT Department of Electrical Engineering and Computer Science (1979).

70. Green W.B., <u>Digital</u> <u>Image</u> <u>Processing</u>: <u>A</u> <u>Systems</u> <u>Approach</u>, Van Nostrand and Co., New York (1983).

71. Gonzalez R.C. and Wintz P., <u>Digital</u> <u>Image</u> <u>Processing</u>, Addison-Wesley, Inc., Reading, MA (1977).

72. Stoffel J.C. and Moreland J.F., "A Survey of Electronic Techniques for Pictorial Reproduction," <u>IEEE</u> <u>Transactions</u> <u>on</u> <u>Communications</u> Vol. **COM-29**(12), pp.1898-1925 (Dec. 1981).

73. Bayer BE., "An optimum method for two-level rendition of continuous tone pictures," <u>IEEE-ICC</u> **1**, pp.26-11 to 26-16.

74. Rosenfeld A. and Kak A., <u>Digital</u> <u>Picture</u> <u>Processing</u>, Academic Press, New York (1976).

75. Barnea D.I. and Silverman H.F., "A Class of Algorithms for Fast Digital Image Registration," <u>IEEE</u> <u>Transactions</u> <u>on</u> <u>Computers</u> Vol. **C-21**(2), pp.179-186 (Feb. 1972).

76. Aggarwal J.K., Davis L.S., and Martin W.N., "Correspondence Processes in Dynamic Scene Analysis," <u>Proceedings</u> <u>of</u> <u>the</u> <u>IEEE</u> Vol. **69**(5), pp.562-571 (May 1981).

77. Lanum J. , "The Damaging Effects of Light on the Retina. Empirical Findings, Theoretical and Practical Implications.," <u>Surv</u> <u>Ophthal</u> Vol. **22**(4), pp.221-49 (Jan/Feb 1978).

78. Ham W., Ruffolo J., Mueller H., and Guerry D., "The Nature of Retinal Radiation Damage: Dependence on Wavelength, Power Level and Exposure Time," <u>Vision</u> <u>Res</u>, pp.1105-11 (1981).

79. Sliney, D., <u>Interaction</u> <u>between</u> <u>Lasers</u> <u>and</u> <u>the</u> <u>Eye</u>, MIT Seminar, May 2, 1984.

80. American National Standards Institute and Food and Drug Administration, "Laser Products: Amendments to Performance Standard," <u>Federal</u> <u>Register</u> (Nov. 30, 1983).

81. American Conference of Governmental Industrial Hygienists, in _A Guide for Control of Laser Hazards_, Paul M. Larrod Co., Baltimore (1976).

82. Dawson et. al., _Data for Biochemical Research II_, 2nd edition 1980.

83. Rothschild J., _Biophysical J._, 1980.

BIOGRAPHICAL NOTE

Jonathan Teich was born October 23, 1955, in New Brunswick, New Jersey, and grew up in the neighboring town of Edison. After graduating from Edison High School in 1972, he attended college at Caltech, receiving the B.S. degree in Biochemistry and Electrical Engineering in 1976. He then pursued a joint M.D.-Ph.D. degree at Harvard Medical School and MIT, receiving his medical degree in 1983. Since then, he has shared time between a residency in internal medicine at the Brigham and Women's Hospital and completion of his graduate studies.

He has received graduate support under a Fannie and John Hertz Foundation Fellowship, and from the Medical Scientist Training Program of the National Institutes of Health. He is a member of Tau Beta Pi, Sigma Xi, and the IEEE.