

Aus der Universitäts-Augenklinik Tübingen  
Abteilung Augenheilkunde II  
Ärztlicher Direktor: Professor Dr. E. Zrenner

**The Flicker Electroretinogram  
in Phase Space:  
Embeddings and Techniques**

Inaugural-Dissertation  
zur Erlangung des Doktorgrades  
der Medizin

der Medizinischen Fakultät  
der Eberhard-Karls-Universität  
zu Tübingen

vorgelegt von

ALBRECHT JOHANNES RILK

aus  
Aalen

2003

Dekan: Professor Dr. C. D. Claussen

1. Berichterstatter: Professor Dr. E. Zrenner

2. Berichterstatter: Professor Dr. F. Nüsslin

# Contents

<b>Table of Contents</b>	<b>6</b>
<b>Abbreviations</b>	<b>8</b>
<b>1 Introduction</b>	<b>8</b>
1.1 Purpose and scope . . . . .	9
1.2 ERG basics . . . . .	10
1.2.1 Normal ERG . . . . .	10
1.2.2 Clinical applications . . . . .	12
1.2.3 ERG varieties . . . . .	14
1.3 Why yet another ERG system? . . . . .	14
1.4 Calypso . . . . .	16
1.4.1 Functional overview . . . . .	16
<b>2 Materials and methods</b>	<b>20</b>
2.1 The delay embedding . . . . .	20
2.1.1 Constructing delay vectors . . . . .	21
2.1.2 The look of a trajectory . . . . .	23
2.1.3 The theoretical point of view . . . . .	24
2.2 Nonlinear projective noise reduction . . . . .	26
2.2.1 NNR basics . . . . .	26
2.2.2 Estimating the parameter values . . . . .	28
2.3 Reducing the recording time . . . . .	29

---

2.4	Experimental setup . . . . .	30
2.4.1	Hardware and software . . . . .	30
2.4.2	Standard ERG recordings . . . . .	31
2.4.3	Safety considerations . . . . .	31
2.5	Participants . . . . .	32
2.5.1	Ethical considerations . . . . .	32
2.6	Examinations . . . . .	34
2.7	Data processing . . . . .	34
2.7.1	Frequency filtering . . . . .	34
2.7.2	Removing artifacts . . . . .	37
2.7.3	Averaging . . . . .	38
<b>3</b>	<b>Results</b>	<b>40</b>
3.1	Standard ISCEV flicker ERG embedding . . . . .	40
3.2	Flicker trajectory topology . . . . .	42
3.2.1	Descriptive approach . . . . .	42
3.2.2	Polar coordinates help quantifying . . . . .	44
3.3	Pathological flicker ERGs . . . . .	46
3.3.1	Usher’s syndrome . . . . .	46
3.3.2	Retinitis pigmentosa . . . . .	47
3.3.3	Juvenile macular degeneration . . . . .	49
3.3.4	Stargardt’s disease . . . . .	50
3.4	Nonlinear Noise Reduction performance . . . . .	51
3.4.1	Generating the test data sets . . . . .	52
3.4.2	Different kinds of noise . . . . .	53
3.4.3	Iterations . . . . .	54
3.4.4	NNR on sinusoids . . . . .	55
<b>4</b>	<b>Discussion</b>	<b>58</b>
4.1	Topographic angle synopsis . . . . .	58

---

4.2	NNR and other noise reduction techniques . . . . .	60
4.2.1	NNR and averaging . . . . .	60
4.2.2	NNR and frequency filters . . . . .	62
4.3	Synergies . . . . .	63
4.4	What can go wrong . . . . .	63
4.4.1	Embedding parameters . . . . .	64
4.4.2	Artifacts . . . . .	68
<b>5</b>	<b>Conclusion</b>	<b>70</b>
5.1	Future outlook . . . . .	71
5.1.1	Follow-ups . . . . .	71
5.1.2	Extending the polar coordinate concept . . . . .	72
5.1.3	Visualizing multi-dimensional spaces . . . . .	72
5.1.4	The multifocal ERG . . . . .	74
<b>A</b>	<b>Noise in different flavors</b>	<b>76</b>
A.1	Signal-independent noise models . . . . .	77
A.1.1	Gaussian noise . . . . .	77
A.2	Signal-dependent noise models . . . . .	78
A.2.1	Spectrally matched noise . . . . .	79
<b>B</b>	<b>Standard deviation, RMS, and SNR</b>	<b>81</b>
<b>C</b>	<b>The Fourier Transform</b>	<b>83</b>
	<b>List of Figures</b>	<b>86</b>
	<b>List of Tables</b>	<b>87</b>
	<b>Bibliography</b>	<b>88</b>
	<b>Danksagung</b>	<b>92</b>

<b>Lebenslauf</b>	<b>93</b>
-------------------	-----------

# Abbreviations

ADC	Analog-to-Digital Converter
Calypso	Computerized Analysis of Physiological Samplings in Ophthalmology
DSP	Digital Signal Processor
ECG	Electrocardiogram
ERG	Electroretinogram
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
IIR	Infinite Impulse Response
ISCEV	International Society for Clinical Electrophysiology of Vision
mfERG	Multifocal Electroretinogram
NNR	Nonlinear Projective Noise Reduction
PCA	Principal Component Analysis
RMS	Root Mean Square
RP	Retinitis Pigmentosa = Retinopathia Pigmentosa
SMN	Spectrally Matched Noise
SNR	Signal-to-Noise Ratio
WGN	White Gaussian Noise

# Chapter 1

## Introduction

About one to three million people worldwide suffer from a form of hereditary retinal degeneration, i. e., *retinitis pigmentosa* (RP); in Germany alone, 30,000 to 40,000 patients are affected. Presently, there is still no cure for RP, which is thus one of the most frequent causes for blindness. Only a few therapies are currently available that allow for partial inhibition of retinal degeneration, but these are applicable to only some rare forms of RP. However, the fact that some treatment does exist encourages us to search for new ways of, if not curing, at least inhibiting degeneration in the more frequent forms of RP.

Since RP was first described in the 19<sup>th</sup> century, several different therapeutic approaches have been attempted. In recent years, molecular genetic studies have been able to shed some light on the pathophysiological disease mechanisms. For example, as of April 2003, there are 134 mapped genes and 90 cloned genes that are thought to be responsible for RP (DAIGER, 2003). In clinical practice, however, the only reasonable approach seems to qualify the pathological changes within the retina in terms of morphology, and, of course, to record these changes by electrophysiological techniques as precisely as possible. A treatment is most effective when applied in the disease's early stages, as is the chance of preserving visual function. By means of objective electrophysiology, we can detect pathological changes even before the patient himself becomes aware of them. But even such early detection needs to be improved upon.

Hence, the task of clinical research is to refine and to extend the electrophysiological methodology with two aims in mind: On the one hand, we need to



preserve the patient's life quality as much as possible by early diagnostics and treatment; on the other hand, meaningful and diagnostically relevant results are mandatory, especially in the disease's final stages, when signals have become very weak or even seemingly extinguished. Achieving such aims would be beneficial not only for RP patients, but for all patients with heavily reduced electroretinograms (ERG), e. g., those suffering from STARGARDT's disease, and from other retinal dystrophies, be they hereditary or acquired.

The diagnostic advantages by state of the art data processing available to ophthalmology are far from being exhausted. A look at the literature reveals several methods that have recently emerged from the nonlinear dynamics sciences, and that have been shown useful for, e. g., cardiological problems. But as of yet, these methods obviously have not been applied to ophthalmological tasks.

## 1.1 Purpose and scope

In recent times, strong evidence was found that the representation of complex, nonlinear data in a phase space embedding reveals more information than the common amplitude-time diagram. It has seemingly not yet been considered if this includes ERG data as well.

Therefore, this thesis will introduce some nonlinear techniques to ERG analysis, and will, in particular, address the question of how ERG data may be embedded in a multi-dimensional space; this is because such an embedding forms the starting point for several modern data processing methods. Moreover, the embedding itself can easily be visualized, showing ERG structures in a different way. Therefore, such ERG embeddings will be explored and discussed, since they might add to the powerful diagnostic capabilities an ERG examination yet provides.

Nonlinear analysis and chaos theory have been developed a bunch of techniques for cases where the examined system is not strictly deterministic, but displays more complexity than can be handled by traditional methods. These nonlinear techniques include, besides plain descriptive aspects, versatile data processing algorithms. Regarding the ERG, its high intrinsic noise level poses presumably the most fundamental problem to ERG analysis; suitable signal recovery techniques are always welcome. Thus, as the second topic of this thesis,

a recently developed noise reduction algorithm based upon the phase space embedding will be evaluated. Its capabilities and limitations will be compared to other noise reduction techniques more familiar to ophthalmology, that is, averaging and Fourier analysis.

The scope outlined above turned out awkward to accomplish with the tools currently on the market, be they hardware or software. It seemed more convenient to have an ERG recording system available that is tailored to the specific requirements, and thus such a system was developed.

## 1.2 ERG basics


The ERG is an electrical potential which arises in the retina after light stimulation. It is detectable all around the eye but, at the surface of the human body, it is largest at the center of the cornea. The ERG represents the composite activity of millions of retinal cells. The photoreceptors appear to make a definite contribution to the negative component, or *a-wave*, of the ERG, while cells in the inner nuclear layer appear to contribute more to the later *b-wave*, or positive component, of the ERG.

### 1.2.1 Normal ERG

The *ganzfeld* ERG allows to monitor the activity of the photoreceptors and the inner nuclear layer (*stratum ganglionare retinae*). The response, i. e., the change of the retina's electrical potential to a light stimulus, shows several distinguishable components, related to their areas of origin, as Figure 1.1 illustrates. In fact, the electrical potential is an integrated mass response made up of a number of independent components (GRANIT, 1947).

Over the years several different classification schemes have been established to describe these components. The classification most often used in clinical practice is given below (NIEMEYER, 1998).

***a-wave*:** Reflecting the mass activity of the photoreceptors, the *a-wave* can be related to phototransduction by assessing changes in its slope (HOOD and BIRCH, 1990).



**Figure 1.1:** The response of the retina to a single flash of light can be projected on a histological slice of the retina itself, indicating how each layer contributes to the generation of each ERG component. Figure E. DELUCA.

***b*-wave:** The *b*-wave is generated by second-order neurons, and reflects, most likely, the activity of the MÜLLER cells. At least in amphibian retina, the responses of the MÜLLER cells to changes in  $K^+$  have been shown to be responsible for the generation of the *b*-wave, driven by a transient increase in extracellular  $K^+$  originating from depolarizing (ON) bipolar cells, and regionally restricted to the outer plexiform layer (NIEMEYER, 1998). Studies by SIEVING *et al.* (1994) yielded evidence for contributions to the normal ERG *b*-wave from both the depolarizing ON and from the hyperpolarizing OFF second order neurons.

***c*-wave:** The *c*-wave of the ERG follows the *a*- and *b*-wave and reflects the balance between the positive response of the retinal pigment epithelium (RPE), and a negative response of the MÜLLER cells to the light-evoked decrease in  $K^+$  ions in the subretinal space (STEINBERG *et al.*, 1980). The *c*-wave can only be recorded by DC-coupled setups due to its slow time course (several seconds); this seems to preclude its routine recording in patients.

**Oscillatory potentials:** These are superimposed on the rising phase of the *b*-wave and originate from the inner part of the *stratum ganglionare*, most likely from amacrine cells at stimulus levels that elicit responses from the rod and from the cone system.

## 1.2.2 Clinical applications

The two photoreceptor systems differ in structure, retinal circuitry, sensitivity to light, and spatial resolution. Evidently, both systems can be affected selectively or together, in a remarkably variable proportion in a wide spectrum of disorders. The pathophysiology of these phenotypically variable conditions determine the patient's individual problems. Therefore, in order to track down retinal disorders to their physiological origins, it is meaningful to examine the rod and cone system separately. Doing so provides also information about the topography of retinal dysfunctions, due to the different spatial distributions of rods and cones.

After dark adaptation, the scotopic rod system of the peripheral retina dominates the ERG response when stimulated with blue light. In light-adapted state,

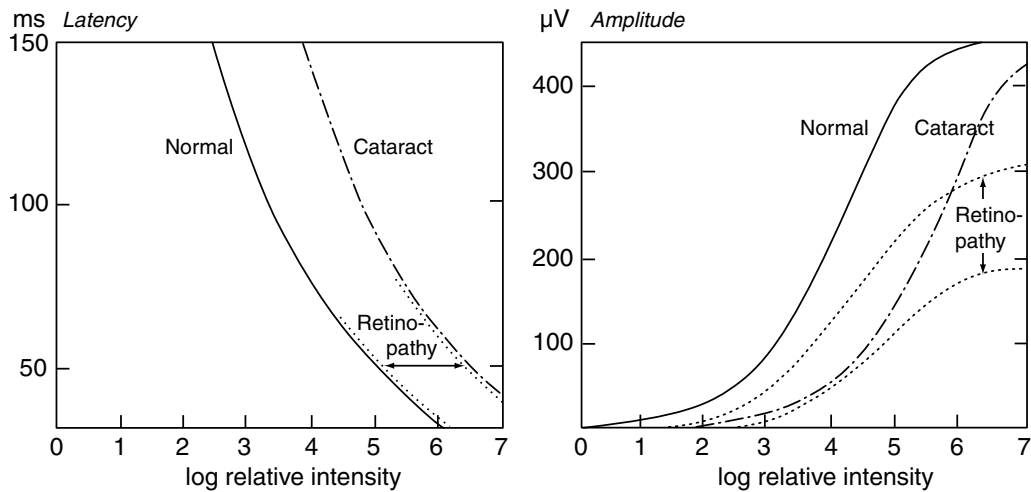
<b>Characteristics of the rod and cone system</b>	
<b>Rods</b>	<b>Cones</b>
<b>Topography and function</b>	
<ul style="list-style-type: none"> <li>• peripheral retina</li> <li>• dark vision</li> </ul>	<ul style="list-style-type: none"> <li>• central retina</li> <li>• light vision</li> </ul>
<b>Frequent dysfunctions</b>	
<ul style="list-style-type: none"> <li>• dark adaptation</li> <li>• peripheral field of view</li> <li>• glare sensitivity</li> </ul>	<ul style="list-style-type: none"> <li>• color vision</li> <li>• central field of view</li> <li>• visus</li> </ul>
<b>Sensitive parameters</b>	
<ul style="list-style-type: none"> <li>• amplitude/latency of the rods' <i>b</i>-waves</li> </ul>	<ul style="list-style-type: none"> <li>• amplitude/latency of the flicker ERG and the cones' <i>b</i>-waves</li> </ul>

**Table 1.1:** Clinically important characteristics of the rod and cone system (ZRENNER, 1983).

the photopic cone system of the central retina dominates the response, showing higher sensitivity to red light.

Changes in the shape of the responses indicate whether a defect resides within the dioptric system before the retina, or in the retina itself. Defects that prevent the light from reaching the retina, as in cataract, cause a general attenuation of the response's amplitude. At first glance, the latencies seem to be affected as well; however, the observed increase of the latencies is most likely due to the reduced intensity level that illuminates the retina after passing through the opaque lens, and does not indicate a dysfunction of the retina itself.

The *ganzfeld* ERG is, by its nature, not able to reveal small local retinal lesions, where most parts of the retina remain unaffected, e. g., chorioretinal scars. On the contrary, panretinal lesions or degenerations may show serious changes in both amplitude and latency of the ERG responses. The ERG is, in fact, the method of choice to verify congenital retinal degenerations like RP. A significant increase of both latency and amplitude can also be observed in retinopathia, as Figure 1.2 shows.



**Figure 1.2:** The relationships (continuous lines) between the latency to the peak of the *b*-wave of the ERG (ordinate left), the amplitude, measured from the peak of the *a*-wave to the peak of the *b*-wave (ordinate right), and the logarithm of the intensity of white light stimulation in relative units (abscissas) for normal dark adapted subjects. The dash-dot lines indicate hypothetically the direction in which the relationships of subjects with retinopathies (dot) and opacities of the dioptric media, such as cataracts (dash-dot), tend to go (GOURAS, 1970).

### 1.2.3 ERG varieties

Since the 1900s, when EINTHOVEN and JOLLY<sup>1</sup> used a string galvanometer to obtain records of the frog eye, ERG techniques nowadays cover at first sight an almost confusing diversity of different diagnostic methods, each of them with its own advantages and pitfalls, for dedicated applications in both clinical and research fields. BACH and KELLNER (2000) give a comprehensive survey of electrophysiological diagnostics in ophthalmology.

## 1.3 Why yet another ERG system?

As stated above, the ERG has become a widely accepted clinical examination to diagnose pathological changes within photoreceptors and the retina. However, these techniques suffer from intrinsic physical and physiological noise phenomena, which may exceed the signal amplitude by many times. Subtle pathological changes are completely buried in many cases, and thus traditional approaches

<sup>1</sup>(EINTHOVEN and JOLLY, 1908)

to recover such changes often fail. Hence, from the clinical point of view, it is desirable to have state of the art signal processing and identification techniques at hand, e. g., for the early diagnosis of rod and cone degenerations. Tiny pathological ERG amplitudes challenge the detection capabilities of commercially available systems, and all too often they escape our attention completely simply because the setup lacks sensitivity. Several attempts to overcome such obstacles have been made. For example, RISPOLI *et al.* (1994) enhanced their detection rate of flickers from RP patients by some 30 %; they achieved this mainly by mechanical and electrical improvements of the electrode/amplifier setup.

Computer aided systems for recording ERGs and VEPs have been established in ophthalmology for more than 20 years, taking advantage of modern data storage and documentation equipment. However, in today's commercially available ERG devices, the processing power of state of the art computers is being almost wasted. We find Digital Signal Processing capabilities only marginally employed, just to the extent required by the standard of the *International Society for Clinical Electrophysiology of Vision* (ISCEV), which has become a standard for clinical research and a benchmark for commercial products. The ISCEV standard (MARMOR and ZRENNER, 1999) has been established to assure equivalence and comparability between recordings taken by very different systems. Though the standard does encourage further technical developments, it is actually some common denominator to cover all currently operating equipment, rather than to promote computationally demanding tasks. Therefore, any modern computer ERG recording system which satisfies the guidelines spends most of its time idle.

In the engineering sciences, a variety of software tools is available for signal processing. Unfortunately, these packages lack either a front-end for directly recording data from real-world sources (because this requires additional hardware), or do not employ the recently developed nonlinear techniques.<sup>2</sup>

---

<sup>2</sup>For an overview of what equipment is currently available, check <http://www.iscev.org/links.html>

## 1.4 Calypso

To overcome the drawbacks mentioned above, I developed a new ERG recording device in order to integrate many of today's desirable features in one system. This system, which will be described below, is referred to as *Computerized Analysis of Physiological Samplings in Ophthalmology*, abbreviated CALYPSO.

### 1.4.1 Functional overview

Despite some variation in details, most modern ERG recording systems share the same basic, chained structure (HOGG and RILK, 2002) listed below; in this respect, as far as the hardware setup is considered, CALYPSO makes no exception.

1. The **stimulator** provides for the light stimuli needed to evoke the retinal responses. Due to very different diagnostic requirements, stimulators can range from simple flashlights for a *ganzfeld* stimulus to sophisticated CRT monitors, which allow the generation of colored and spatially resolving stimuli.
2. **Electrodes** have to maintain a proper, electrically conductive, and mechanically stable, contact between the subject and the recording unit. Although they might seem to be the simplest (and cheapest) part of the system, they determine the quality of the signal in an essential manner: they have to pick up the tiny fractions of a volt we are interested in, and run them unchanged through meters of copper cable to the amplifier. Exposed to mechanical stress and chemical detergents, they rarely live long.
3. The **amplifier** increases the signal level from the fractions of millivolts the retina delivers to the range of volts that are required by the ADC stage. Besides, higher signal levels are less prone to artifact pickup from various sources when running through lengthy cables. Electrophysiological amplifiers generally feature a *differential input* configuration that sets them apart from, e. g., audio-amplifiers; such a configuration is very useful for suppressing a nasty sort of noise, called common-mode noise. Most often, pre-filter stages are included prior to the amplifier to prevent the amplifier



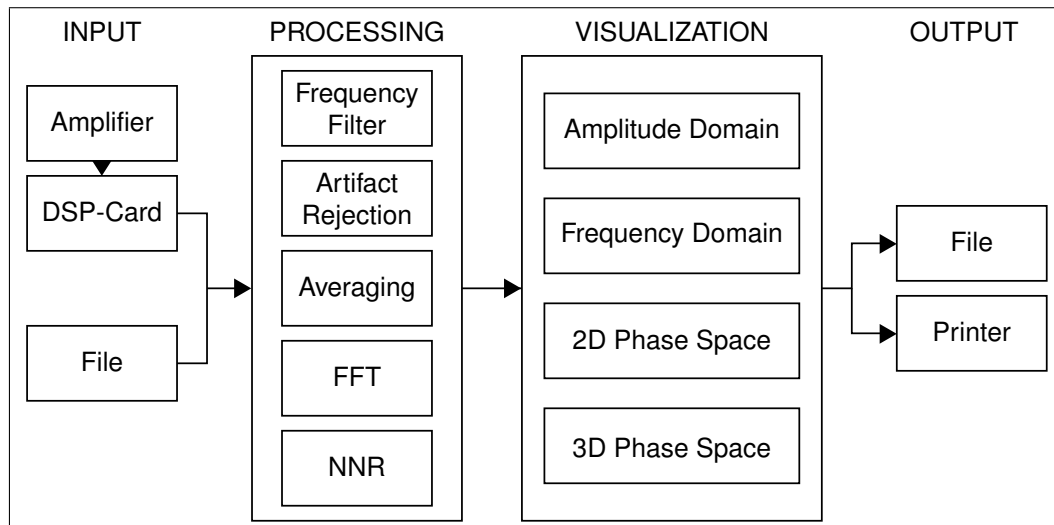
---

from saturating, to reject artifacts, and to narrow the frequency bandwidth to the interesting frequency range.

4. The **ADC** stage translates the so far analog signals into a digital representation the computer can understand. An ERG examination can of course be performed without an ADC, and it was done so until the advent of computer technology, but in still doing so we would sacrifice all advantages of modern data analysis.
5. The operation of a contemporary ERG system is controlled by means of a **computer**, which also provides the tools to process, to display, and to store the acquired data.

Before going into CALYPSO's own specific details, a feature list may be convenient. Some of the listed items may not yet be familiar to clinical researchers; these will be explained later on. Since the present work is supposed to be a medical thesis, I will not stress too exhaustively the plain technical matters of hardware and programming (although they are as interesting as the scientific aspects), but rather outline some selected topics that are important to grasp how the system basically works.

- Hardware front-end capable of sampling two channels with variable rates from 600 Hz to 7800 Hz.
- A default set of software frequency filters, easily expandable.
- Averaging (coherent averager with user selectable number of resulting output responses).
- Artifact rejection tools.
- Frequency analysis by means of the Fast Fourier Transform (FFT).
- Phase space visualization plots.
- Nonlinear noise reduction (NNR).
- Print utility and ASCII based data export function.



**Figure 1.3:** Architecture of CALYPSO. The arrows indicate the data flow through the system. Refer to Figure 2.5 for a different, more hardware-oriented system scheme.

The system architecture, which is quite similar to the one of a commercial system, is shown in Figure 1.3. We will begin our brief “guided tour” at the left, where the data enter the system. This may be done in two ways: either offline from a previously recorded file, or by an online ERG recording session. The latter case requires a stimulator and physiological amplifier. Both devices are standard electrophysiological equipment, readily available commercially; they will be described in Section 2.4.

The next stage following the amplifier is the Digital Signal Processor (DSP) plug-in board, that contains, among other things, the ADC stage explained above. Here the analog data are translated into a digital representation. The DSP also performs some data buffering and critical real-time processing tasks, for example, trigger evaluation and synchronisation to the stimulus events.

The analog data are then digitized by the ADC, and are made available to the next stage, i.e., the personal computer. This computer hosts CALYPSO’s processing and visualization software, which, as well as the DSP board, was developed by the author. Once the data are stored in PC memory, the standard ERG processing techniques may be applied, for instance, frequency filtering, averaging, and rejection of artifacts. Furthermore, some advanced techniques are provided, e.g., Fourier transform and a nonlinear noise reduction algorithm. These items will be discussed in detail below. The results of each manipulation

can be visualized by means of a familiar ERG amplitude diagram, and additionally by some advanced visualization methods, namely phase space plots. A frequency domain representation is also available that allows some frequency filter analysis. A save dialog provides for saving both raw data and processed data to the system's hard disk in portable ASCII format. A printer routine is also included.

The aim of this thesis is to explore and discuss phase space embeddings of the flicker ERG, and, to compare new noise reduction techniques basing on these embeddings to techniques that are yet common to ophthalmology. The clinical applicability will be illustrated by some examples.

# Chapter 2

## Materials and methods

The nonlinear methods here applied consider geometric properties of the data, rather than relying on a more algebraic approach as other common signal recovery algorithms do. This chapter covers the basic ideas of nonlinear time series analysis, as far as they match the scope of this thesis. First, the issues of constructing an embedding from ordinary time series data will be addressed. After we will have seen how an ERG embedding looks like, the principles of the Nonlinear Projective Noise Algorithm as an application of nonlinear techniques will be explained.

### 2.1 The delay embedding

In recent times, nonlinear analysis has been successfully applied to data from such different fields as astronomy, meteorology, economics, and biology. Moreover, several papers have been published addressing its use on medical data, especially electrophysiological data from the electrocardiogram (FELL *et al.*, 2000), but also on neurological topics (DINGWELL and CUSUMANO, 2000), human breath rate, and epilepsy (QUYEN *et al.*, 1997). Surprisingly, a survey of the ophthalmologic literature finds no evidence that these methods have ever been applied to ophthalmological tasks, for example, to improve the diagnostic significance of the ERG. Although, in most cases, the high intrinsic noise level discourages us from doing a comprehensive<sup>1</sup> nonlinear analysis of ERG data, there

---

<sup>1</sup>Here, "comprehensive" means the "classic", standard set of nonlinear tools used to analyze the data. These include LYAPUNOV exponents, the correlation sum, the correlation integral,

are still some nonlinear techniques available which can be used to improve data quality significantly, for example, nonlinear algorithms that address the noise problem. CAWLEY and HSU (1992), and SAUER (1992) have independently introduced the noise reduction algorithm that will be applied here.

### 2.1.1 Constructing delay vectors

When performing electrophysiological diagnostics, we rely on the assumption that the retina's responses to uniform stimuli will be similarly shaped. Ideally, i. e., under noise-free conditions and with a completely known physiology, all responses to uniform stimuli are supposed to be identical. Although this cannot be taken *a priori* for granted, it seems nevertheless a reasonable assumption to make, because otherwise the ERG technique would provide little diagnostic reliability. We can state this underlying assumption by other words: If we know the entire initial state of the system at a given time, we can predict the system state for any time. Such systems are referred to as *deterministic*.

In order to gather information about a dynamic system like the retina, we usually measure some time dependent properties of the system. The change in the electrical potential of the retina when stimulated by a light flash is well suited for this task. This change in potential is the result of the interaction of several non-observable internal variables that uniquely determine the system's behaviour. Thus, examining only one variable (for instance, the electrical potential), does not recover all information about the system. In the physical notion, the system dynamics can be modeled by assigning a geometric dimension to each independent system variable. A simple example might be a swinging pendulum. To describe its dynamics, we first might want to draw a two-dimensional coordinate plot. Then we can assign the pendulum's position at a given time to the abscissa and its momentum to the ordinate. Two dimensions are enough to define the motion of a simple pendulum, since such a pendulum has only two independent variables: its position and its momentum. The higher is the number of independent variables, the more dimensions will be necessary to describe the system's dynamic behaviour. A one-dimensional representation is obviously not sufficient for representing such complex biological systems like

---

and the fractal dimension. Unfortunately, most of these techniques are quite sensitive to noise contaminations and will fail if the data set is too noisy.

the human retina.


Consequently, if we manage to reconstruct the multi-dimensional properties from the single (scalar) time series, we could gain more information about the underlying system. In fact, such a reconstruction is possible; we do have tools to recover even a full-dimensional representation of the underlying system from a scalar time series only (TAKENS, 1981). One of the techniques to accomplish this entered common usage some twenty years ago; it is known as the *delay coordinate reconstruction* (PACKARD *et al.*, 1980).

What has to be achieved is a transformation of the time series into vectors, or points<sup>2</sup>, respectively, in a high-dimensional coordinate space, called the *phase space*, or *state space*. The technical term for this transformation is *embedding* the system. We will see shortly how these vectors can be constructed from the time series. Each vector corresponds to an unique system state, and this is why they are called *state vectors*. If the state changes, the vector coordinates will change, too, constituting a new state vector. Thus, the entire set of state vectors of a dynamic (i. e., state changing) deterministic system will form a curve called *trajectory* in phase space; the trajectory closes to a loop if the system is periodic. Think of the trajectory as a LISSAJOUS figure that sometimes can be viewed on an oscilloscope screen.<sup>3</sup> To get an idea of how a trajectory looks like, examine Figure 2.2.

How are the state vectors created? Let  $U(t)$  be the continuous voltage response from the retina.  $U(t)$  is sampled by means of a computer with the sampling frequency  $f_s$ . The result is a time series that consists of single data points  $U(t_i)$ , equally spaced in time by  $t_s = 1/f_s$ . Next, we have to choose two parameters: first, a fixed time interval, the *delay*  $\tau$ , which has to be an integer multiple of  $t_s$ . The second parameter is the *embedding dimension*  $m$  of the phase space the vectors are to be embedded in. We will discuss the proper choice of  $\tau$  later on, together with how we choose  $m$ ; let us simply assume here that such a choice has already been made. To build the  $m$ -dimensional state vector  $\vec{d}(t_i) = \vec{d}_i$  associated with  $U(t_i)$ , we simply pick  $U(t_i)$  as the first vector element. Then we go back in time by  $\tau$  to get  $U(t_i - \tau)$  as the second element, i.e., the second

<sup>2</sup>Because vectors and points are essentially the same in this context, the terms "vector" and "point" will be used interchangeably.

<sup>3</sup>In fact, LISSAJOUS figures are just simple trajectories in a two-dimensional phase space, and they are physically generated from a sinusoid just like state vectors are constructed from ERG signals.



**Figure 2.1:** Construction of the state vector  $\vec{d}_{15}$ . Here, the delay is chosen to  $\tau = 3$ , with the embedding dimension  $m = 4$ .

vector element is the voltage that happens to be by  $\tau$  earlier in time than  $U(t_i)$ . The third element is  $U(t_i - 2\tau)$ , being  $2\tau$  earlier in time than  $U(t_i)$ , and so on, until all  $m$  elements are found. Figure 2.1 illustrates this way of constructing a state vector.

As equation, the state vector  $\vec{d}_i$  at a given time  $t_i$  with the delay  $\tau$  and the dimension  $m$  reads:

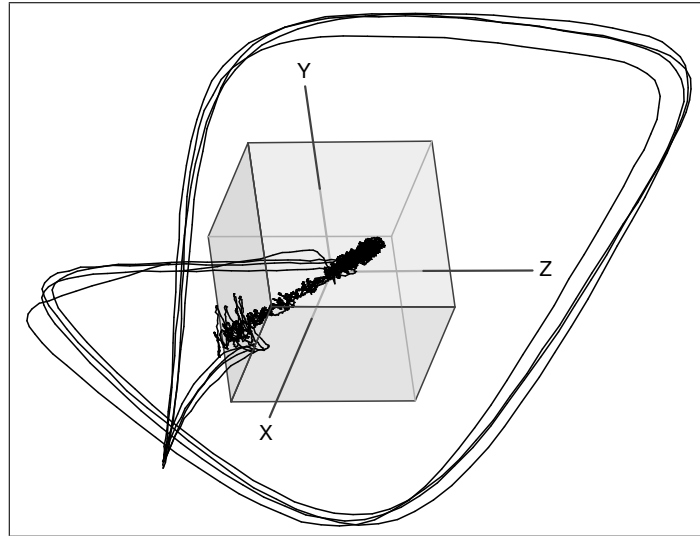
$$\vec{d}_i = \left[ U(t_i), U(t_i - \tau), U(t_i - 2\tau), \dots, U(t_i - (m - 1)\tau) \right]$$

After performing this procedure for all  $U(t_i)$ , we end up with a complete set of  $m$ -dimensional state vectors, one for each  $U(t_i)$ , constituting a trajectory in an  $m$ -dimensional phase space.

Let us have another look at Figure 2.1. The elements of  $\vec{d}_i$  cover a time range of  $(m - 1)\tau$ , which is designated the *embedding window*. We can think of each vector capturing the unique signal dynamics, or the *pattern*, within its embedding window; and we can watch the pattern changing as we slide the embedding window along the time series to build new vectors. This picture provides the first insight why the whole concept should work. The mathematical justifications will be outlined below.

### 2.1.2 The look of a trajectory

Fortunately, we are not constrained to explain everything only with algebra—state vectors are easy to visualize (at least their two-dimensional projections are). In Figure 2.2, trajectories from a cone ERG are embedded in a three-dimensional phase space, projected on the two-dimensional paper; the cube has been



**Figure 2.2:** Four successive cone ERG responses embedded in a three-dimensional phase space.

added to indicate the three axes  $x, y, z$ . We can identify four different signal periods. The trajectories actually consist of single discrete points (pointed to by their state vectors), rather than being continuous curves as shown; however, in order to be able to distinguish between different trajectories, adjacent points have been connected by lines. The enclosing loop represents the fast fluctuating patterns within the ERG's  $a$ - and  $b$ -waves and contains most of the relevant clinical information. The bar-shaped, slightly oscillating structure near the origin is due to the slow baseline fluctuations when, after a response, the ERG returns to the baseline level.

### 2.1.3 The theoretical point of view

Now that we have seen how a trajectory in phase space is constructed, it is important to realize that the axes do *not* address dependencies between different physical quantities like time, voltage, sensitivity, angles, etc., in contrast to other, more familiar, data plots. Instead, all axes are scaled in the unit of the ordinate of the former time series, that is: voltage. This may be the most serious drawback, since we are commonly used to judge measurements by how their different physical parameters relate; we are not used to have all axes scaled in the same measure. On the other hand, thanks to TAKEN's delay embedding theorem (TAKENS, 1981), the big advantage of such a representation is that it can



show the complete dynamic behaviour of the underlying system. Or, as MEISS (2000) states:

“Phase space is the collection of possible states of a dynamical system. [...] Implicit in the notion is that a particular state in phase space specifies the system completely; it is all we need to know about the system to have complete knowledge of the immediate future.”

What is more, the phase space representation is not only just another visualization tool, but it allows the use of a wide range of new techniques for analyzing and processing data. In Section 2.2 we will turn to an application of one of these techniques in detail.

So far we have concentrated on the plain process of constructing trajectories from time series data. But how can we be sure that this method is applicable to ERG data, without loss or, perhaps even worse, misleading distortion of the underlying information? As mentioned above, embedding techniques form a fundamental starting point for many applications, including physiological ones, e. g., the electrocardiogram (ECG). However, surely some more thoughts are required than merely to state that, if these methods are suitable for the ECG, this must be true for the ERG as well.<sup>4</sup>

Fortunately, there are a number of theorems dealing with the theoretical issues of when and how embedding leads to reasonable results. Though it may seem surprising that a delay embedding of a simple scalar time series and a “true” state space representation (with eventually much more variables to measure) bear the same information, the fundamental proofs were given already one or two decades ago. In the case of ideally deterministic systems, the embedding theorems by TAKENS (1981) and their generalization by SAUER *et al.* (1991) assure the equivalence between the space constructed by delay coordinates and the true phase space the system lives in, under fairly broad conditions.<sup>5</sup> Moreover, STARK (1999) was recently able to further extend the embedding theorems to even cover a wide variety of *periodically forced systems*. This provides the

---

<sup>4</sup>At least there is a small, but mathematically important, difference between the ERG and the ECG: The ERG is driven by a well-known constant stimulus frequency, while the driving mechanism of the ECG is some kind of “black box”; it is *self-driven*. Curiously enough, the mathematics concerning a self-driven system is much easier to handle as the one of a periodically forced system, even if the driving frequency is constant and exactly known.

<sup>5</sup>Refer to OTT *et al.* (1994), Chapter 5, for more on the theory of embedding.

theoretical justification of the approach taken in this thesis; recall what the ERG actually is: the response of a physiological system forced by a periodic sequence of light stimuli.


## 2.2 Nonlinear projective noise reduction

Glancing back to Figure 2.2, it is remarkable that the trajectories of each response are similarly shaped and close to each other. This suggests that the dynamics is basically deterministic, which yields regular cycles, while the noise causes the trajectories to deviate from the basic cycle, the *attractor*. Now, if we manage to unveil the attractor, we will be able to reconstruct the noise-free time series signal. In fact, we can more or less precisely approximate the attractor by the *Nonlinear Projective Noise Reduction* (NNR) technique described by SCHREIBER and RICHTER (1999), which has already been applied to fetal ERG signals fairly successfully (RICHTER *et al.*, 1998). The following section will explain its basics. For a detailed description see KANTZ and SCHREIBER (1998), and SCHREIBER and RICHTER (1999).

### 2.2.1 NNR basics

As discussed before, the phase space representation is valid as long as the embedding dimension  $m$  is large enough. Basically, there is no upper limit for  $m$ , while a minimal dimension  $m_0$  exists that  $m$  must not fall below of. If this rule is followed, the phase space will contain the entire dynamic behaviour of the *noise free* system. Any further increase of  $m$ , say up to  $n$  dimensions ( $n > m_0$ ), will, at least in theory, not unveil additional information. However, if we nevertheless observe new structures—and we do in any real-world noisy system—, these new features are caused only by uncorrelated noise. In the following discussion, we will call the first  $m_0$  dimensions *valid dimensions*, and the remaining  $n - m_0$  “noisy” dimensions the *nullspace*.

The idea of the NNR is as follows. First, we determine the minimal dimension  $m$  as described below. Then we add some more dimensions (how many, depends on the data) and embed the data as usual. As stated above, the trajectory characteristics in the additional dimensions represent the noise. We thus can get



**Figure 2.3:** A neighborhood of a point  $\vec{p}$  (filled circle) in a 3-dimensional space. Assume the Neighborhood Radius (which defines a sphere that is not shown here) chosen large enough that the neighborhood contains all points depicted. The neighborhood is distributed along some principal directions that equal the axes of an ellipsoid. Each principal direction corresponds to the longer axis of one single ellipse, and the length of this axis is termed the *size* of this dimension. The largest principal directions form the subspace on which the point  $\vec{p}$  is projected to, in order to get an estimate of the local attractor. Note that the principal directions are perpendicular to each other, and that the center of the ellipsoid is not necessarily identical to  $\vec{p}$ .

rid of the noise by projecting the points on the subspace that consists of the valid  $m_0$  dimensions. Unfortunately, we cannot tell *a priori* which  $m_0$  dimensions out of the set of  $n$  dimensions are valid, and which are not. This has to be decided for each  $n$ -dimensional point separately and is governed by the local shape of the attractor in the vicinity of the point considered. Because the attractor itself is simply a set of points, examining the local attractor basically means examining all points within the *neighborhood* around the point under consideration.<sup>6</sup>

In the mathematical notion, a neighborhood of the point  $\vec{p}$  is the subset of points  $\vec{p}_i$  with distances  $r_i$  to  $\vec{p}$  smaller than, or equal to, a given distance  $r_0$ . This distance is referred to as the *Neighborhood Radius*, defining a sphere (or a *hypersphere* in a multi-dimensional space) around  $\vec{p}$ . Depending on how large we choose  $r_0$ , the number of elements within the neighborhood will change. The crucial point is that if we make  $r_0$  larger than the noise level, the neighborhood will reveal a geometrical structure along one or more principal directions (see Figure 2.3). The idea to be understood is that, provided the noise level is not too large, the  $m_0$  *largest* directions are an estimate of the local noise-free attractor geometry

<sup>6</sup>In fact, the idea of neighborhoods is very common to nonlinear analysis; almost all nonlinear methods exploit local neighborhood relationships.

near the point  $\vec{p}$ , and the  $n - m_0$  *smallest* dimensions build up the nullspace. For example, consider a set of points in a  $n$ -dimensional coordinate system. Actually, they are all supposed to inhabit a, say, two-dimensional plane defined by two principal directions, but happened to be distorted by some stochastic process, i. e., by noise. Thus, rather than populating their original plane alone, the points have been scattered off the plane into all  $n$  dimensions, as illustrated in Figure 2.3. Note that, since this deviation is uniformly distributed, both the *center of mass* of all points and the principal directions remain unaffected. We can recover the true position of the points within the plane by back-projecting them orthogonally on the plane, but in order to do this, we have to know the orientation of the plane beforehand. Geometry tells us that a plane is uniquely defined by one point within the plane (here: the center of mass), and two vectors (here: the two largest principal directions). Therefore, we have to compute the center of mass and the two largest principal directions to get an estimate of the desired plane.

The calculation of the center of mass is straightforward, and the principal directions can be derived by the method of *Principal Component Analysis*, sometimes also referred to as *Singular Value Decomposition*, or *Empirical Orthogonal Functions*. Once we know these directions, we have an estimate of the "true", i. e. noise-free, position of  $\vec{p}$ . Performed on all state vectors, we get the desired cleaned trajectory, which can be easily transformed back into a time series. By the way, NNR is not constrained to a delay embedding, but rather works with any valid embedding of a stationary data set.

### 2.2.2 Estimating the parameter values

The embedding theorems by TAKENS (1981) and STARK (1999) guarantee that, under certain genericity conditions, a delay embedding is valid whenever the embedding dimension  $m$  is chosen large enough. In practice, we determine  $m_0$  by means of some supplementary nonlinear techniques, for instance, the *method of false neighbors* (KANTZ and SCHREIBER, 1998). Refer to Section 4.4.1 for more on finding the proper parameter values.

The estimation of the delay  $\tau$  is done in a similar way. Here, an appropriate technique is to plot the *autocorrelation function* of the data vs. the delay and then to choose the delay value where the autocorrelation function drops to  $1/e$ . We

can also estimate  $\tau$  simply by visually inspecting the trajectory, because the trajectory's global shape is quite sensitive to changes in delay. As it turns out, a "good-looking" unfolded trajectory, obtained by tuning  $\tau$ , is quite a good criterion and matches the results revealed by the autocorrelation method.

NNR likes very noisy data better when the number of dimensions on which the trajectory is back-projected is chosen smaller than the number of valid dimensions  $m_0$ . This means in practice that we have to distinguish between the proper, analytically derived  $m_0$ , which provides an overall embedding free of trajectory intersections, and another, smaller, number of dimensions, referred to as *projecting dimension*  $m_{\text{proj}}$ , that yields good NNR results. Surprisingly, best results are obtained when  $m_{\text{proj}}$  is set to the lowest reasonable number possible,  $m_{\text{proj}} = 1$ . Why so?

It is because a trajectory looks different at a global scale and at a microscopic, local scale. While the macroscopic shape of a trajectory may be rather complex and may require a lot of embedding dimensions  $m_0$ , it can always be approximated at microscopic scales by many short but straight lines, given that the scale is small enough. The inherent dimension of such a line is 1, regardless of how it is oriented in its  $m_0$ -dimensional embedding space. The projection of points on such a line is therefore a valid approximation of the true trajectory; it is exactly what is done by NNR when the Neighborhood Radius is small and  $m_{\text{proj}} = 1$ .

## 2.3 Reducing the recording time

A major concern about ERG recordings is their intrinsic high noise level, even in healthy, relaxed subjects. Moreover, in retinopathies, the amplitude drops compared to the noise as the retinal dysfunction proceeds. The reason why we do not see any signal in the final state is presumably not that the ERG response has been extinguished completely, but the decreasing signal-to-noise ratio (SNR), supposed that the noise level remains constant.

We can overcome this problem by adding more *redundancy* to the data. The classic approach, signal averaging, simply takes as many responses as possible into account. If we assume that the signal shape of each response is constant in both amplitude and phase, the uncorrelated noise will cancel out when we

average over all responses. Unfortunately, the efficiency of averaging increases only with the square root of the number of responses, i. e., to improve the SNR by, say, a factor of ten, it is necessary to evaluate a hundred times responses more. This means that the total recording time will increase considerably. The drawbacks are obvious: ERG baseline drifts, rising electrode impedance, and poor patient compliance—and each of them affects signal quality.

Increasing the recording time is only one way of improving signal redundancy. Another way, of which we will take advantage here, is not to acquire more responses, but to sample more values of each response, i. e., to choose a higher sampling rate. In doing so we gather more information about the signal, but not at the expense of prolonging the total recording time. Note that increasing the sampling rate does not improve the efficiency of averaging, and this may be why commercially available ERG equipment rarely provides for sampling rate adjustments.

Now, what is the benefit of a higher sampling rate? From the NNR point of view, it is desirable to have as many points as possible in a neighborhood. This assures that the true shape of the local attractor is estimated correctly. The more neighbors are available, the smaller is the statistical error of the estimation, that is, the deviation of the estimation from the expected value. By a higher sampling rate, we increase the number of data points per ERG response and, correspondingly, the number of points in state space, which in the end leads to an increased neighbor density.

## 2.4 Experimental setup

### 2.4.1 Hardware and software

Besides the software suite, the CALYPSO system contains several hardware components. The physiological pre-amplifier was chosen to be the *Universalverstärker* #224701 by Jaeger-Toennies<sup>7</sup>, Höchberg, Germany. Analog data input to the PC is performed by a Digital Signal Processor (DSP) plug-in card, hosting a 33 MHz DSP (ADSP-2181 by Analog Devices<sup>8</sup>, Norwood, USA) with

---

<sup>7</sup><http://www.jaeger-toennies.com/jaeger.html>

<sup>8</sup><http://www.analog.com>

fixed point arithmetic. The software itself is hosted on a Personal Computer equipped with a Pentium III processor clocked at 600 MHz, and with 128 MB RAM.

The filter engines and the user interface were programmed using the Delphi programming language, version 4.0, provided by Borland Corp.<sup>9</sup>, Scotts Valley, USA. Some of the graphics output routines and higher mathematical functions are based upon the SDL Delphi Component Suite by Software Development Lohninger<sup>10</sup>, Wien, Austria. The DSP software was written in ADSP assembly language, available at Analog Devices.

All CALYPSO-specific hard- and software has been developed by the author. To get an idea how CALYPSO looks like, a screenshot of the user interface is depicted in Figure 2.4.

Some additional analysis and data plots were performed using the Scilab package by INRIA Rocquencourt<sup>11</sup>, France. Scilab is a free scientific software package for numeric computations.

### 2.4.2 Standard ERG recordings

It turned out convenient to have the ERG stimuli generated by a commercially available device. This is the PRIMUS system by Tomey Europe<sup>12</sup>, Erlangen, Germany, well suited for providing the stimuli for the standard ISCEV *ganzfeld* recordings. An optical sensor for stimulus detection attached to the PRIMUS *ganzfeld* dome allows for synchronization between the PRIMUS system and CALYPSO's data acquisition.

### 2.4.3 Safety considerations

Since CALYPSO runs from the mains power supply and is coupled to the human eye, safety was of major concern. To make sure that there is no hazardous voltage breakthrough, the amplifier was selected to fulfill the requirements of the Medical Devices Directive 93/42/EEC, Annex IX. In order to not deteriorate

---

<sup>9</sup><http://www.borland.com>

<sup>10</sup><http://www.lohninger.com>

<sup>11</sup><http://www-rocq.inria.fr/scilab/>

<sup>12</sup><http://www.tomey.de>

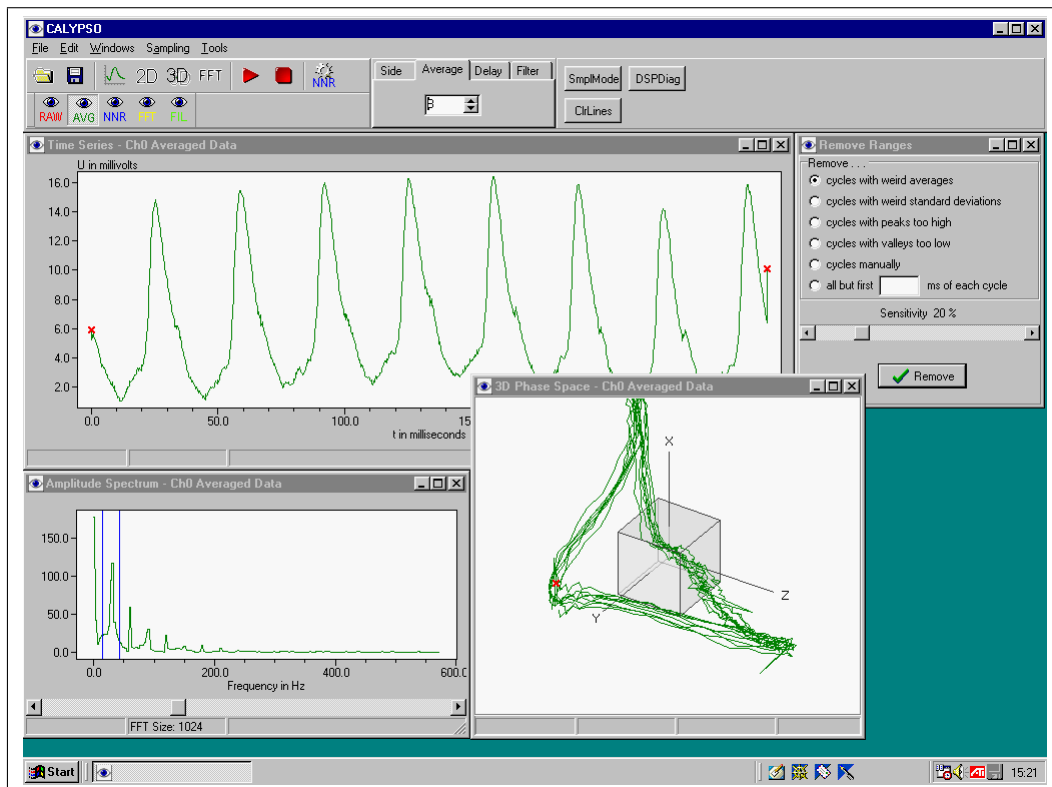


Figure 2.4: Screenshot of the CALYPSO user interface, displaying a normal 30 Hz flicker ERG.

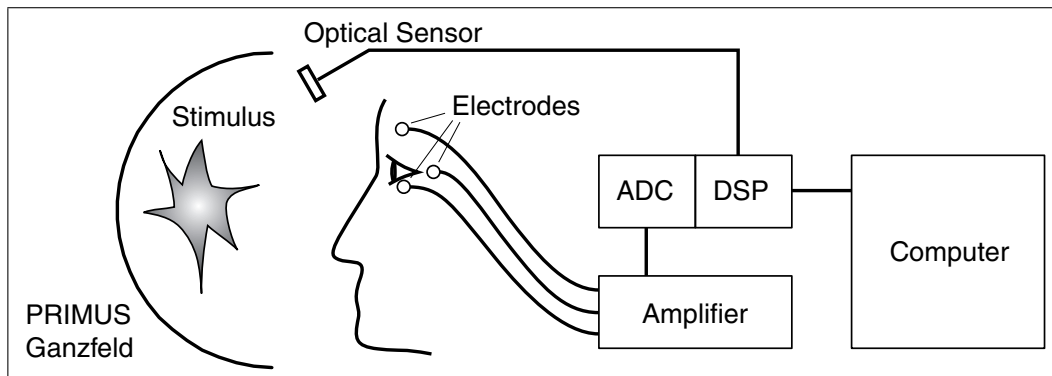
this safety concept, all connections to other stimuli generating equipment are decoupled by optical sensors which provide for galvanic isolation. Thus, the entire setup fulfills the requirements of the Medical Devices Directive 93/42/EEC, Annex IX, as certified by the *Gerätezentrum* of the University Hospital of Tübingen.

## 2.5 Participants

### 2.5.1 Ethical considerations

This study was conducted in accordance with the tenets of the Declaration of Helsinki. The study design was reviewed and approved by the Ethics Committee in human experimentation of the University Hospital of Tübingen. Normal subjects were recruited via advertisement in the University Eye Hospital in Tübingen. Patients were asked if they would like to participate when visiting





**Figure 2.5:** Hardware setup for the standard ISCEV recordings. For safety reasons, the CALYPSO system is decoupled from the stimulus generating equipment by an optical sensor which provides for galvanic isolation.

the RP outpatients' department at the hospital.

### Inclusion Criteria

Persons were allowed to participate in the study only if they met *all* of the following criteria:

1. Age between 18 and 50 years
2. Ability to provide written informed consent
3. Ability to fixate

### Exclusion Criteria

Participants were excluded from the study for *any* of the following reasons:

1. Disposition to increased intraocular pressure in the participant's history or in his/her family history
2. Known photopic ERG amplitude of both eyes less than 50% of a normal photopic ERG amplitude

The study has been designed as a case study. The number of normal subjects was  $N = 10$ , the number of patients was one or two for each specific retinal disease.

## 2.6 Examinations

DTL electrodes (DAWSON *et al.*, 1979) have been used as recording electrodes rather than contact lens electrodes, due to the advantages of a lower risk of corneal lesions (JACOBI *et al.*, 1993), and not requiring mandatory local anesthesia.

The examination procedure was as follows. First, the pupils were dilated by means of a mydriatic (Tropicamide 1 %, Mydriaticum Stulln®). The skin at the temples and the forehead was cleaned with a 70 % ethanol solution to accomplish good electrical transduction for the electrodes; these consisted of a DTL electrode which was placed contacting the bulbus within the conjunctival sac of each eye, a skin electrode (reference electrode) attached to the ipsilateral temple of each eye, and a skin electrode (ground electrode) attached to the forehead. Before attaching, each skin electrode was covered with conducting paste. In case the subject felt inconvenient with the DTL electrode, a local anesthetic (Novesine®) was applied on request. As the next step, the electrodes were connected to the amplifier. Now the participant was asked to take place in front of the *ganzfeld* dome, where the flicker stimuli were presented according to the ISCEV standard.

Note that all responses were recorded "as they are", i. e., except for a wide band-pass frequency filtering (0.3 . . . 1500 Hz) to reject DC voltages and aliasing effects, no additional signal processing, in particular artifact rejection and averaging, was performed online. Therefore, the signal to be analyzed offline was a time series, including all responses and their related stimulus events (and, besides, all artifacts). The sampling frequency was set to 3600 Hz. The total examination time including preparation took about 30 minutes.

## 2.7 Data processing

### 2.7.1 Frequency filtering

While averaging is not very sensitive to drift effects, phase space techniques are. Averaging is almost unaffected by slight remaining drift artifacts, but trajectories may suffer seriously. "Drift" is just another term for signal fluctuations of

very low frequencies, i. e., frequencies much lower than the stimulus frequency. Because an ERG response is phase-locked to the evoking stimuli, any frequency content below the stimulus frequency can be regarded as artifact. As a highly unwanted property, ERGs do show a considerable amount of drift, which can easily exceed the ERG response amplitude by many times.

Thus, the comparatively simple high-pass filters of common ERG equipment, which work just fine in combination with averaging, have been substituted by other, more efficient, filters in the CALYPSO setup. Care had to be taken not to affect the frequency range of interest, neither in terms of amplitude nor in terms of phase shifts. Since latencies are of major concern in ERG diagnostics, phase shifts may turn out as troublesome side-effects. ERG latencies are easily altered by carelessly applying filters of unknown phase behaviour; as a consequence, latency data may get corrupted and lose its value as fundamental diagnostic parameter.

A possible solution to the problem is to filter nevertheless, but by filters of both very steep filter slopes and well-known phase behaviour. This is fortunately an extensively addressed topic in signal processing theory, and, accordingly, very powerful software filter concepts have been developed (LYONS, 1997). One of the most effective ones are the *infinite impulse response* (IIR) filters. They include an internal feedback mechanism, and this turns out to be an advantage and a pitfall: While being powerful in terms of computing time and filter slope, they may suffer from instabilities and overflow errors that depend on the input data. Their effect on the output is therefore not fully predictable.

For these reasons, I decided to apply a related, but slightly less fancy form of filters: the *finite impulse response* (FIR) filter. Although FIR filters may require much more computing time than IIR filters do, their behaviour is fully predictable, and they are mathematically guaranteed to be stable, independently from the input data. What is more, they feature a constant group delay<sup>13</sup>, which is the key advantage where preserving signal latencies is mandatory. All filters in CALYPSO are digital FIR filters, except for the very first high pass and low pass filters at the input of the first amplifier stage. Their purpose is to keep the

---

<sup>13</sup>The group delay is actually the amount of time a signal is delayed while passing from the input to the output of the filter. If it is constant, that is, independent of the signal frequency, a broad-band signal will not suffer from phase shifts after being filtered. FIR filters have a constant group delay if their coefficients are symmetrical. All CALYPSO FIR filters are designed to fulfil this requirement. For more on digital filters, see, e. g., LYONS (1997).

<b>Processing parameters for the flicker ERG</b>	
Total number of responses	570
High pass filter (Hz)	10
Low pass filter (Hz)	—
Avg. responses	5
Delay (ms)	5.0
$m_0$	2
$n$	5
$m_{\text{proj}}$	1


**Table 2.1:** Synopsis of the processing parameters of the standard ISCEV flicker ERG. The first line gives the total number of raw responses considered for processing. Responses corrupted by artifacts are excluded and removed after the frequency filter operations. The next two lines contain the frequency filter settings for the high pass filter and the low pass filter. The line labeled "Avg. Responses" shows how many responses resulted after the averaging process. This number is normally 1, that is, the averager returns only one resulting response. However, for the flicker ERG, the raw data were averaged to five resulting responses to allow for further NNRing. The last three lines of the table provide the parameter settings for the phase space representations and NNR, respectively, as described in Section 2.2.

amplifier from saturating, which cannot be accomplished by any subsequent filter stage. Refer to Table 2.1 for the actual FIR filter settings.

FIR filters are basically a more or less large table of coefficients. Once the filter algorithm was programmed, new filter characteristics are easily invoked by just loading a new coefficient table.<sup>14</sup>

When a FIR filter's center frequency is required to drop below around 1 Hz, the coefficient tables become inconveniently large. This holds true especially for the high pass filters (DC blockers) that are used for the rod, maximum response, and cone ERGs. In these special cases, CALYPSO's Fourier transform routines provided a neat, alternative way to cut off the very low signal frequencies.

<sup>14</sup>A tool for automatically calculating such coefficients is readily available at <http://www.ittc.ku.edu/~rvc/projects/firfilter/normal/>



**Figure 2.6:** Synopsis of an unfiltered flicker sequence from a normal subject, consisting of a total of 570 responses. The large abscissa time scale leaves the single flicker responses unresolved in the plot. The sequence is actually supposed to cover an amplitude range of only about  $100\ \mu V$ ; artifacts from muscle activity are responsible for spreading the signal over a range ten times wider. Note the large blink reflex spike at the onset of the stimulus sequence.

### 2.7.2 Removing artifacts

There are few human beings that stand a vividly bright, rapidly oscillating flicker light over twenty seconds without any muscular reaction, even if no mydriasis was applied. The electric footprints of such muscular reflexes show up clearly in a diagram of the entire flicker time series, as illustrated in Figure 2.6. It was a matter of course to equip CALYPSO with an artifact rejection tool, because it would be a really annoying task to check each single flicker response by hand for possible artifact contamination.

From Figure 2.6, we see that these artifacts reveal themselves by their large amplitudes. The simplest way to remove them is therefore just to cut off all responses that contain voltage values above (or below, respectively) some pre-selected maximum (or minimum) limits. This is, in fact, the kind of artifact detection featured by several commercial ERG systems; however, reasonable cut-off values may change remarkably between different subjects, and even between different examinations of the same subject. A general cut-off limit is difficult, if not impossible, to agree on. With a fixed limit, we might remove 10% of all responses in one case, and 90% in another.

A more general approach (the one CALYPSO uses) requires a numeric property that can be computed for each response, for example, each response's average amplitude. The idea is to find a property that distinguishes just by its value an artifact-contaminated response from a valid one. If the property is numeric,

we can uniquely sort the responses, and if it is cleverly chosen, the rank of each response after sorting reflects the response's degree of contamination.

The sorting of the responses with respect to the property yields a *distribution*. The shape of this distribution depends of course on the selected property and may take any imaginable form, but often it may turn out to be two-lobed around a mean value. That is, the responses with the smallest and with the largest property values might be regarded as most severely contaminated. Once we have such a distribution, it is convenient to remove a given fraction, say, 5% of the responses with the most extreme property values. This approach ensures that always a fixed fraction of all responses is removed, regardless of how noisy the time series actually is.

For this purpose, CALYPSO offers currently four different properties to select from: Average amplitude, RMS amplitude, maximum voltage, and minimum voltage. I found that using the average amplitude criterion provided best results, at least on the flicker data yet tested.

It is quite easy to add new properties according to which the single responses are ordered. One might think, for example, even of sophisticated properties, like morphology or frequency considerations, as long as they allow the responses to be sorted in a unique order.

### 2.7.3 Averaging

Averaging is a simple, yet a powerful and robust, method of signal recovery. Many cycles of a noisy signal are added, and then the sum is divided by the number of cycles. By the way, this is exactly how the statistical measure *mean* is calculated.

The averaging process used here<sup>15</sup> requires the time phase of each cycle to be identical, otherwise the cycles would not correlate. Fortunately, the ERG provides this feature by its very nature, because the start of each response is determined uniquely by the external stimulus event. Note that this is by no means a matter of course—consider, for example, the ECG; but if this condition is met, the signal will always average to its true amplitude value regardless of the noise

---

<sup>15</sup>If the signal cycles are phase correlated (as is the ERG), the averaging process is known as *coherent, predetection, or vector averaging*.

level, if enough periods are included. The uncorrelated noise, however, will average to zero; the effect is an improvement of the signal-to-noise ratio (SNR). Here is a summary of the advantages of averaging:

- Averaging by itself will not add processing artifacts that are not already part of the input signal. The process of averaging is stable in this sense.
- Provided that the noise is uncorrelated to the signal and averaging is done on a sufficiently large number of responses, averaging is guaranteed to converge to the true signal, regardless how high the noise level was initially. Moreover, the performance is then even independent of the actual noise content.

It can be shown that the noise reduction  $R$  depends on the number  $N$  of averaged periods (LYONS, 1997). More precisely, the SNR after averaging improves to  $R_{\text{ave}} = \sqrt{N}$ , or in decibels:  $R_{\text{ave}} = 20 \log_{10} \sqrt{N}$  dB. The square root dependency causes averaging to lose its power somewhat at large  $N$ ; when  $N$  is already high, considerably more cycles are required to increase  $R_{\text{ave}}$  only slightly. Despite of this drawback, averaging plays an essential role in ERG processing. Note, by the way, that it is not necessary to always average in such a way that only one response results; an averaged signal of two, three, or even more responses is easily generated by a longer average interval. Of course, in doing so we would decrease the overall noise reduction performance, and so the benefit is not yet obvious. However, a subsequent NNR analysis works better with a couple of cycles when each of them is slightly deviated from the expected value, than it does on only one cycle, where the deviation is no longer apparent. We can anyway average the outcome down to one cycle as soon as NNR has finished. The averager of CALYPSO can therefore be preset to return any arbitrary number of resulting cycles, i. e., ERG responses. Again, refer to Table 2.1 for the settings.

# Chapter 3

## Results


This chapter presents what the previously described techniques are capable of when applied to flicker ERGs. Because an embedding of experimental data either can be used solely for a topographic analysis of the data structure, or may serve as base for further data processing, this chapter is divided into two parts. First, we will examine the appearances of normal flicker ERGs in a delay embedding, and try to describe their morphology in a meaningful way. Embeddings of some pathological flicker responses will be shown, and we will compare their different morphology to the normal ERG embedding. Second, we will see whether the application of the NNR method to flicker ERGs is feasible, and, if it is, on which conditions.

### 3.1 Standard ISCEV flicker ERG embedding

In a phase space representation we are mainly interested in the shape of the trajectory, rather than in absolute amplitude values; the analysis of amplitude differences is well covered by the standard ERG examination in terms of amplitudes and latencies.

$N = 10$  normal 30 Hz flicker ERGs of 10 different subjects were obtained according to Section 2.4. The RMS amplitudes of each flicker ERG were determined and averaged; the mean RMS amplitude was found to  $\bar{a}_{\text{RMS}} = 32.4 \mu\text{V}$ . Now the individual flickers were scaled, so that each individual flicker's RMS amplitude equals  $\bar{a}_{\text{RMS}}$ ; the idea was to clarify differences in the trajectory shape,






**Figure 3.1:** Different representations of a normal 30 Hz flicker ERG, averaged to five responses. The upper panel shows the familiar amplitude-time diagram. The lower left panel contains a two-dimensional delay embedding of the same data set. A three-dimensional embedding is displayed in the lower right panel; the orientation was arbitrarily chosen to show additional structures that are not visible in the embedding on the left. The delay was set to 5.0 ms in both embeddings. The circles indicate the position of the vector related to the stimulus event.

rather than differences in their absolute amplitudes, because a comparison of the shape of two graphs seems easier if the graphs have roughly the same scaling. Moreover, when we aim at finding new meaningful diagnostic parameters, it seems reasonable to start at maximum independence from the ones already used.

The result is shown in Figure 3.2. The mean  $\bar{s}$  of all flickers is represented by the bold trajectory, while the dashed lines are the individual flickers. Note that, despite the remarkable inter-individual variability, the flickers all seem to follow a basic shape that is fairly well represented by  $\bar{s}$ . A standard deviation is deliberately not given, because the ERG is considered to be non-parametrically



**Figure 3.2:**  $N = 10$  normal 30 Hz flicker ERGs (thin lines) and their mean  $\bar{s}$  (bold line), roughly centered within their bounds, which are also emphasized by bold lines. The circles denote where the stimulus event occurred on the trajectories; as time proceeds, the trajectories evolve counterclockwise. As a measure of the individual flickers' variation around their mean trajectory, a standard deviation was avoided, because the ERG is considered to be non-parametrically distributed (JACOBI *et al.*, 1993). Instead, the bounds enclosing all flickers have been plotted to provide at least a simple, rough estimation of the amount of scatter.

distributed (JACOBI *et al.*, 1993).

## 3.2 Flicker trajectory topology

As Figure 3.2 indicates, an embedding of a normal flicker ERG shows a typical basic structure, despite of inter-individual variations. This section will attempt to establish a reasonable nomenclature that is suitable to describe the morphology of normal and pathological ERG trajectories.

### 3.2.1 Descriptive approach

Figure 3.3 illustrates one way of describing a flicker trajectory topography. The roughly triangular trajectory evolves counterclockwise in time, starting at the

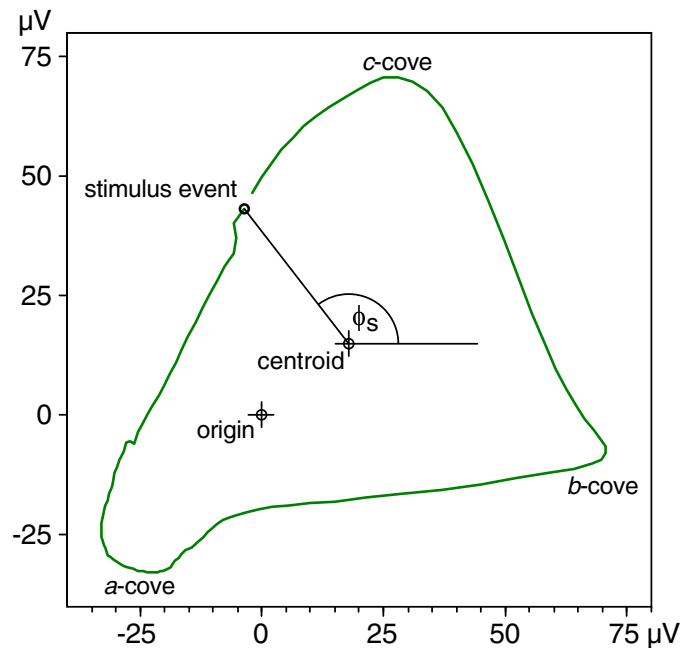


Figure 3.3: The topography of a normal mean flicker ERG.

stimulus event (circle). First, it moves quite straight down to the lower left quadrant to approach its smallest ordinate and abscissa values. Right there it forms the first arc turning right. There are three arcs, or *coves*, altogether. The term *cove* was selected to denote those arcs for their appearance, prefixed by letters *a*, *b*, *c*, that distinguish between different *coves*; the letters emphasize the analogy to the familiar ERG terminus *wave* and its own letter prefixes.

After the trajectory passed the *a-cove*, it moves to the right lower quadrant to increasing abscissa values. Here it bends to the *b-cove*, which is sharper than the *a-cove*, and here the abscissa maximum of the trajectory is located. Now the trajectory rises again, heading to the last arc, the *c-cove*, which surrounds the ordinate maximum of the trajectory; finally, the loop closes at the stimulus event.

Some aspects are helpful to keep in mind for the following discussion. First, the *coves* can be fairly well separated from the "straight" parts of the trajectory (although one can obviously argue about where a *cove* precisely starts and ends). Second, the stimulus is located roughly at one-third between the *c-cove* and the *a-cove*.

Figure 3.2 reveals that the shape of the normal's individual flickers seems to vary mostly around the *c-cove* area, followed by the *b-cove* region. This is in

contrast to the *a*-cove, whose shape appears inter-individually quite robust. The *a*-cove's nozzle-like, down-bended shape shows up in all trajectories.

Right after the stimulus event, and before the *a*-cove, some trajectories suffer from a small dip that continues almost to the start of the *a*-cove. The source for this is presumably an artificial one, because of its instantaneous occurrence after the flash; electrical crosstalk from the flash supply to the electrodes can cause such effects, for instance.

This so far pure phenomenologic description does not yet satisfy, because it lacks attributes that allow to compare different trajectories; what we need is an at least rough quantitative access to the trajectory structure.

### 3.2.2 Polar coordinates help quantifying

Let us glance back to Figure 3.2 again and consider the distribution of the stimulus events. The events of all normals' individual flickers do display some scatter, but they appear fairly concentrated in the upper right quadrant, rather than uniformly distributed along the trajectories. On the other hand, pathologic ERGs may deviate significantly from this finding, as we will see soon. To capture the position of the stimulus event on the trajectory, an angle measure, the *stimulus angle*  $\phi_s$ , seemed convenient. We define  $\phi_s$  to be the angle contained by a horizontal line and a line to the stimulus event. In fact, we can uniquely locate any trajectory point by its specific  $\phi$ , as long as the trajectory is not too complicated. The set of angles that describe a given flicker trajectory shall here be termed *topographic angles*.


Using angles as a trajectory's defining parameters offers another advantage: Angles are invariant to amplitude scalings of the underlying time series. This means that one gains complete independence of the ERG amplitude value as diagnostic measure, which is exactly what we demanded in Section 3.1. In other words, if topographic angles turn out to be of good selectivity with regard to retinal diseases, we can be sure that the diagnostic information is in fact a new one and adds to our understanding, rather than just being redundant to the results of amplitude analysis.

By introducing an angular measure we have already gone halfway from the *cartesian* representation to *polar coordinates*. The definition of an angle is, how-

ever, not yet sufficient; to unambiguously define the locus of a point in a two-dimensional plot, a *pole* and a distance  $r$  from the pole to the point are additionally required. The choice of the pole is ours, and so the question arises which point to favor from the infinitely many ones possible. At first sight, it is tempting to pick the appealing *origin* (0,0) of the phase space plot as the vertex, because it is simply the crosspoint of the cartesian coordinate axes and is thus easy to locate. Moreover, the trajectory is guaranteed to be always at least near the origin; that is because of the ERG's alternating nature, having always positive and negative values. However, Figure 3.3 illustrates that the origin is quite off-centered with respect to the area enclosed by the trajectory. The origin as vertex would systematically bias the angular resolution to the trajectory parts that are nearer to the origin. In the case of the mean flicker, this is the region of the *a-cove*.

For this reason, a point somewhere "in the middle" of the trajectory area would obviously suit better as the pole. One can of course argue about where the "middle" of a flicker trajectory actually is; somewhat arbitrarily, I decided on assigning the pole to the *centroid* of the trajectory. The centroid equals the center of mass of the area the trajectory encloses, if the area were cut out from a thin homogeneous material, for example, a piece of cardboard. For a comparison, refer to Figure 3.3, where both the origin and the centroid are depicted. The centroid seems to match the intuitive "middle" better than the origin does. Given the centroid as the new pole, the stimulus angle of the mean flicker was determined to  $\phi_s = 127^\circ$ .

At least in the mean flicker trajectory, the structures termed *coves* show up strikingly. Unfortunately, this is all but true for other flicker ERGs; the coves may lose shape and become quite ambiguous in the pathological examples, and even in some of the normals' trajectories. In these cases, polar coordinates do another good. They help to define the location of the coves in a more concise, objective way by means of  $r$  and  $\phi$ : *A cove is where  $r$  has a local maximum with respect to  $\phi$ .* Figure 3.4 displays this by the example of the mean flicker trajectory. Note how the maxima of  $r$  on the right panel correspond to the cove positions on the left panel.



**Figure 3.4:** Mean flicker trajectory in polar coordinates. The lefthand panel illustrates how the trajectory is transformed from a cartesian representation  $(x, y)$  into polar coordinates  $(r, \phi)$ . The same information  $(r, \phi)$  is given on the right-hand panel, but in order to facilitate quantitative statements about the local maxima of  $r(\phi)$ , a convenient cartesian representation of  $r(\phi)$  was chosen.

### 3.3 Pathological flicker ERGs

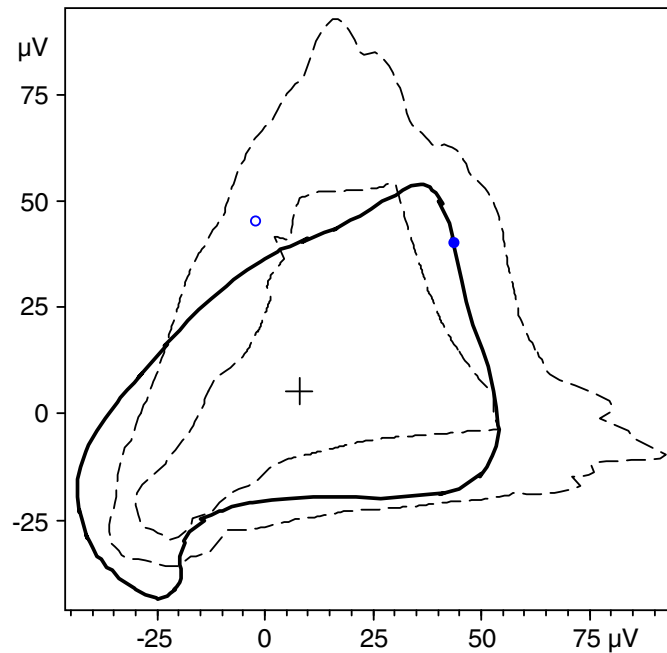
In the following sub-sections, the embedding of the normal flicker ERG will be compared to examples of different, clinically verified pathologies. Note that these examples are single cases, and therefore they are neither capable of, nor intended to, provide significance of whatever kind.

Each diagram shows the actual trajectory in bold printing. Additionally, the bounds from Figure 3.2 that enclose the normal ERGs are given as dashed lines. Before being embedded, the pathological flickers were amplitude-scaled to match the mean RMS amplitude  $\bar{a}_{\text{RMS}} = 32.4 \mu\text{V}$  of the mean flicker, for reasons explained above. The cross at the middle of each plot marks the trajectory's centroid.

#### 3.3.1 Usher's syndrome

USHER'S syndrome is a combination of hearing impairment and retinitis pigmentosa. Being first described by A. VON GRAEFE in 1858, it was named for CH. USHER, who believed that this condition was inherited.

The flicker depicted in Figure 3.5 is from a 46-year old female patient suffering




**Figure 3.5:** Flicker ERG example of a patient suffering from USHER's syndrome (bold), and the bounds of the normal flicker ERG distribution (dashed) from Figure 3.2. The filled circle denotes the very first point of the USHER trajectory, that is, the vector associated with the stimulus event. As time passes on, the trajectory evolves counterclockwise. The empty circle marks the location of the stimulus event of the mean flicker trajectory (which is not plotted here). The trajectory's centroid is marked by a cross.

from USHER's syndrome whose subtype was identified as type II, according to Seeliger *et al.* (2001). Her flicker trajectory shows three coves at obviously roughly the same angular positions as in the mean flicker. The topographic angles were found to  $\phi_a = 223^\circ$ ,  $\phi_b = 344^\circ$ , and  $\phi_c = 58^\circ$ . The position of the stimulus event, however, changed to considerable smaller angles ( $\phi_s = 45^\circ$ ); it moved even beyond the *c*-cove:  $\phi_s < \phi_c$ . Note that the trajectory leaves the normal's range outwards at the *a*-cove, but inwards at the *c*-cove.

### 3.3.2 Retinitis pigmentosa

Figure 3.6 displays the trajectory from a 35-year old female RP patient. Her disease was first diagnosed in 1995, and has become slightly worse since then. When examined in 2002, her visual acuity was 0.8 on both eyes, and the visual field was concentrically restricted to 5–10° on both eyes. The fundus showed atrophies of the retinal pigment epithelium and the choroidea. The




**Figure 3.6:** Flicker ERG example of a patient with retinitis pigmentosa (bold). Refer to Figure 3.5 for an explanation of the symbols.

dark adaptation threshold after 30 minutes was increased by three logarithmic units. The ISCEVstandard *ganzfeld* cone ERGs were also significantly affected (scotopic: *b*-wave amplitude  $50 \mu V$ , *b*-wave latency 60–70 ms on both eyes; photopic: *b*-wave amplitude  $< 10 \mu V$ , *b*-wave latency  $> 35$  ms on both eyes). The photopic flicker ERG amplitude was  $5.7 \mu V$  RMS on the examined eye. Results from molecular biology regarding the genotype were unfortunately not available.

The trajectory obviously suffers from a relatively high noise content, due to the highly reduced ERG amplitude typically found in this disease. Compared to the normal flicker, the coves seem much sharper, especially the *c*-cove, which forms an edge rather than a radius. If only the coves were smoother, the trajectory would resemble remarkably the one of USHER's syndrome from the previous example; this is even true for the areas where the trajectory exceeds the normal range.

The topographic angles of the coves were found to fall within the distribution of the normals:  $\phi_a = 225^\circ$ ,  $\phi_b = 342^\circ$ , and  $\phi_c = 55^\circ$ . However, considering the stimulus angle  $\phi_s$ , it moved considerably. It not only passed the *c*-cove, but went even further crossing the horizontal  $0^\circ$  line, and it is now located way






**Figure 3.7:** Flicker ERG example of juvenile macular degeneration (bold). Figure 3.5 explains the symbols.

within the fourth quadrant at  $\phi_s = 322^\circ$ . This is almost opposite to the location of a normal stimulus.

### 3.3.3 Juvenile macular degeneration

The term “juvenile macular degeneration” refers actually to a number of different forms of retinal dystrophy. As their most important symptom, they share a reduction in central vision while preserving peripheral vision. The name indicates that these diseases affect mainly the macula as a comparably small retinal area. One can expect a *ganzfeld* ERG to look fairly unsuspecting, since it examines the retina as a whole. Consequently, the trajectory in Figure 3.7, which was recorded from a 41-year old male patient with juvenile macular degeneration, seems within the normal range; the shift of the trigger event towards the *c*-cove ( $\phi_s = 84^\circ$ ) is the only hint for a pathology here.

The patient reported a continuous decrease of vision since ten years; juvenile macular degeneration was diagnosed in 1995. At the time of the examination, he needs a magnifying glass when reading. The clinical examination showed a visual acuity of 0.2 on both eyes. In Conventional perimetry, there was an



**Figure 3.8:** Flicker ERG example of STARGARDT's disease (bold). Refer to Figure 3.5 for an explanation of the symbols.

absolute central scotoma in the visual fields of both eyes; the peripheral visual field is only slightly affected. In both maculae, changes of the retinal pigment epithelium were found. The patient has a normal dark adaptation threshold after 30 minutes. The scotopic *ganzfeld* ERG shows no anomalies, while the amplitudes of the photopic cone ERG are decreased. The photopic flicker ERG amplitudes of the examined eye were  $16.3 \mu\text{V}$  RMS.

### 3.3.4 Stargardt's disease

STARGARDT's disease, also known as *fundus flavimaculatus*, is a subtype of the juvenile macular degenerations, or better, macular dystrophies. It is usually diagnosed in individuals under the age of 20 when they first become aware of decreased central vision.

The example in Figure 3.8 was recorded from one of those patients, a 18-year old woman, whose STARGARDT's disease was diagnosed at the age of 13. Currently she is suffering from a visual acuity of 0.1 of both eyes, caused by absolute central scotoma detectable in the visual field of both eyes. Fundoscopy was able to confirm a fundus flavimaculatus. Interestingly, her ISCEV ERG exami-

nations were normal, which is in some contrast to the findings of the following trajectory analysis. A multifocal ERG examination yielded further evidence for STARGARDT's disease.

The overall trajectory shape is similar to the normal mean; even the convex shape between  $b$ -cove and  $c$ -cove was occasionally found in normals. Each of the cove angles falls within the corresponding normal range. Thus, it is once again the stimulus angle that discloses the disease: As in all pathological examples,  $\phi_s = 53^\circ$  has passed clockwise, and is here located to the right of the  $c$ -cove.

### 3.4 Nonlinear Noise Reduction performance

The performance of nonlinear noise reduction (NNR) methods has been extensively demonstrated on artificial data; a classic example of such data is the HÉNON map (HÉNON, 1976), which yields reproducible, often irregular, and sometimes even (in the physical sense) chaotic data. The advantage over real-world data is that, in the artificial case, the properties of the data are precisely known, and so the effects of NNR can be just as precisely described. However, this thesis aims to evaluate the NNR performance on real-world data, and this will be done using steady-state flicker responses, under different conditions.

The basic idea is to contaminate a well-known data set by artificial noise of different properties, and to see to what extent NNR is capable of removing the noise again. It was decided to perform the evaluation based on one single ERG response model data set. The reason for this was to keep the amount of data within a manageable range, and to provide comparability between the results. Now the question arises: Which one of the standard ISCEV ERGs serves best for such an evaluation?

Even when processing "natural" data we have to rely on some type of model. For reasons explained in Appendix A, it is highly desirable to have an artificial response (i. e., a data set) at hand that both equals the expected "real" response, and is precisely known *a priori*. A look at the ISCEV standard reveals that the flicker ERG is set apart from the other ones by its high number of responses; a twenty-second time series of a 30 Hz flicker contains 600 responses, compared to 3...10 responses of the other standard ERGs. As explained in Section 2.7.3,

averaging yields the most reliable results when the number of ERG responses is large. The flicker ERG was thus selected as model, because then the difference between the model and the true value can be expected to be minimal.<sup>1</sup>

### 3.4.1 Generating the test data sets

#### Source data

The source data set  $s$  consists of five identical successive flicker responses of a normal subject. First, the entire time series of a normal subject responding to 600 flicker events ( $f_{\text{flicker}} = 30$  Hz) was high-pass filtered at 10 Hz and then averaged to one single resulting response. This averaged response was duplicated four times; the duplicates were appended to the original averaged response, in order to obtain five identical successive responses. To ensure smooth transitions within the resulting data set  $s$ , the last four data values of each response were slightly adjusted. The RMS amplitude  $a_{\text{src}}$  of  $s$  was determined to  $a_{\text{src}} = 62.3 \mu\text{V}$ .

#### Noise data

The generation of the noise sets  $n_i$  was straightforward once  $a_{\text{src}}$  was known. A Gaussian random number generator was used to create the noise data sets.<sup>2</sup> Since the aim was to investigate NNR performance for different initial signal-to-noise ratios  $\text{SNR}_{\text{in}}$ , each noise set was amplitude-scaled to match the desired  $\text{SNR}_{\text{in}}$ . In order to provide a measure of the variation of the NNR performance,  $N = 10$  different noise sets  $n_i$ , where  $i = 1, \dots, N$ , were generated for each  $\text{SNR}_{\text{in}}$ . The resulting NNR performance was assumed to be normally distributed, and the standard deviation  $\sigma$  was calculated for each  $\text{SNR}_{\text{in}}$ ; this  $\sigma$  is represented by the error bars in the diagrams below. To minimize systematical errors, the 10 noise data sets were generated anew for each  $\text{SNR}_{\text{in}}$ .

<sup>1</sup>Besides, the well-behaviour of the flicker ERG with respect to averaging means that the flicker examination is best suited for the detection of very weak ERG amplitudes. If a flicker ERG does not show a significant response, other ERGs most probably won't, either.

<sup>2</sup>Actually, in order to mimick different properties of real biological noise, two different models were used to generate the noise: white Gaussian noise and a type of spectrally matched noise; these models will be introduced below and will be discussed in Appendix A. To avoid confusion, let us assume for the moment that "noise is noise", and that its properties don't matter as long as it is random.

### Contaminated data

The final contaminated data sets  $q_i$  were obtained simply by adding the noise sets  $n_i$  to the source  $s$ , that is,  $q_i = s + n_i$ , where  $i = 1, \dots, 10$ . This means that there were  $N = 10$  contaminated data sets  $q_i$  available for each  $\text{SNR}_{\text{in}}$ .

### 3.4.2 Different kinds of noise

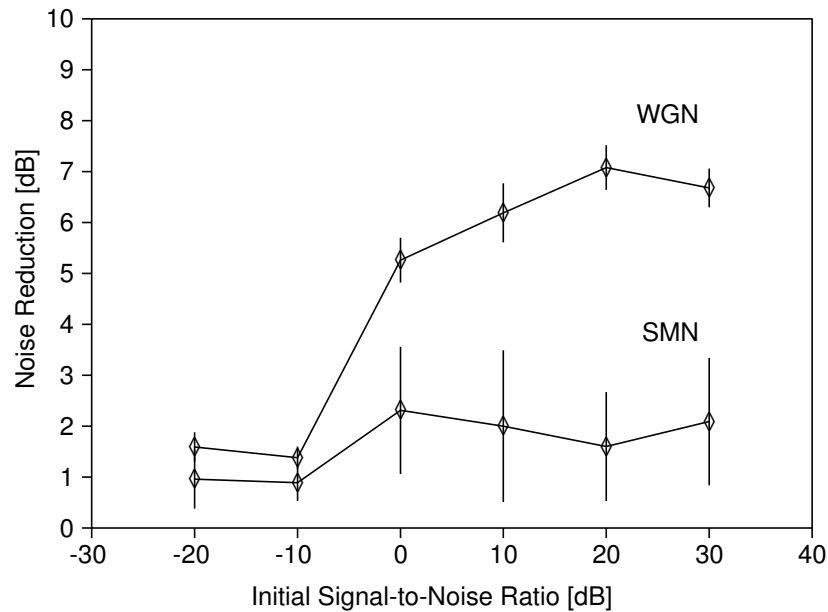
When it comes to generate artificial noise and to select a noise model, it is white Gaussian noise (WGN) that is used most commonly, because it is both easy to compute and its properties are well understood. However, WGN does not necessarily mimic real physiological noise very well; I therefore decided to do an additional analysis by applying another noise model which matches the signal not only in terms of RMS amplitude (as WGN does), but also in its spectral distribution. The idea is to provide a noise that shares some more properties with the original signal, because noise in a physiological context might be expected very similar to the signal of interest. Appropriate noise models have been described by, e. g., THEILER *et al.* (1992). A variation was used here, in the following referred to as *spectrally matched noise* (SMN). Refer to Appendix A for details on these models.

As can be expected, a spectrally matched noise is more difficult to remove from the source than plain white noise is. Figure 3.9 shows a comparison of the results.

At the left part of the diagram, at negative  $\text{SNR}_{\text{in}}$ , the noise power exceeds the signal power by far. This is where the NNR's noise reduction  $R$  is barely noticeable, regardless of the noise type. Embeddings do not reveal any signal-related structure at such high noise levels, and any phase space technique is likely to fail under such circumstances.

Considering WGN, the performance rises significantly as  $\text{SNR}_{\text{in}}$  reaches 0 dB, i. e., when the noise's power roughly equals the signal's. The performance culminates at an  $\text{SNR}_{\text{in}}$  of 20 dB,  $R \simeq 7$  dB, and seems to drop again when  $\text{SNR}_{\text{in}}$  is further increased.

The performance on SMN, however, is centered around some 2 dB. Although this seems next to nothing, it indicates that the algorithm basically works also with such complex noise models like SMN, where other filter techniques, even




**Figure 3.9:** Nonlinear projective noise reduction performed on different kinds of artificial noise: white Gaussian noise (WGN) and spectrally matched noise (SMN). The ordinate is scaled in  $\text{SNR}_{\text{in}}$ , while the abscissa denotes the noise reduction performance  $R$ . The error bars indicate the standard deviation  $\sigma$ .

sophisticated ones like Fourier transform, inherently fail. Because this type of noise equals the signal in terms of frequencies, it is a difficult task for many filter algorithms.

### 3.4.3 Iterations

If we can apply NNR to raw data in order to get a noise-reduced output, why not reapply NNR to the output and reduce noise even more? In fact, NNR can be used iteratively, and in most cases it turned out good practice to do so.

Figure 3.10 illustrates how NNR performance depends on the number of iterations. Again, the results base upon 10 different noise sets, applied to a 30 Hz flicker response of a normal subject. Initial noise level was 0 dB white Gaussian noise (WGN). The solid line shows how the noise reduction (left ordinate) depends on the number of NNR iterations  $i$  (abscissa) performed. At the given initial noise level of 0 dB, the performance curve indicates an optimum result at about 4 iterations; increasing  $i$  causes the performance to drop again. The relative error (dashed) becomes steadily worse with increasing  $i$ , but in a rather



**Figure 3.10:** NNR iterated on flicker data. The abscissa gives the number  $i$  of performed NNR iterations. The arithmetic mean  $\mu(i)$  of the noise reduction  $R$  (diamond symbols), obtained as the result of the last iteration, is shown at the right ordinate. The error bars denote the standard deviation  $\sigma(i)$ . The crosses represent the relative error  $\sigma(i)/\mu(i)$ , scaled at the right ordinate.


nonlinear way; for  $i \geq 5$ , the relative error climbs rapidly.

### 3.4.4 NNR on sinusoids

Sinusoids are the most promising data to be analyzed by NNR, because they are known to form a perfect circle in phase space. The probability of false neighbors is minimal, and any deviating structure is easily detectable and is due to noise only. Although one wouldn't judge a flicker ERG to look completely different from a sinusoid, NNR performs apparently much better on sinusoids than it does on flicker data (RILK and SHARPE, 2000). One would expect the NNR results to improve, if only the flicker data could be rendered to look more sinusoidal.

This is, in fact, a quite easy task. We just have to add a pure sinusoid (with  $f_{\text{sin}} = f_{\text{flicker}}$ ) to our flicker data prior to NNRRing, and remove it afterwards.

The sinusoid's two degrees of freedom, its RMS amplitude  $a_{\text{sin}}$  and its phase  $\phi_{\text{sin}}$ , determine how much the addition influences the resulting time series. The higher is  $a_{\text{sin}}$ , the more the artificial sinusoid dominates the time series. There is presumably a reasonable upper limit for  $a_{\text{sin}}$ , with respect to the flicker's RMS



**Figure 3.11:** Results improve remarkably when the sinusoid trick is used. The same data were used as in Figure 3.10, but now a sinusoid with the same RMS amplitude as the noise was added to the signal before the noise reduction process, and was subtracted afterwards. Again, the abscissa gives the number of NNR iterations  $i$  performed, and the arithmetic mean  $\mu(i)$  of the noise reduction  $R$  (diamond symbols) as the result of the last iteration is shown at the right ordinate. The error bars denote the standard deviation  $\sigma(i)$ . The crosses represent the relative error  $\sigma(i)/\mu(i)$ , scaled at the right ordinate.

amplitude. An  $a_{\text{sin}}$  too high dwarfs the flicker features, and probably causes NNR to eliminate them altogether with the noise, leaving the plain sinusoid. The phase  $\phi_{\text{sin}}$  should obviously match the phase of the flicker response. A proper phase may be found, for example, by inspecting the phase of the flicker response's zero-order harmonic in the frequency domain.

In order to just demonstrate the basic idea, the same data as in Section 3.4.3 was taken, that is, a flicker contaminated by 10 different WGN data sets at  $\text{SNR}_{\text{in}} = 0$  dB. A sinusoid was added prior to NNR; the sinusoid's amplitude was set to  $a_{\text{sin}} = a_{\text{src}} = 62.3 \mu\text{V}$ , and its phase to  $\phi_{\text{sin}} = 180^\circ$ , estimated by a visual inspection. The values were selected somewhat arbitrarily, and a comprehensive analysis might yield better ones. The same NNR processing scheme as in Section 3.4.3 was then applied. After NNR was completed, the sinusoid was subtracted from the outcoming data.

Figure 3.11 shows the results. The performance of NNR improved remarkably, compared to Figure 3.10, in terms of both noise reduction and relative error as a



---

measure of predictability. Noise reduction rises steeply, tops 8 dB at 6 iterations, and decreases only marginally for higher numbers of iterations. The relative error stays fairly constant over the entire range, nowhere exceeding 0.12. One might get even better results, considered that the ones given here were obtained with non-optimized  $a_{\text{sin}}$  and  $\phi_{\text{sin}}$ .

# Chapter 4

## Discussion

**W**hile it could be proved in theory that a delay embedding reveals more information than does the plot of the underlying time series, the processing of ERG data in such a manner has not been addressed in the literature so far. Thus the question now to ask is: *Cui bono?* What is the use of it?

This chapter will first discuss the properties of the previously given embedded flicker ERG examples. Then, the results of the NNR technique will be related to other signal recovery techniques, in order to point out how the different methods compare. We will also consider some pitfalls that can occur when applying embedding techniques, and how they are best avoided.

### 4.1 Topographic angle synopsis

With the outline of a trajectory topology in Section 3.2, we aimed for a tool that is able to distinguish between normal and pathological flicker trajectories. In particular, the topographic angles were introduced with the idea of an empirical quantitative measure of trajectory structures. We will now discuss whether this topology meets the expectations, and, if it does, to what extent. For this purpose, Table 4.1 provides a synopsis of the topographic angles that were found in the ERG trajectories from Chapter 3.

The angular positions  $\phi_a$  of the  $a$ -cove, given in the first data column, turned out surprisingly robust in all of the examined flicker ERGs, be they from nor-

Flicker trajectory topographic angles				
Flicker type	$\phi_a$	$\phi_b$	$\phi_c$	$\phi_s$
Normals <i>median</i>	224°	336°	84°	123°
<i>lowest...highest</i>	219°...229°	330°...345°	48°...102°	103°...164°
USHER's syndrome	223°	344°	58°	45°
RP	225°	342°	55°	322°
Juv'Mac'Degeneration	222°	336°	79°	84°
STARGARDT's disease	224°	331°	68°	53°

**Table 4.1:** Synopsis of the topographic angles of the flicker trajectories, as found according to the polar coordinate method from Section 3.2.2.

mals or from patients. The examined normals cover a range of only 10°, from 219° to 229°, with a median of 224°. The results for the pathological ERG examples were almost identical; note how precise they match the range of the normals. The same picture, except for slightly wider ranges, was found for the other topographic angles  $\phi_b$  and  $\phi_c$ . Thus it seems that the cove positions are not suitable to distinguish between normal and ill retinal conditions. This is disappointing, despite the small data base which only allows preliminary statements. The cove positions are obviously insensitive to pathological changes in the retinal function.

The situation seems quite different for the topographic angle  $\phi_s$ , that is the angular locus of the stimulus event, presented in the last column of Table 4.1. The normals' range is 103°...164°, and their median 123°. Although the range covers 61° and is thus wider than any of the other ones, this time none of the pathological  $\phi_s$  falls within. Assuming that the rather small sample of 10 normals provides trustworthy results, each one of the pathological examples discloses itself by its odd  $\phi_s$ . The Figures 3.5–3.8 illustrate how the embeddings reflect this finding: All pathological stimulus events more or less shifted on their trajectories in clockwise direction. Most of them even moved beyond their trajectory's *c*-cove, a phenomenon that was never observed in any of the normal flickers.

In spite of these encouraging results the question still remains whether  $\phi_s$  is selective enough to discriminate between different diseases, or not. Having only single cases available, one naturally cannot give a satisfying answer by now.

Definitive statements can only be provided on the base of more experimental data, which may be gathered by a follow-up clinical study that includes more participants with different diseases in different stages (which in turn requires elaborate criteria for the assignment of patients to different disease stages).

## 4.2 NNR and other noise reduction techniques

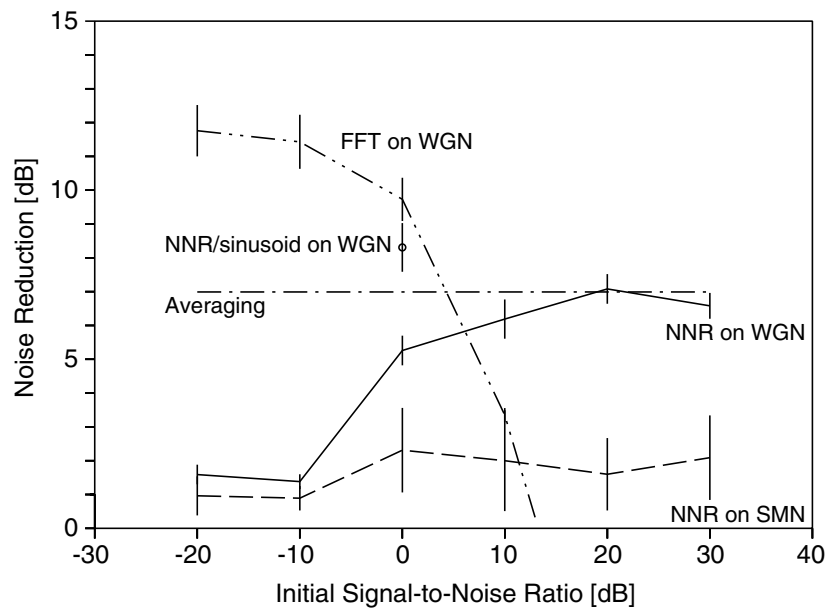
Classical noise reduction techniques are concerned with the removal of low amplitude noise from a time series. There are a lot of different filtering schemes; among them the Fourier power spectrum based routines are a typical method. The noise spectrum may be assumed, for example, to be white, equally distributed over the entire frequency range. A bandpass filter centered at the region of interest might then be applied to produce a cleaned version of the signal. A comparison between such “classic” methods and NNR is of course mandatory and will be provided below.

### 4.2.1 NNR and averaging

As mentioned above, averaging is fairly well-behaved in terms of predicting its noise reduction performance  $R_{\text{ave}}$ ; on stationary data,  $R_{\text{ave}}$  depends only on the number of averaged periods, and can even be written as a simple equation:  $R_{\text{ave}} = 20 \log_{10} \sqrt{N}$  dB, as explained in Section 2.7.3. Considering the five-response flicker data set from Section 3.4.1, we can calculate the noise reduction to  $R_{\text{ave}} = 20 \log_{10} \sqrt{5}$  dB = 6.99 dB. Note that it does not depend on the initial  $\text{SNR}_{\text{in}}$ —this is why the noise reduction performance of coherent averaging is a straight horizontal line. Moreover, it is even independent from the noise’s spectral content, as long as the noise is uncorrelated. If this condition is met, all other noise properties do not matter. All these benefits make averaging a pretty good and stable signal recovery tool.

Figure 4.1 provides a synopsis how the different noise reduction methods perform with respect to  $\text{SNR}_{\text{in}}$ . It is obvious why averaging is everybody’s darling: It pleases with a satisfying, constant, predictable performance over the entire range.

The NNR performance  $R_{\text{NNR}}$ , on the other hand, can seemingly not compete



**Figure 4.1:** Performance of different noise reduction methods on the flicker data from Section 3.4.1. Shown are the curves of NNR applied to white Gaussian noise (solid line), NNR applied to spectrally matched noise (dashed), FFT applied to white Gaussian noise (dash-dot-dot), and averaging (dash-dot). The dot is the best result of the NNR/sinusoid combination on WGN. The test data sets were the same as in Figure 3.9; again,  $\text{SNR}_{\text{in}}$  is given at the ordinate, while the abscissa denotes the noise reduction performance  $R$  for the different methods. Both axes are scaled in decibels. Refer to Section 3.4.1 for how the test data sets were generated. Error bars for averaging are deliberately not given, since, in this special case, the expected value is exactly known from theory.

with averaging. Only at  $\text{SNR}_{\text{in}} = 20$  dB its WGN performance curve scratches the one of averaging. It is even worse when applied to SMN. However, these are the outcomes without optimization; the result obtained by the sinusoid trick trial seems more promising. In Figure 4.1, it is represented by a point at 0 dB abscissa and 8.3 dB ordinate. This is clearly above averaging and improves the results of plain NNR on WGN by more than 3 dB.

The sinusoid trick can be expected to be particularly useful at a low  $\text{SNR}_{\text{in}}$ , whereas it presumably fails at a  $\text{SNR}_{\text{in}}$  of 10 dB and higher. If a NNR/sinusoid combination is applied rigorously, it might smooth the NNR performance curve and boost it even above the range of the performance of averaging—given that the contamination is plain white Gaussian noise only.

### 4.2.2 NNR and frequency filters

The performance of filters that operate in the frequency domain was estimated by means of the Fast Fourier Transform (FFT) method as a representative. For this purpose, CALYPSO contains a FFT-based frequency filter tool with user-selectable passband and amplification settings. For the evaluation the passband limits were set to 25 Hz and 150 Hz, such that the lower limit was well below the flicker frequency of 30 Hz, and the upper limit excluded frequencies above the fourth harmonic of the fundamental frequency. The passband amplification was set to 60 dB, and the FFT size to 1024. No window function was applied. The performance was tested using the flicker data from Section 3.4.1 once more, noised by WGN of different amplitudes.

Again, refer to Figure 4.1 for the results. FFT is obviously hard to beat at a low  $\text{SNR}_{\text{in}}$ . It has by far the highest  $R$  of all noise reduction algorithms considered, with  $R_{\text{FFT}}$  approaching 12 dB at  $\text{SNR}_{\text{in}} = 20$  dB. What is more, its error bars are pleasingly small, which indicates a performance whose predictability is only superceded by that of averaging. However, as  $\text{SNR}_{\text{in}}$  increases above 0 dB, we encounter a unexpectedly steep decline in performance, which even drops below zero at about  $\text{SNR}_{\text{in}} \simeq 13.5$  dB. This finding was a surprise. As a possible cause, a careless FFT parameter setting was initially suspected, but some checks showed the phenomena to be quite invariant against changes of the parameter values. One does not want to find negative performances at all, because this means that the output data is actually worse than was the input. However, an ERG with an SNR of more than 15 dB is already quite good; no examiner would worry about noise reduction if all ERGs were that clear.

The application of frequency filters is, by their nature, limited to noise that can be separated by its different spectral content. We can all but guarantee this restriction to be true in a physiological environment, and thus the findings for FFT on WGN in Figure 4.1 may constitute a sort of ideal case. For this reason, the flicker ERG is well suited for a FFT analysis, because its frequency spectrum is dominated by well-separated peaks at the flicker frequency and its harmonics. Other standard ERGs may show a more continuous spectrum that resembles white noise more than a flicker ERG does;  $R_{\text{FFT}}$  would drop in such cases.

A negative performance seems not a problem that NNR has, too. The NNR results presented here and in the previous chapter are based upon a total of

many hundred NNR runs on different noisy data sets. As long as the NNR parameters were kept within fairly reasonable limits, an outcome below 0 dB was only encountered at a handful of sets, even when the noise's power was thirty times higher than the signal's. This finding might indicate that NNR does not corrupt an already noisy signal even more, regardless of the noise level.

Due to the FFT's steep roll-off at higher  $\text{SNR}_{\text{in}}$ ,  $R_{\text{NNR}}$  crosses at  $\text{SNR}_{\text{in}} \approx 6$  dB. On the left of this point,  $R_{\text{FFT}}$  is way above  $R_{\text{NNR}}$ , while at the right  $R_{\text{NNR}}$  recovers.

### 4.3 Synergies


While, in some cases, we can observe a slight performance advantage of NNR over averaging, at least from Figure 3.11, an interesting question might be: is this finding also clinically relevant? Do a few meager dB of better performance justify the undoubtedly high demands of NNR in terms of complexity, computing time, and user skills required?

At first glance it seems not. We have just learned from Figure 4.1 that the performance of NNR nowhere really exceeds the performance of averaging. The sinusoid trick from Section 3.4.4 surely helps, but it introduces another two parameters that require values correctly set.

Fortunately, the three noise reducing methods considered here operate in quite different ways. They can therefore easily be combined to enhance each other's benefits, without having to worry much about redundancy effects, as illustrated in Figure 4.2. It is, for example, possible, and even advisable, to first average the approximately 600 flicker responses to 5 resulting responses, and further clean them by NNR, or by FFT, if desired. If the outcome still does not satisfy, another averaging step may be applied that yields one final response. All steps contribute independently to the total amount of noise reduction.

### 4.4 What can go wrong

Although no significant evidence was found that NNR renders the signal worse than it was before, the results cannot provide a guarantee that the method will always work, even in the limited circumstances that have been described here;



**Figure 4.2:** Noise reduction performance can be enhanced by combining the different methods. Ordinate: noise reduction in dB. Left bar: performance of NNR alone; middle bar: performance of FFT alone; right bar: performance of the combination of FFT and NNR. The flicker data from Section 3.4.1 were again taken as test set, noised by 0 dB WGN. The error bars denote the standard deviation, assumed that the performance is normally distributed.


NNR shares this drawback with most algorithmic methods based on nonlinear dynamics. We have to remark that one of the most serious annoyance of nonlinear theory is the obvious lack of reliable methods to determine the proper parameter settings. Considerable human judgement is still required to achieve optimal performance. However, at least to the extent that it has been tested, the method was in fact found to perform robustly, as long as some constraints are observed.

#### 4.4.1 Embedding parameters

##### Dimension

So far the actual value of the embedding dimension  $m$  has been just taken for granted, rather than having been strictly deduced. The embedding theorems guarantee a delay embedding to be valid, but only if  $m$  is chosen large enough. In other words, there must be a number of dimensions  $m_0$  for which a delay reconstruction is an embedding, while for  $m_0 - 1$  it is not. What happens when





**Figure 4.3:** The method of false neighbors applied to the flicker data of Section 3.4.1. The result suggests  $m_0 = 2$ , as described in the text.

passing from  $m_0$  to  $m_0 - 1$ ? One simply projects along one coordinate and thus maps different parts of the trajectory onto each other. In a macroscopic view, the trajectory intersects itself. Because the neighborhoods are mixed up, the rate of *false neighbors* increases at the crosspoints, indicating that determinism is violated.<sup>1</sup> The fraction of false neighbors is estimated by the *method of false neighbors*; when increasing  $m$ , starting at small values, we can thus detect the minimal embedding dimension  $m_0$  by finding no more false neighbors. Figure 4.3 presents the results of this analysis performed on the flicker data from Section 3.4.1. The fraction of false neighbors drops rapidly when passing from  $m = 1$  to  $m = 2$ , and stays nearly constant for larger  $m$ . This finding is consistent with the result of a visual inspection, that, for  $m = 2$ , the flicker trajectory unfolds without intersections. One would expect the fraction to drop to zero once  $m_0$  has been reached; however, it levels at some 5%, which is due to the inherent limitations of the method.

## Delay

Consider the case when the delay is zero,  $\tau = 0$ . Then the embedding window is also zero, and all elements of any vector are identical. All vectors accumulate

<sup>1</sup>Figuratively spoken, determinism in phase space means that there is always an unique direction from one point to the next. A crosspoint on the trajectory contradicts determinism because one is no longer sure about which way to go.

around the bisectrix of the embedding space, regardless of the structures of the underlying time series. Such a representation would obviously reveal only very little information.

As the delay increases from zero, the trajectory unfolds more and more. In fact, there exists a delay value for which the trajectory is maximally unfolded; any further increase of the delay beyond this optimal value causes the trajectory to collapse again and any informations to be harder to unveil. So far, so good, but how can we reliably determine the proper delay setting for a given data set? Since  $\tau$  has no relevance in the mathematical framework,<sup>2</sup> there exists no rigorous way of determining its optimal value, and it is even unclear what properties the optimal value should have. The *estimation* of  $\tau$ , however, has been a greatly emphasized issue in the literature, but the suggested methods yield optimal results for selected systems only, and perform just as average for others.

Fortunately, the more intuitive way that uses the discriminatory capabilities of the human eye yields satisfying results. The slight variation of the estimated delay due to the individual observer's bias does not matter much; the delays provided by different algorithmic methods vary, anyway.


Figure 4.4 demonstrates how the trajectory morphology changes depending on the delay. Fortunately, at least NNR is more robust to changes in  $\tau$  as one might suspect, but the overall trajectory shape is severely affected, as obvious in the diagram; examine, for example, the position of the stimulus event on the left panel to the one on the right panel. Trajectories from the same time series, but with different delays, are therefore difficult, if not impossible, to compare. In consequence, once a particular delay value is agreed on, it has to be kept constant in all further analysis, and for all data sets to be compared. The delay value chosen here for the ISCEV 30 Hz flicker ERG was  $\tau = 5.0$  ms.

### Neighborhood Radius

Obviously, the NNR outcome depends on the actual size chosen for the neighborhoods. Too small neighborhoods will contain too few points for good statistics, and all directions will be dominated by noise only. Too large neighbor-

---

<sup>2</sup>This is because the data are assumed to be of infinite precision; the delay can then be chosen arbitrarily without sacrificing accuracy. In real-world applications, however, the proper choice of  $\tau$  is quite important.




**Figure 4.4:** The appearance of a trajectory depends very sensitively on the delay. The left panel shows the trajectory from Figure 3.3, where the delay was set to 5 milliseconds (as it was throughout the analysis). The trajectory on the right panel is exactly the same data set, but now the delay was doubled to 10 milliseconds.

hoods cancel out local attractor features as the spatial resolution drops, and they increase computing time considerably. As it turns out, setting the Neighborhood Radius visually to roughly twice the noise level is a reasonable choice.

A number of at least five neighbors per dimension is desirable, a lesser number of neighbors likely deteriorates the neighborhood statistics and will therefore yield poor NNR results. This is the reason why the sampling rate has been set to 3600 Hz, because then we can guarantee a sufficiently large number of neighbors, even when the Neighborhood Radius is small.

Determining the proper size of a point's neighborhood depends on many factors. At the present state, at least one method has been proposed which may simplify the selection of the proper Neighborhood Radii by involving local noise analysis (KERN *et al.*, 2000), but it unfortunately fails at the very high noise levels the ERG is subject to. Thus, a more empirical and straightforward solution was established here that led to acceptable results. The proper neighborhood size apparently depends quite linearly on the RMS noise level. Therefore, CALYPSO includes some "distance" measuring tools: after clicking on the two endpoints of a distance to be measured in the phase space plots, the program returns the length of the distance in  $\mu V$ . These tools can, among other things, be used for a rough estimation of the noise level as it appears in the phase space



**Figure 4.5:** Measuring tools provide a convenient way to roughly estimate the noise level in the phase space plots. The bold line marks the distance to be measured; it is  $11.5 \mu\text{V}$  here.

representations; this is illustrated in Figure 4.5. It turned out that twice this size provides a reasonable estimation for the proper Neighborhood Radius.

#### 4.4.2 Artifacts

Artifacts in electrophysiological settings can originate from many, and very different, sources. Despite their various origins, they often share a common morphology: a sharp peak at very high amplitudes, compared to the expected ERG signal level. This class of artifacts constitutes vectors that usually point far away from the main trajectory, due to their high amplitudes. They therefore do not pose as large problems to NNR as they do to averaging, because their neighborhoods do usually not interfere with those of the main trajectory. This means that the artifacts will stay quite unaffected by NNR, and, after an inverse transformation into a time series, they will be still recognizable as such. Nevertheless, it is of course advisable to remove them prior to NNRing.

Artifacts that show up as a slow drift are a different matter. A slow changing offset will shift the trajectories with respect to each other; points that actually belong to the same neighborhood will not share it any longer. Instead, they will

---

literally drift apart, which introduces a systematical error to the PCA calculation. Fortunately, such artifacts can be removed in many cases by a high-pass frequency filter of a frequency just below the base frequency of the signal. If this is impossible for whatever reason, another remedy may work: The search for neighbors of a point may be restricted to the immediate past (or the immediate future) of the point, excluding neighbor candidates that are actually member of the neighborhood, but were constituted from measurements too long ago (or too far in the future). We can thus minimize the effects of drift artifacts by carefully selecting such a temporal cut-off.

# Chapter 5

## Conclusion

Phase space methods may offer remarkable capabilities of analyzing and processing complex data; however, their application is all but simple. It is still a long way to go to make nonlinear methods at least as half as foolproof as, e. g., averaging, and it is even likely that such a robustness will never be accomplished. After all, their behaviour is much less predictable, just as the behaviour of the nonlinear dynamics they are intended to examine. Because few experiences has been available with real-world data by now, and none at all with ERGs in particular, the approach taken here has been governed by a feasibility point of view.

We have seen how to construct a delay embedding, that is, a multi-dimensional representation of a system of which information is only available as a scalar time series. Although a delay embedding is not exactly a state space in the physical sense, it was proved elsewhere that both are equal with respect to the completeness of the represented information. From a more general point of view a delay embedding and a state space representation are just two members of the set of all possible embeddings. It turned out that *ganzfeld* flicker ERGs, be they from normals or from patients, are well suited for an embedding, because they require only two embedding dimensions, which eases a visual inspection considerably. Moreover, all the examined normals' ERGs share the same basic topography, that is, they look similar enough to encourage the definition of a standard. By means of such a standard, one may undertake a comparative analysis of different pathological ERGs. Polar coordinates help to obtain a quantitative description of the trajectory topography, because the flicker trajectories

are closed loops, basically. This has motivated the introduction of some angular measure that is independent from amplitude considerations; for instance, the *stimulus angle* that defines the position of the stimulus event. The findings indicate that the stimulus angle is very sensitive to pathological changes, while maintaining a fairly constant value under normal conditions.

Embedding techniques establish access to advanced data processing methods, of which the Nonlinear Projective Noise Reduction (NNR) is an example. NNR considers the state vectors as points that actually populate a low-dimensional subspace, but had been scattered off their hosting subspace by a noising process into all dimensions. One can estimate the original local subspace by analyzing the local neighborhood of each point, and, by projecting each point on its estimated subspace, the noise-free trajectory is reconstructed. The performance of NNR has been evaluated with a normal averaged flicker that was contaminated by artificially generated noise. The results could not be stated as briefly as the performance of averaging, because they depend on several circumstances, for example, the initial noise level, the type of noise, and the number of iterations. The essence from the Figures 3.9–4.1 is that it seems reasonable to expect a noise reduction of about 5...8 dB in a data set of 5 responses.

A comparison of NNR to other noise reduction methods has shown each to have its pitfalls and pearls. Averaging acts equally and predictably with a medium performance, frequency filters are very good at low signal-to-noise ratios (SNRs) while failing completely at high SNRs, and NNR is a low to medium performer over the SNR range, but it is not restricted to white noise and even cleans spectrally matched noise. Thanks to their different functional principles, all these techniques may be combined to improve overall performance.

## 5.1 Future outlook

### 5.1.1 Follow-ups

Evidence has been found that the topographic angle  $\phi_s$  may serve to discriminate pathological from healthy conditions. The question whether this finding is also true for more elaborate problems like separating between pathologies, or the discrimination between different stages of a disease, cannot be answered by

the single cases that have been considered here. Thus, additional studies are required to shed more light on the practical value of a flicker topology in general, and of topographic angles in particular. The questions to address further may be, for example:

- What happens to a flicker trajectory when the stimulus intensity is varied?
- How do other ISCEV ERG embeddings look like? Are topographic angles still feasible?
- Is a trajectory topology selective enough to distinguish between different stages of a disease?

### 5.1.2 Extending the polar coordinate concept


It has been discovered in Section 3.2.2 that a polar coordinate representation of a flicker trajectory eases the task of assigning and quantifying topographic angles. Moreover, the polar coordinate idea might bear some more potential than has yet been discussed. When one wants to investigate other ISCEV-type ERGs by nonlinear analysis, one most likely encounters cases that demand more embedding dimensions than the flicker ERGs presented here. Figure 5.1 displays an ISCEV ERG trajectory of such kind. Now, can a multi-dimensional embedding still be described in polar coordinates? In other words: Are polar coordinates scalable with respect to the number of dimensions?

Fortunately, this is not really a problem. In the case of three dimensions, the so constructed system is known as *spherical coordinates*, which are defined by a distance  $r$  and two angles  $\theta, \phi$ . In fact, by adding the appropriate number of independent angle measures, one compensates for any arbitrary number of additional dimensions. Such angular coordinates are by their nature just as versatile as cartesian ones, although less common.

### 5.1.3 Visualizing multi-dimensional spaces

Fortunately, all the flicker ERGs discussed so far reveal a structure that is easily embedded in a two-dimensional space, which relieves us of the need to juggle with three, or even more, dimensions. This is by no means a matter of course






**Figure 5.1:** An oscillatory potential ERG of a normal subject embedded in a two-dimensional space. Note that the trajectory intersects itself, indicating that the embedding is not appropriate.

in the case of other ERG types, as Figure 5.1 illustrates. The data embedded here is the oscillatory potential ERG of a normal subject, measured according to the ISCEV standard. Without going into any details let us just note that the trajectory intersects itself, which indicates that the embedding dimension  $m = 2$  was chosen too small, as is pointed out in Section 4.4.1.

When we want to visualize embeddings, we run into the fundamental problem that a piece of paper is restricted to only two dimensions. Any higher-dimensional features can only be shown in a projection, that is, a two-dimensional slice. Of course, we may add some perspective aids like the cube in Figure 2.2 to make the brain believe that it sees a 3D image, but this is not really satisfying, and anyway it does not support a quantitative access to the information the plot contains. In order to nevertheless provide at least a three-dimensional information on a two-dimensional medium, a false color scheme was recently implemented that allows a color encoding of the third dimension. Figure 5.2 gives an example.



**Figure 5.2:** A flicker ERG in a two-dimensional embedding. Its third dimension is encoded in false colors. The color bar at the ordinate assigns the colors to numeric values.

### 5.1.4 The multifocal ERG

The multifocal ERG (mfERG) is an exciting new method which was introduced some years ago by SUTTER and TRAN (1992), and which has been already widely approved. While the conventional ERG can take us far in diagnosing global retinal diseases, it lacks one essential thing: spatial resolution. A *ganzfeld* ERG integrates the response from the entire retina. While this is very useful for examining panretinal degenerations, focal retinal lesions will go undetected—think of macular degenerations, where a flash ERG may turn out completely normal, even when the central vision is severely degraded. In such cases, the mfERG is the method of choice. The stimulus is a hexagonal flicker pattern on a computer monitor, providing for spatial resolution, and, by a clever mathematical scheme, the response associated with the small retinal area illuminated by each hexagon is extracted from the retina’s integral response.

An interesting question to ask may be whether NNR is applicable to the mfERG. This is not trivial to answer, because the time series of a mfERG does not disclose its spatial information all too easily. It requires an advanced mathematical method, termed *m-transform*, to extract the responses elicited by each hexagon out of the raw time series that the amplifier delivers. Although one ends up basically with ERG curves that are comparable to a *ganzfeld* examination, the data

post-processing scheme is entirely different. One cannot just grab the output of the amplifier as in a full-field ERG. Nevertheless, NNR and phase space techniques are best applied on steady-state (i. e., stationary) responses. As a mfERG does produce such data, it may be a challenging task to investigate whether a NNR analysis of an mfERG is feasible.

# Appendix A

## Noise in different flavors

**A** main task of this thesis is to evaluate and to compare the performance of different noise reduction schemes. This requires a pretty precise knowledge of what the term *noise* actually means in contrast to the term *signal*. A simple definition of noise is "everything of a measurement that is not the signal itself". Unfortunately, this will not get us far, since we need to know the plain signal beforehand, which is however exactly what we would like to uncover by the highly non-trivial noise reduction process. How can we be sure that the algorithms do not affect the pure signal when applied to noisy experimental data? The short answer is: We never can, because we don't know the signal precisely. However, this is quite a common problem in real signal analysis, and so a number of approaches exist that ease the situation.

Let us recall how a signal analysis process works: First, the measurement data is acquired, second, an algorithm is applied, which yields, third, the signal and the noise (hopefully) well separated. For performance evaluation of the algorithm, we initially will go the reverse direction: First, we generate a *model* of the clean signal, and, second, some noise of different kinds. Third, signal and noise are added together; the outcome will serve as artificial measurement data with properties precisely known. We thus are able to compare algorithm performance, because now we know exactly how the output should look like.

Obviously we encounter a problem with this method: There are different suitable noise models, and the results depend on the model actually considered. This appendix will address the question which kind of noise is best and explain the noise models that have been used here.

Noise is an umbrella term for a wide variety of phenomena, usually connotated with a process that corrupts an otherwise clear structure. Noise in a pure electrical environment can be often modeled by a few straightforward physical equations. Obviously, such "simple" noise plays a role in the electronic equipment involved by the measurement process, but the main fraction arises from electrophysiology, and here's where the trouble comes in. Electrophysiological voltage noise is very diverse, while a mathematical description is required in order to model such noise by means of a computer. Fortunately, such noise models do exist.

Somewhat arbitrarily, noise models can be divided into two subclasses, depending on to what amount the original signal is involved in the noise generation process. Here, they will be termed *signal-independent noise models*, and *signal-dependent noise models*, respectively.

## A.1 Signal-independent noise models

Signal-independent noise models rely on a sequence of random numbers, most often generated by a software pseudo-random generator. Once generated, these number sequences can be shaped according to the desired distribution, e.g., a Gaussian one (SMITH, 1999). Note that information from the signal to be corrupted is usually not involved in the generating process, except for scaling purposes to adjust the noise level to the desired signal-to-noise ratio (SNR). The structure of such noise is therefore completely independent from the signal.

### A.1.1 Gaussian noise

Gaussian noise (normally distributed noise) is seen widely in nature. The term *Gaussian* refers to the shape of noise's amplitude distribution, the well-known bell curve that shows up whenever many different random forces are interacting. A Gaussian noise generator is easily implemented in software (SMITH, 1999). For convenience, an equation is given below that yields a sequence of  $N$

random numbers  $n_i$ , whose probability density function is Gaussian:

$$n_i = \sum_{j=1}^m k_j \quad \text{where}$$

- $i = 1, \dots, N$
- $k_j$ : uniformly distributed random numbers, e. g., output of a pseudorandom generator
- $m$ : determines the accuracy of the Gaussian shape. The larger is  $m$ , the better is the Gaussian shape.  $m = 12$  is a good choice in most cases.


If we consider the power spectrum of uncorrelated Gaussian noise, we will find it to be a flat line over the entire spectral range. In other words, no frequency dominates the time series. Such a spectrum is called *white*. Intuitively, we can see that this implies the samples in the time series to be independent of each other. As to why *white* is used to signify a flat power spectrum, think of white light. White light is composed of equal amounts of blue, red, orange, indigo, etc., light. Colored light has a non-flat distribution of frequencies, i. e., one color dominates.<sup>1</sup>

Because Gaussian noise is both ubiquitous in nature and easily generated artificially, it seems convenient to use such noise to evaluate the performance of noise reduction algorithms. However, keep in mind that this does not imply such noise to be the only (and the best) choice for the test of algorithms, particularly of those that deal with real-world biological noise.

## A.2 Signal-dependent noise models

The noise that contaminates real ERG signals is, besides technical artifacts, to a large extent due to a continuous background activity of non-retinal structures, e. g., skeletal muscle activity, especially if it occurs in the proximity of only one electrode (FROEHLICH and KAUFMAN, 1992). It is what the clinician encounters as blink, eye movement, and poor fixation artifacts (HECKENLIVELY and ARDEN, 1991). The effects can clearly be noticed in the ERG pictured in Figure A.1. Such noise can, in terms of amplitude and spectral distribution, appear

<sup>1</sup>The terms *white* and *Gaussian* are often confused. If noise is Gaussian, its *amplitude distribution* is bell-shaped. If noise is white, then its *power spectrum* is flat. Gaussian noise does not have to be necessarily white—most probably it is not, if, for example, some kind of amplitude correlation is involved. Noise can be Gaussian, or can be white, or can be both, or can be neither.



**Figure A.1:** ERG artifacts from excessive muscle activity. Flash ERG recorded from a dog that was not sufficiently sedated. Arrows indicate muscle activity following the peak of the *b*-wave (KOMÁROMY *et al.*, 2002).


surprisingly similar to the signal of interest, and it is therefore much more difficult to separate than, for example, power line noise with its sharp peak at the power line frequency.

It would be therefore convenient to have artificial noise available that mimicks the spectral distribution of (i. e., having the same Fourier power spectrum as) the signal. This surely provides a powerful and reasonable test for any noise reduction algorithm in the bio-physiological field. In other words, the generated noise has to preserve certain properties of the original data, but has to be random otherwise. In order to generate such noise we need to take information from the original signal into account.

### A.2.1 Spectrally matched noise

A couple of years ago, nonlinear scientists came up with a method, referred to as the *method of surrogate data*, that makes heavy use of the *fast Fourier transform* algorithm (FFT). This algorithm estimates the Fourier power spectrum of a given signal and yields random noise with a power spectrum rendered to match that of the signal.<sup>2</sup> Here, it is therefore referred to as *spectrally matched noise* (SMN). Once such kind of noise is added to its parent signal, it can never

<sup>2</sup>This is basically achieved by Fourier-transforming the signal, randomizing the Fourier phases while preserving the corresponding amplitudes, and a final inverse transform. The method is not foolproof; THEILER *et al.* (1992) provide a survey of such techniques along with their specific pitfalls. For an excellent, readable introduction of the computational issues of FFT, see LYONS (1997).



**Figure A.2:** An ERG signal and its corresponding SMN noise. Upper left: 30 Hz flicker response. Lower left: power spectrum of the flicker response. Upper right: SMN noise, generated from the flicker response. Lower right: power spectrum of the SMN noise. The SMN noise was scaled to match the variance of its parent data.

be removed again by any frequency filtering technique, neither by convolution filters (LYONS, 1997), nor even by Fourier filtering.

Figure A.2 shows a synopsis of a flicker sequence and of the constructed SMN noise data sets. Note the similar frequency distribution of both sets. A close inspection reveals slight differences between the corresponding power spectra: The peaks in the noise spectrum seem to have widened and decreased in amplitude. In the signal processing sciences, this is called *leakage*; it can be minimized, but not avoided completely, and is a result of the finite frequency resolution. Moreover, the phase randomisation in the noise generation process adds some slight frequency redistribution, as observed by THEILER *et al.* (1992), presumably due to the resulting phase discontinuities.



## Appendix B

### Standard deviation, RMS, and SNR

The *standard deviation*  $\sigma_x$  of a time series  $x$  is calculated as follows. Let  $x$  be the time series that contains  $N$  elements  $x_j$ , where  $j = 1, 2, \dots, N$ . The average  $\bar{x}$  of  $x$  is given by  $\bar{x} = \frac{1}{N} \sum_{j=1}^N x_j$ . The equation for  $\sigma_x$  then reads

$$\sigma_x = \sqrt{\frac{1}{N} \sum_{j=1}^N (x_j - \bar{x})^2}$$

It is interesting to compare the equation above to the calculation of the *root mean square* (RMS)  $x_{\text{RMS}}$  of the time series  $x$ :

$$x_{\text{RMS}} = \sqrt{\frac{1}{N} \sum_{j=1}^N (x_j)^2}$$

Obviously both equations differ only in how they consider  $\bar{x}$ ; the former does, the latter ignores it. What happens if  $\bar{x}$  approaches zero? Then both equations will yield the same result. In our electrophysiological context, a nonzero  $\bar{x}$  means that the input ERG response still has a nonzero DC offset (which amounts exactly to  $\bar{x}$ ). Such a static offset is not regarded as valid part of the signal, because what we expect is a time-dependent, dynamic response to a stimulus. This is one reason why in most ERG setups the DC component is rejected by means of built-in hardware high pass filters. We can therefore state that—for our purpose—the two properties of a signal *standard deviation* and *RMS* denote the same matter.

Ratio	dB	dB	Ratio
		-40	0.01
		-30	0.032
		-25	0.056
0.01	-40	-20	0.1
0.02	-34	-15	0.18
0.05	-26	-10	0.32
0.1	-20	-5	0.56
0.2	-14	-3	0.71
0.5	-6	-1	0.89
1	-0	0	1
2	6	1	1.1
5	14	3	1.4
10	20	5	1.8
20	26	10	3.2
50	34	15	5.6
100	40	20	10
		25	18
		30	32
		40	100

**Table B.1:** Some useful logarithmic relationships

The noise content of a signal can be described as a fraction termed *signal-to-noise ratio*, or SNR:

$$\text{SNR} = \frac{\sigma_{\text{signal}}^2}{\sigma_{\text{noise}}^2}$$

Besides this representation, another, nonlinear, representation of the SNR turned out to be more useful: it is the base 10 logarithm of the ratio above. The effect of this nonlinearity is to provide a greater resolution when the ratio is small, making it easier to recognize very small differences in the power levels of signal and noise.

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \left( \frac{\sigma_{\text{signal}}^2}{\sigma_{\text{noise}}^2} \right) = 20 \log_{10} \left( \frac{\sigma_{\text{signal}}}{\sigma_{\text{noise}}} \right)$$

Although a logarithm cannot have a physical dimension, in this special case a unit of measure, *bel*, named after A. G. BELL, was assigned to emphasize the difference between the bare SNR fraction and its logarithm. Because the unit of bel was soon found to be inconveniently large, one-tenth of it, the *decibel* (dB), now effectively replaced the unit bel.

# Appendix C

## The Fourier Transform

The Discrete Fourier Transform and its descendant, the *Fast Fourier Transform*<sup>1</sup> (FFT) provide a powerful way for data analyzing and processing in the frequency domain. The FFT enables us to analyze, manipulate, and synthesize signals in ways not possible with continuous (analog) signal processing. While once populating the engineering sciences as their most important application fields, FFT techniques continue to flourish as their utility becomes more widely understood.


FFT techniques in electrophysiological settings has been discussed in detail by, e. g., BACH and MEIGEN (1999). Although the scope of this thesis is not supposed to cover FFT techniques exhaustively, I refer to them in the main text here and there, and so this section is intended to give a brief overview.

The FFT is a mathematical procedure used to determine the harmonic, or frequency, content of a discrete signal sequence—for instance, a digitally sampled ERG response, which consists actually of a number of voltage values  $U(n)$  equally spaced in time. Let  $N$  be the number of samples in the sequence, and  $\tilde{U}(m)$  be the frequency content of  $U(n)$ . Then the calculation reads

$$\tilde{U}(m) = \sum_{n=0}^{N-1} U(n) \exp\left(\frac{-j2\pi nm}{N}\right) \quad \text{with} \quad \begin{array}{l} \bullet \exp(\dots) = e^{(\dots)}; e \text{ is the} \\ \text{base of natural logarithms} \\ \bullet j = \sqrt{-1} \end{array}$$

---

<sup>1</sup>The FFT is a subtype of the more general Discrete Fourier Transform (DFT). A FFT is much more faster to calculate compared to the DFT, but requires the analyzed time series to meet some conditions. Fortunately, these conditions are easy to fulfil, and so the FFT is almost always the method of choice when a frequency spectrum of a signal is asked for.



**Figure C.1:** The FFT tool window of CALYPSO. Shown is the frequency spectrum of a 30 Hz flicker ERG of a normal subject. The hairlines enclose a user-selectable frequency range, referred to as *region of interest* (ROI); the frequencies within the ROI can be scaled by a mouse move either to amplify or to attenuate the selected frequency range. The effect is a filter operation in the frequency domain. An inverse transform on the modified spectrum applies the desired frequency filtering to the original time-domain sequence.

While the above equation may seem a bit intimidating, a lot of FFT computer programs are available, which keeps the mess of implementing the math away from users who are only interested in applying its benefits. Anyway, the result of the FFT equation is a representation of the frequency content of the original time series.

Even if it were restricted to this application, FFT would yet be undoubtedly a very useful tool. But it works not only in the time-to-frequency direction, but the reverse direction as well. For example, if only the frequency spectrum of a signal is available, we can reconstruct the original time series by applying the counterpart of FFT, the *inverse FFT*. What is more, we can even manipulate the frequency content of a time sequence in the frequency domain, and apply these changes to the time series by the inverse FFT. Besides its plain analyzing capabilities, the Fourier transform can obviously also serve as powerful and fast frequency filter.

# List of Figures

1.1	Functional histology of the inner retina . . . . .	11
1.2	Latency and amplitude changes in cataract and retinopathy . . .	14
1.3	CALYPSO system architecture . . . . .	18
2.1	Construction of a state vector . . . . .	23
2.2	Cone ERG in 3D phase space . . . . .	24
2.3	Neighborhood with principal directions . . . . .	27
2.4	Screenshot of the CALYPSO user interface . . . . .	32
2.5	Hardware setup for the standard ISCEV recordings . . . . .	33
2.6	Artifacts within a flicker sequence . . . . .	37
3.1	Normal 30 Hz flicker ERG in different representations . . . . .	41
3.2	Synopsis of normal 30 Hz flicker ERGs . . . . .	42
3.3	Topography of a flicker trajectory . . . . .	43
3.4	A flicker trajectory in polar coordinates . . . . .	46
3.5	Flicker ERG of Usher’s syndrome . . . . .	47
3.6	Flicker ERGs of retinitis pigmentosa . . . . .	48
3.7	Flicker ERG of juvenile macular degeneration . . . . .	49
3.8	Flicker ERG of Stargardt’s disease . . . . .	50
3.9	NNR performance on different artificial noise . . . . .	54
3.10	NNR performance at different iterations . . . . .	55
3.11	NNR performance using the sinusoid trick . . . . .	56

4.1	Performance of different noise reduction methods . . . . .	61
4.2	Combining different noise reduction methods . . . . .	64
4.3	False Neighbors of a flicker ERG . . . . .	65
4.4	Effects of delay variations . . . . .	67
4.5	Estimating the noise level from a phase space plot . . . . .	68
5.1	2D embedding of an oscillatory potential ERG . . . . .	73
5.2	False color encoding of a flicker trajectory's third dimension . . .	74
A.1	Muscle activity corrupting an ERG . . . . .	79
A.2	Signal and corresponding SMN noise . . . . .	80
C.1	FFT window of CALYPSO . . . . .	84

# List of Tables

- 1.1 Characteristics of the rod and cone system . . . . . 13
- 2.1 Synopsis of the processing parameters . . . . . 36
- 4.1 Flicker trajectory topographic angles . . . . . 59
- B.1 Logarithmic relationships . . . . . 82

# Bibliography

- BACH, M., KELLNER, U. (2000): Elektrophysiologische Diagnostik in der Ophthalmologie. *Der Ophthalmologe* **97**(12): 898–920.
- BACH, M., MEIGEN, T. (1999): Do's and don'ts in Fourier analysis of steady-state potentials. *Documenta Ophthalmologica* **99**(1): 69–82.
- BIRCH, D. G., SANDBERG, M. A. (1996): Submicrovolt full-field cone electroretinograms: artifacts and reproducibility. *Documenta Ophthalmologica* **92**(4): 269–280.
- BRIGELL, M., BACH, M., BARBER, C., KAWASAKI, K., KOOIJMAN, A. (1998): Guidelines for Calibration of Stimulus and Recording Parameters Used in Clinical Electrophysiology of Vision. *American Journal of Physiology – Heart and Circulation Physiology* **95**(1): 1–14.
- BROWN, E. L. (1985): *The Flux Capacitor*. Universal Press, Los Angeles.
- CARR, R. E., SIEGEL, I. M. (1982): *Visual Electrodiagnostic Testing*. Williams & Wilkins, Baltimore.
- CAWLEY, R., HSU, G. H. (1992): Local-geometric-projection method for noise reduction in chaotic maps and flows. *Physical Review A* **46**(6): 3057–3082.
- DAIGER, S. P. (2003): RetNet: Summaries of Genes Causing Retinal Diseases. URL: <http://www.sph.uth.tmc.edu/RetNet/sum-dis.htm>
- DAWSON, W. W., TRICK, G. L., LITZKOW, C. A. (1979): Improved electrode for electroretinography. *Investigative Ophthalmology & Visual Science* **18**(9): 988–991.
- DINGWELL, J. B., CUSUMANO, J. P. (2000): Nonlinear time series analysis of normal and pathological human walking. *Chaos* **10**(4): 848–863.
- EINTHOVEN, W., JOLLY, W. (1908): The form and magnitude of the electrical response of the eye to stimulation at various intensities. *Quarterly Journal of Experimental Physiology* **1**: 373–416.



- FELL, J., MANN, K., ROSCHKE, J., GOPINATHAN, M. S. (2000): Nonlinear analysis of continuous ECG during sleep. I. Reconstruction. *Biological Cybernetics* **82**(6): 477–483.
- FROEHLICH, J., KAUFMAN, D. I. (1992): Improving the reliability of pattern electroretinogram recording. *Electroencephalography and Clinical Neurophysiology* **84**(4): 394–399.
- GOURAS, P. (1970): Electroretinography: Some basic principles. *Investigative Ophthalmology & Visual Science* **9**(8): 557–569.
- GRANIT, R. (1947): *Sensory Mechanisms of the Retina With an Appendix on Electroretinography*. Oxford University Press, London.
- HECKENLIVELY, J. R., ARDEN, G. B. (1991): *Principles and practice of visual electrophysiology*, chap. 49. Mosby-Year Book, Inc., St. Louis, 384–390.
- HÉNON, M. (1976): A two-dimensional mapping with a strange attractor. *Communications in Mathematical Physics* **50**(1): 69–77.
- HOGG, C., RILK, A. J. (2002): Technische Aspekte bei der Ableitung von ERGs und VEPs. In: *TED – Tübinger Elektro-Diagnostikkurs 2002*, Universitäts-Augenklinik Tübingen.
- HOOD, D. C., BIRCH, D. G. (1990): A quantitative measure of the electrical activity of human rod photoreceptors using electroretinography. *Visual Neuroscience* **5**(4): 379–387.
- INTERNATIONAL SOCIETY FOR CLINICAL ELECTROPHYSIOLOGY OF VISION (2001): Visual Electrodiagnostics: A Guide to Procedures.  
URL: <http://www.iscev.org/standards/proceduresguide.html>
- JACOBI, P. C., RÜTHER, K., MILICZEK, K.-D., VÖLKER, M., ZRENNER, E. (1993): Klinische Elektroretinographie: Standardprotokoll und Normwerte. *Klinische Monatsblätter Augenheilkunde* **202**(1): 27–42.
- KANTZ, H., SCHREIBER, T. (1998): *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge.
- KERN, A., STEEB, W.-H., STOOP, R. (2000): Projective noise cleaning with dynamic neighborhood selection. *International Journal of Modern Physics C* **11**(1): 125–146.
- KOMÁROMY, A. M., BROOKS, D. E., DAWSON, W. W., KÄLLBERG, M. E., OLLIVIER, F. J., OFRI, R. (2002): Technical issues in electrodiagnostic recording. *Veterinary Ophthalmology* **5**(2): 85–91.

- KRETSCHMANN, U., BOCK, M., GOCKELN, R., ZRENNER, E. (2000): Clinical applications of multifocal electroretinography. *Documenta Ophthalmologica* **100**(2–3): 99–113.
- LYONS, R. G. (1997): *Understanding digital signal processing*. Addison Wesley Longman, Reading, Mass.
- MARMOR, M. F., ZRENNER, E. (1999): Standard for clinical electroretinography (1999 update). *Documenta Ophthalmologica* **97**(2): 143–156.
- MEISS, J. D. (2000): Frequently asked questions about nonlinear science, version 1.5.  
URL: <ftp://amath.colorado.edu/pub/dynamics/papers/sci.nonlinearFAQ.pdf>
- NIEMEYER, G. (1998): Selective rod- and cone-ERG responses in retinal degenerations. *Digital Journal of Ophthalmology* **4**(4).  
URL: <http://www.djo.harvard.edu/OA/NIEMEYER/INDEX.html>
- OTT, E., SAUER, T., YORKE, J. A. (1994): *Coping with chaos*. Wiley, New York.
- PACKARD, N. H., CRUTCHFIELD, J. P., FARMER, J. D., SHAW, R. S. (1980): Geometry from a time series. *Physical Review Letters* **45**(9): 712–716.
- QUYEN, M. L. V., MARTINERIE, J., ADAM, C., VARELA, F. J. (1997): Unstable periodic orbits in human epileptic activity. *Physical Review E* **56**(3): 3401–3411.
- RICHTER, M., SCHREIBER, T., KAPLAN, D. T. (1998): Fetal ECG extraction with nonlinear state-space projections. *IEEE Transactions on Biomedical Engineering* **45**(1): 133–137.
- RILK, A. J., SHARPE, L. T. (2000): Noise reduction in the ERG by nonlinear signal processing. *Investigative Ophthalmology & Visual Science* **41**(4): S498.
- RISPOLI, E., IANNACCONE, A., VINGOLO, E. M. (1994): Low-noise electroretinogram recording techniques in retinitis pigmentosa. *Documenta Ophthalmologica* **88**(1): 27–37.
- SAUER, T. (1992): A noise reduction method for signals from nonlinear systems. *Physica D* **58**(1–4): 193–201.
- SAUER, T., YORKE, J. A., CASDAGLI, M. (1991): Embedology. *Journal of Statistical Physics* **65**(3–4): 579–616.
- SCHREIBER, T., RICHTER, M. (1999): Fast nonlinear projective filtering in a data stream. *International Journal of Bifurcation and Chaos* **9**(10): 2039–2045.

- SEELIGER, M. W., ZRENNER, E., APFELSTEDT-SYLLA, E., JAISSE, G. B. (2001): Identification of Usher Syndrome Subtypes by ERG Implicit Time. *Investigative Ophthalmology & Visual Science* **42**(12): 3066–3071.
- SIEVING, P. A., MURAYAMA, K., NAARENDORP, F. (1994): Push-pull model of the primate photopic electroretinogram: a role for hyperpolarizing neurons in shaping the b-wave. *Visual Neuroscience* **11**(3): 519–532.
- SMITH, S. W. (1999): *The Scientist and Engineer's Guide to Digital Signal Processing*. California Technical Publishing, San Diego.  
URL: <http://www.dspguide.com/pdfbook.htm>
- STARK, J. (1999): Delay embeddings for forced systems. I. Deterministic forcing. *Journal of Nonlinear Science* **9**(3): 255–332.
- STEINBERG, R. H., OAKLEY, B., NIEMEYER, G. (1980): Light-evoked changes in  $K_0^+$  in retina of intact cat eye. *Journal of Neurophysiology* **44**(5): 897–921.
- SUTTER, E. E., TRAN, D. (1992): The field topography of ERG components in man – I. The photopic luminance response. *Vision Research* **32**(3): 433–446.
- TAKENS, F. (1981): Detecting strange attractors in turbulence. In: D. A. RAND, L. S. YOUNG (editors): *Dynamical systems and turbulence, Warwick 1980 (Coventry, 1979/1980)*, Lecture Notes in Mathematics, Vol. 898, Springer, Berlin, New York, 366–381.
- THEILER, J., EUBANK, S., LONGTIN, A., GALDRIKIAN, B., FARMER, J. D. (1992): Testing for nonlinearity in time series: the method of surrogate data. *Physica D* **58**(1–4): 77–94.
- ZRENNER, E. (1983): Grundlagen elektroophthalmologischer Untersuchungsmethoden in der Augenheilkunde. In: *Hauptreferate der XVIII. Essener Fortbildung für Augenärzte*, Enke, Stuttgart, 129–140.

# Danksagung

Die vorliegende Arbeit entstand zwischen 1999 und 2002 an der Universitäts-Augenklinik in Tübingen in der Abteilung für Neuroophthalmologie und Pathophysiologie des Sehens. Diese Abteilung verfügt traditionell über große Erfahrung bei der Eigenentwicklung von Geräten für die physiologische Forschung – ein Umstand, der mir bei vielen Aspekten dieser Dissertation in idealer Weise entgegenkam. Ich danke deshalb allen Kolleginnen und Kollegen für ihre Unterstützung. Mein besonderer Dank gilt

dem Ärztlichen Direktor *Herrn Prof. Dr. E. Zrenner* für seine jahrelange wohlwollende, großzügige Förderung und für die Betreuung dieser Arbeit.

*Herrn Dipl.-Ing. Stefan Beuel*, denn sein immenses „know how“, sein Ideenreichtum und sein Sinn für Humor ließen die täglichen Hürden, die den Entwicklungsprozess begleiteten, leichter nehmen. Dies gilt sowohl für die Umsetzung der mathematischen Algorithmen in funktionierende und robuste Programme, als auch für das Design der diversen Elektronik zur Datenakquisition, ohne die CALYPSO sozusagen taub und blind geblieben wäre.

*Herrn Dr. med. Dipl.-Phys. Jens Reinhard*, *Herrn Prof. Dr. Lindsay T. Sharpe*, *Herrn Dipl.-Ing. Wilhelm Durst* und *Herrn Dr. Herbert Jäggle* für Tipps und Anregungen. *Jens* insbesondere danke ich für seine unermüdliche Motivation.

*Frau Ulrike Kefßler* und *Frau Annegret Edeko* für die Hilfe bei der Rekrutierung der Patienten, *Frau Regina Nicolaidis, MA*, und *Frau Doina Ruppert, MA*, für ihr freundliches und pragmatisches Management.

*Meinen Eltern* danke ich für ihren weisen Rat, ihre beständige Unterstützung und ihr Vertrauen. Mein Lebensweg und die beiden Studiengänge wurden dadurch erst möglich.

Ganz besonders danke ich *Ines*. Ihre liebevolle Unterstützung und Geduld habe ich durch mehr als nur diese Arbeit beansprucht. Ich danke ihr dafür, dass ich jederzeit auf sie zählen kann.

Ich danke herzlich der *Waldtraut und Sieglinde Hildebrandt-Stiftung*, die mir für diese Promotion ein Stipendium gewährte.

# Lebenslauf

03.12.1969	geboren in Aalen
1976–1989	Schulzeit in Königsbronn und Heidenheim/Brenz
Mai 1989	Abitur am Hellenstein-Gymnasium in Heidenheim/Brenz
1989–1990	Zivildienst
Oktober 1990	Beginn des Studiums der Physik an der Universität Tübingen
Oktober 1992	Vordiplom
1995	Praxissemester bei Carl Zeiss, Oberkochen
1996	Diplomarbeit über das Thema <i>Detektorsystem zum Test eines Hardwareprozessors des Satelliten INTEGRAL</i>
Januar 1997	Diplom in Physik
April 1997	Beginn des Studiums der Medizin an der Universität Tübingen
April 1999	Ärztliche Vorprüfung (Physikum)
April 2000	Erstes Staatsexamen
April 2002	Zweites Staatsexamen
April 2002–März 2003	Praktisches Jahr
Dez. 2002–März 2003	Unterassistentzeit in der Abteilung für Chirurgie des Spitals Zofingen, Schweiz
23.05.2003	Drittes Staatsexamen
1996–2003	Tätigkeiten an der Universitäts-Augenklinik Tübingen, u. a. im RetinaImplant-Projekt (Prof. Dr. E. Zrenner), im Elektronik-Entwicklungslabor (Dipl.-Ing. S. Beuel), und in der AG für Visuelle Psychophysik (Prof. Dr. L. T. Sharpe)
1999–2002	Anfertigung der medizinischen Dissertation über das Thema <i>The Flicker Electroretinogram in Phase Space: Embeddings and Techniques</i>