
Albert-Ludwigs-Universität Freiburg
Fakultät für Physik

Maximum likelihood estimation in dynamical systems

Inauguraldissertation
zur Erlangung des Doktorgrades
der Fakultät für Physik
der
Albert-Ludwigs-Universität Freiburg im Breisgau

vorgelegt von
Werner Horbelt
aus Nürnberg
Juni 2001

Dekan: Prof. Dr. K. Königsmann

Leiter der Arbeit: Prof. Dr. J. Honerkamp

Referent: Prof. Dr. J. Honerkamp

Koreferent: Prof. Dr. N.N.

Tag der Verkündigung des Prüfungsergebnisse: 16.07.2001

Publications

Horbelt, W., Timmer, J., Bünner, M., Meucci, R., and Ciofini, M. (2001b). Identifying physical properties of CO₂ laser by dynamical modeling of measured time series. *To be published in Phys. Rev. E*, 64

Timmer, J., Rust, H., Horbelt, W., and Voss, H. (2000b). Parametric, non-parametric and parametric modelling of a chaotic circuit time series. *Phys. Lett. A*, 274:123–134

Gombert, A., Rose, K., Heinzl, A., Horbelt, W., Zanke, C., Blasi, B., and Wittwer, V. (1998b). Antireflective submicrometer surface-relief gratings for solar applications. *Solar Energy Materials and Solar Cells*, 54:333–342

Gombert, A., Glaubitt, W., Rose, K., Dreiholz, J., Zanke, C., Blasi, B., Heinzl, A., Horbelt, W., Sporn, D., Doll, W., Wittwer, V., and Luther, J. (1998a). Glazing with very high solar transmittance. *Solar Energy*, 62:177–188

Horbelt, W., Timmer, J., Bünner, M., Meucci, R., and Ciofini, M. (2001a). Dynamical modelling of measured time series from a Q-switched CO₂ laser. *Chaos, Solitons & Fractals*, submitted

Conference Contributions

Horbelt, W., Timmer, J., and Melzer, W. (1998b). Estimating parameters in differential equations with application to physiological data. In Osipenko, G., editor, *Differential equations and applications*, pages 23–33. State Technical University, Saint Petersburg

Horbelt, W., Timmer, J., Bünner, M., and Honerkamp, J. (1998a). Estimating parameters in nonlinear differential equations with application to Q-switched lasers. In *Technische Anwendungen von Erkenntnissen der Nichtlinearen Dynamik*, pages 311–314. VDI, Frankfurt

Horbelt, W., Müller, T., Timmer, J., Melzer, W., and Winkler, K. (2000). Analysis of nonlinear differential equations: parameter estimation and model selection. In Brause, R. and Hanisch, E., editors, *Medical Data Analysis*, volume 1933 of *Lecture Notes in Mathematics*, page 152. Springer, Berlin

Rust, H., Timmer, J., Horbelt, W., and Voss, H. U. (2000). Parametric, nonparametric and parametric modelling of a chaotic circuit time series. In Oishi, S. and Schwarz, W., editors, *Proceedings of the 2000 International*

Symposium of Nonlinear Theory and its Applications, NOLTA 2000, volume 2, pages 713–716, Dresden. w.e.b.-Univ.-Verl

Timmer, J., Horbelt, W., Bünner, M., Ciofini, M., and Meucci, R. (2000a). Estimating parameters in differential equations with application to laser data. In D.S. Broomhead, E.A. Luchinskaya, P. M. and Mullin, T., editors, *Stochaos: Stochastic and Chaotic Dynamics in the Lakes*, pages 617–623, Melville, NY, USA. American Institute of Physics

Contents

1	Introduction	1
2	Methods	5
2.1	Estimating Parameters	5
2.1.1	Solution of the Optimisation Problem	6
2.1.2	Statistics	8
2.1.3	Multiexperiment Analysis	8
2.2	Ordinary Differential Equations	8
2.2.1	Initial-Value Approach	9
2.2.2	Multiple Shooting Approach	9
2.2.3	Calculating Sensitivities	12
2.3	Time-discrete Systems	14
2.4	Delay Differential Equations	15
2.4.1	Direct Problem	16
2.4.2	Calculating Sensitivities	16
2.4.3	Multiple Shooting	17
2.5	Other classes	20
3	Time-discrete Systems	23
3.1	Introduction	23
3.2	Logistic Map without Noise	25
3.3	Shift Map	28
3.4	Logistic Map with Noise	29
3.5	Correlation	34
3.6	Conclusion	37
4	The Q-switched CO₂ Laser	39
4.1	Experimental Setup	40
4.2	Preprocessing	42
4.2.1	Baseline Correction	42
4.2.2	Pulse Merging	43
4.2.3	Calibration Procedure and Observation Equation	47
4.3	Theoretical Investigation	49

4.3.1	Dynamical Model	49
4.3.2	Laser Startup Process	53
4.3.3	Consistency of the Initial State Vector	56
4.4	Simulation Study	57
4.5	Experimental Results	57
4.5.1	Low Currents	59
4.5.2	Medium and High Currents	63
4.5.3	Estimation of n_ξ and K_1	67
4.6	Conclusion	67
5	Delay Differential Equations	69
5.1	The Delay Chain Approach	69
5.1.1	An Example from Infectious Disease Modelling	70
5.1.2	Modelling with ODEs	72
5.1.3	Modelling with the Delay Chain Model	72
5.1.4	Discussion	75
5.2	The Mackey-Glass System	76
5.2.1	Known Time Lag	77
5.2.2	Unknown Time Lag	79
5.2.3	Systematic Test with Random Starting Guesses	81
5.2.4	Summary	84
5.3	Application to Measured Data	84
6	Summary	89
A	Equivalence of Maps	93
	Acknowledgements	95
	Bibliography	97

Symbols

$\hat{\boldsymbol{\theta}}, \hat{a}, \hat{x}_0, \dots$	Estimator of $\boldsymbol{\theta}, a, x_0, \dots$
\dot{x}	Time derivative of x
$\langle . \rangle$	Expectation value
\sim	... is proportional to ...
\approx	... is approximately equal to ...
χ^2	Objective function
δ	Dirac delta function, population inversion
δ_{ij}	Kronecker delta symbol
λ	Lyapunov exponent or laser net gain
$\boldsymbol{\theta}$	Vector of unknown variables (parameters)
$\Delta\boldsymbol{\theta}$	Update step
σ	Standard deviation
τ	Delay, time lag
D	Separation of multiple shooting intervals for DDEs
\boldsymbol{f}	Right-hand side of a dynamic equation
g	Observation function
h_j	Initial curve of a DDE in j th multiple shooting interval
K	Number of data points per multiple shooting interval
M	Number of multiple shooting intervals
m^e	Number of equality constraints
m^g	Number of inequality constraints
N	Number of data points
\boldsymbol{p}	Vector of dynamic parameters
s_j	Spline variable in j th multiple shooting interval
\boldsymbol{S}	Sensitivity matrix
t	Time
Δt	Sampling interval
T_0	Beginning of fitting interval
T_j	Beginning of j th multiple shooting subinterval
T	Length of fitting interval
$\boldsymbol{x}(t), \boldsymbol{x}_i$	State vector for continuous and discrete systems
x_i	Component of \boldsymbol{x}
\boldsymbol{x}^0	Initial values (not to be confused with starting guesses)
x_τ	Retarded state variable in DDEs
y_i	Measurement

Chapter 1

Introduction

*The eye of the cat is blind to the mouse
when the mouse is still.*

Time series analysts expect to be able to extract more information from dynamical processes than from steady states. From the observation of dynamical systems they hope to obtain insight into the laws ruling their time evolution and thus to gain knowledge of the systems themselves.

Fig. 1.1 on page 3 shows an overview of important types of deterministic dynamical systems, classified according to the mathematical equations used to model them. Though a rather broad range of mathematical theory is touched by them, they have some common properties. There is a state vector evolving in time according to some dynamic equation. Typically the dynamics depends on unknown parameters and the trajectory is only partially observed. The aim is to estimate the parameters from noisy measurements of the time course and to construct the unobserved components of the system.

Considering the right-hand side of a dynamical system, one can distinguish between truly *parametric models*, in which all parameters are supposed to have a physical meaning, and *black box models* in which the right-hand side belongs to a rather general class of functions, e.g. polynomials or radial basis functions (Crutchfield and McNamara 1987; Giona et al. 1991; Weigend et al. 1990; Casdagli 1991). While the latter models do a good job when only secondary properties of the dynamical system like Lyapunov exponents or in-sample predictions are sought, parametric models should be preferred when the aim is to gain insight into the process itself or to predict the behaviour of the system in regions of the state space that are not covered by measured data.

Sometimes educated guesses are made for the parameters and the resulting model predictions are compared with measured data by visual inspection (Arecchi et al. 1988; Ciofini et al. 1993; Meucci et al. 1992; Weiss et al. 1995;

Zehnlé et al. 1992; Gouesbet et al. 1996). In the absence of a systematic adaptation of the mathematical models to the data, it is not possible to distinguish between discrepancies that stem from wrongly chosen parameters in a sufficient model on the one hand and shortcomings of the model itself on the other. A fair comparison between two different models requires that the models be optimised with respect to their parameters first.

Many methods estimate the parameters systematically but regard each data point as independent of all but the immediately preceding one (Crutchfield and McNamara 1987; Giona et al. 1991; Corrêa et al. 2000; Gouesbet et al. 1996; Hegger et al. 1998; Cremers and Hübler 1987). They neglect important information about the data, namely that they all stem from a single underlying trajectory. As a consequence the estimates are less precise than they could be.

This thesis considers maximum likelihood methods for the estimation of the parameters and the unobserved components. These methods take into account the entire information about the deterministic nature of the underlying true trajectory, i.e., they determine the model trajectory that satisfies the dynamical equations and comes closest to the observed data with respect to an appropriately defined cost function. Maximum likelihood methods are asymptotically unbiased and most efficient with respect to the variance of the estimates. The difficulty is to find the global minimum of the cost function. Since the solution of the dynamic equations is rather sensitive to the parameters, there will usually exist numerous local minima apart from the global one. For ordinary differential equations (ODEs), this problem is overcome by means of the *multiple shooting* technique developed by Bock (1981, 1983). An equivalent method will be developed and tested for delay differential equations (DDEs).

This thesis can by no means treat all of the systems in Fig. 1.1. It considers time-discrete systems, ODEs and DDEs in some detail. Chapter 2 describes methods to estimate parameters in dynamical systems in general and in these three classes in particular. The following three chapters are each devoted to one of the classes. Chapter 3 examines the strengths and limits of the maximum likelihood approach for time-discrete systems, using the logistic map as an example. While in this chapter simulated data are used, Chapter 4 describes an application to measured data from a CO₂ laser, using a five-dimensional ODE. In order to emphasise the fact that experimental data represent a much greater challenge than theoretical simulations, this chapter describes comprehensively all steps undertaken to obtain the results.

Delay differential equations are considered in Chapter 5. Sometimes DDEs are approximated by ODEs with additional compartments. The usefulness of this technique is examined with an example from infectious disease modelling. Finally, the multiple shooting method for DDEs is applied to the Mackey Glass system and to measured data from an electronic circuit.

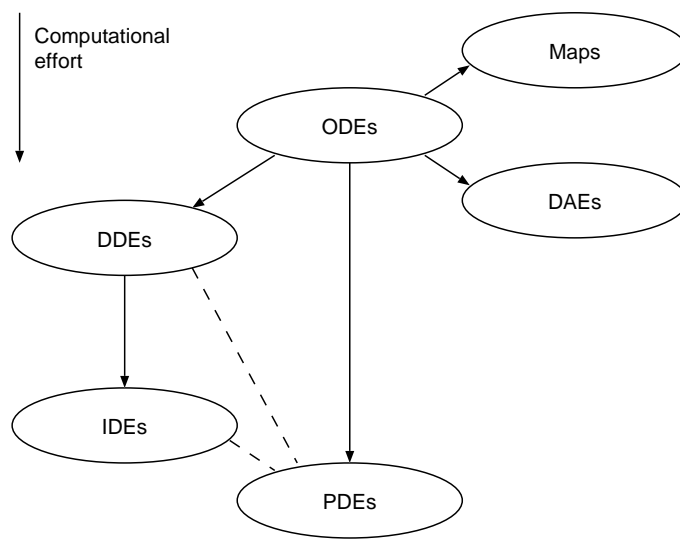


Figure 1.1: Overview of classes of deterministic dynamical systems. ODEs: ordinary differential equations; DDEs: delay differential equations; IDEs: integro-differential equations; DAEs: differential algebraic equations; PDEs: partial differential equations. The vertical arrangement reflects the computational effort associated with them. The arrows between the bubbles are to be read as “are a special case of”. The dashed lines indicate that DDEs and IDEs can also be seen as PDE-like systems.

Chapter 2

Methods

This chapter provides the methods needed for the estimation of parameters in the applications presented in Chapters 3–5. Though the classes outlined in Fig. 1.1 are qualitatively different, they share some properties and so do the methods used to treat them. Section 2.1 describes these common methods. Hereafter, methods are described that are specific to the different classes. The last section makes some remarks concerning the classes not treated here.

2.1 Estimating Parameters in Deterministic Dynamical Systems

Besides the specific form of the dynamical equations there is always a state vector \mathbf{x} evolving in time, thereby defining a trajectory that depends on a set of unknown parameters $\boldsymbol{\theta}$:

$$\mathbf{x} = \mathbf{x}(t, \boldsymbol{\theta}), \quad t \in I, \quad (2.1)$$

where I is some subset of \mathbb{R} . In general, $\mathbf{x}(t, \boldsymbol{\theta})$ is not known explicitly. For instance, I may be a time interval and $\mathbf{x}(t, \boldsymbol{\theta})$ the solution of an ordinary differential equation for the given parameters. The aim is to estimate the unknown parameters from noisy measurements of the state vector.

The very general Eq. (2.1) relates the state vector \mathbf{x} at a time t to the unknown parameters $\boldsymbol{\theta}$. As a second part, the complete description of a model contains the *observation equation*, the relation between the state vector and the measured data. Normally not all components of \mathbf{x} can be accessed experimentally and the observables, i.e. those that can be accessed, are often subject to nonlinear transformations that contain additional unknown parameters. For ease of notation the following description is restricted to scalar observations. The generalisation to multivariate data is straightforward. The time series $\{y_i\}$ is a measurement of the state vector at discrete

times t_i via the observation function g :

$$y_i = g(\mathbf{x}(t_i, \boldsymbol{\theta}), \boldsymbol{\theta}) + \eta_i, \quad i = 1, \dots, N. \quad (2.2)$$

Here, η_i denotes independent normally distributed random numbers with zero mean and variance σ_i^2 , accounting for measurement noise.

It must be emphasised that the observation equation is just as important as the dynamical equations of the model underlying Eq. (2.1). Both together define the *direct problem*. The aim here is to estimate the parameters from measured data, i.e. to solve the *inverse problem*.

A well established estimator for $\boldsymbol{\theta}$ is the maximum likelihood estimator. It is defined as the vector that minimises the likelihood of the measured data, given $\boldsymbol{\theta}$:

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} L(\{y_i\}|\boldsymbol{\theta}).$$

In the case of white Gaussian observation noise, it is equivalent to minimise the *objective function*, which is the sum of squared residuals between the data and the model trajectory, weighted with the inverse variance of the noise:

$$\chi^2(\boldsymbol{\theta}) = \sum_{i=1}^N \left(\frac{y_i - g(\mathbf{x}(t_i, \boldsymbol{\theta}), \boldsymbol{\theta})}{\sigma_i} \right)^2 \quad (2.3)$$

Often additional knowledge about the parameters can be formulated as equality or inequality constraints. For instance, rate constants in compartment models are non-negative. In the example of the CO₂ laser in Section 4.3.3, relations between the initial values of the ODE and the parameters are used to reduce the number of degrees of freedom. Finally, in the context of the multiple shooting approach described later on, the continuity of the final trajectory is ensured by means of equality constraints.

All these constraints and the least squares minimisation together amount to a nonlinear optimisation problem with nonlinear equality and inequality constraints: find the vector of parameters $\boldsymbol{\theta}$ such that

$$\chi^2(\boldsymbol{\theta}) = \sum_{i=1}^N R_i^a(\boldsymbol{\theta})^2 = \min, \quad (2.4a)$$

$$R_i^e(\boldsymbol{\theta}) = 0, \quad i = 1, \dots, m^e, \quad (2.4b)$$

$$R_i^g(\boldsymbol{\theta}) \geq 0, \quad i = 1, \dots, m^g. \quad (2.4c)$$

2.1.1 Solution of the Optimisation Problem

Since at least R^a and R^e are nonlinear in general, iterative algorithms must be used for the solution. Gill et al. (1981) give a good overview of various approaches. For an efficient optimisation, at least first derivatives with respect to the parameters (*sensitivities*) should be provided. For the applications reported later on, two well-established optimisation algorithms were used.

The first one is a generalised Gauss-Newton method. It is well suited to least squares minimisation because it exploits the structure of χ^2 as a sum of squared residuals. In the k -th iteration step, the nonlinear functions R_i^a , R_i^e and R_i^g are approximated by their first order Taylor expansions around the current parameter estimate $\boldsymbol{\theta}_k$, resulting in the following linear optimisation problem with linear equality and inequality constraints:

$$\sum_{i=1}^N (R_i^a(\boldsymbol{\theta}_k) + \frac{\partial R_i^a(\boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}} \cdot \Delta \boldsymbol{\theta})^2 = \min, \quad (2.5a)$$

$$R_i^e(\boldsymbol{\theta}_k) + \frac{\partial R_i^e(\boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}} \cdot \Delta \boldsymbol{\theta} = 0, \quad i = 1, \dots, m^e, \quad (2.5b)$$

$$R_i^g(\boldsymbol{\theta}_k) + \frac{\partial R_i^g(\boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}} \cdot \Delta \boldsymbol{\theta} \geq 0, \quad i = 1, \dots, m^g. \quad (2.5c)$$

It is solved by the code LSEI (Hanson and Haskell 1982).

The solution of the linear system yields an update step $\Delta \boldsymbol{\theta}_k$ to the parameters. The improved estimate is given by $\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k + \lambda \Delta \boldsymbol{\theta}$, where $\lambda \leq 1$ is an appropriately chosen *damping stepsize* (Bock 1987). λ is smaller than 1 when the problem is highly non-linear and equal to 1 during the final iterations. The rank-deficient case, e.g. when parameters or combinations of parameters cannot be identified reliably, is treated by an appropriately chosen (and manually adjustable) rank decision criterion. When the linearised problem is rank deficient, the corresponding linear combination of parameters is considered undeterminable and is not changed. When the algorithm signals rank deficiency at the convergence point, this indicates that the model is over-parameterised.

The difference between the left hand side of Eq. (2.5a) and the second order Taylor approximation of $\chi^2(\boldsymbol{\theta}_k + \Delta \boldsymbol{\theta})$ is the term

$$\sum_i \Delta \boldsymbol{\theta}^t \mathbf{G}_i \Delta \boldsymbol{\theta} R_i^a$$

with the matrix

$$\mathbf{G}_i = \frac{\partial^2 R_i^a}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^t}.$$

For correctly specified models, the residuals R_i^a are small at the convergence point and the method shows a nearly quadratic convergence behaviour. However, in practical applications, one often accepts some degree of model mis-specification for the sake of simplicity or because the measurements do not contain enough information for the estimation of the most comprehensive model.

Furthermore, \mathbf{G}_i is large, if R^a is extremely nonlinear, as in the case of chaotic maps. Then even small residuals can render the convergence linear or even result in a repelling behaviour at the minimum of χ^2 . For these cases a *quasi-Newton* method with *BFGS-update* was used, implemented in the NAG routine E04UNF (Gill et al. 1981; NAG 1997).

2.1.2 Statistics

At the convergence point, the matrix

$$\mathbf{C} = \left(\frac{1}{2}\mathbf{H}\right)^{-1}$$

is computed, where

$$H_{ij} = \frac{\partial^2 \chi^2(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j}$$

is the *Hesse matrix* of the objective function. If some prerequisites are fulfilled, the errors of the parameters are multivariately normally distributed with covariance matrix $\mathbf{C} = \{c_{ij}\}$ (Honerkamp 1994). Independent 95% confidence intervals for θ_i are given by $\hat{\theta}_i \pm 1.96 \cdot (c_{ii})^{1/2}$. The prerequisites are that the model is specified correctly, the global minimum of $\chi^2(\boldsymbol{\theta})$ is found and its quadratic approximation holds in a large enough region around the solution point.

An eigenvalue analysis of \mathbf{C} reveals parameters or linear combinations of parameters that are not identifiable from the given data. This happens, for instance, when there is a continuous ambiguity in the model, i.e. , the model is invariant under a continuous family of transformations that alter the state vector and the parameters, but not the observations. χ^2 itself serves as a criterion for the ability of the model to describe the given experimental observations (*goodness of fit*).

Overfitting A good correspondence between the model trajectory and the measurements has little significance if it has been achieved by adjusting a large number of free parameters. *Overfitting* is when the number of degrees of freedom is not in a reasonable relation to the amount of information contained in the data. This can be detected by disproportionately large confidence intervals.

2.1.3 Multiexperiment Analysis

Often there are several data sets that one hopes to describe by the same model and at least partially with the same parameters. The method is able to take into account all data sets simultaneously. Each parameter can either be varied independently, forced to be the same for all records or fixed to a given value. This is called a *multiexperiment analysis*.

2.2 Ordinary Differential Equations

Ordinary Differential Equations (ODEs) are the canonical way of describing time-continuous dynamical systems. Consider a dynamical process described by n nonlinear ordinary differential equations of first order that

depend on the unknown dynamical parameters \mathbf{p} . The *initial value problem* is defined by the ODE together with an initial condition:

$$\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}, \mathbf{p}), \quad \mathbf{x} \in \mathbb{R}^n, \quad t \in [T_0, T_0 + T], \quad (2.6a)$$

$$\mathbf{x}(T_0) = \mathbf{x}^0 \quad (2.6b)$$

In general, the initial values are themselves unknown parameters. They are contained in the vector of unknown parameters $\boldsymbol{\theta}$:

$$\boldsymbol{\theta} = (\mathbf{p}, \mathbf{x}^0) \in \mathbb{R}^{p+n}. \quad (2.7)$$

The function \mathbf{f} is called the *right-hand side* of the differential equation. If it is linear in \mathbf{x} , the solution can be given analytically in terms of exponential functions and sine waves. Nonlinear ODEs, on the other hand, can give rise to completely irregular behaviour and a wide spectrum of interesting phenomena even in as few as three dimensions, as Lorenz showed with his famous example (Lorenz 1963).

2.2.1 Initial-Value Approach

If the problem of parameter identification in nonlinear ODEs is treated in a straightforward way, one is led to the *initial-value approach*: starting from some initial guess of the parameter vector, the dynamic equations are solved over the entire fitting interval and $\chi^2(\boldsymbol{\theta})$ is evaluated. Then $\chi^2(\boldsymbol{\theta})$ is minimised iteratively as described in Section 2.1.1.

Compared with regression, where an explicit function is used as a model, this task is more difficult because the ODE trajectory $\mathbf{x}(t, \boldsymbol{\theta})$ is rather sensitive to the parameters, especially in the case of complex dynamical systems. Badly chosen initial estimates can preclude most methods from yielding any solution at all because the trial trajectory might diverge. Since $\chi^2(\boldsymbol{\theta})$ shows a highly non-linear dependence on $\boldsymbol{\theta}$, it will usually have numerous local minima apart from the global one that corresponds to the true parameters. For these reasons the initial-value approach is often not successful for parameter estimation in nonlinear ODEs.

To demonstrate these problems, a time series of the Lorenz system (Hornbelt et al. 1998b) of length $T=4$ was simulated with the standard parameters $\sigma = 10$, $r = 46$, $b = 2.667$ and the initial values $x_0 = 5.7654$, $y_0 = 10.50547$ and $z_0 = 30.58941$. The sampling interval was 0.04 and Gaussian noise with a noise level of 20% was added. Then the parameters were estimated with the initial-value approach. As starting guesses the parameters have been set to twice the true values. The outcome can be seen in Fig. 2.1 on the next page. The optimisation procedure is trapped in a local minimum.

2.2.2 Multiple Shooting Approach

The technique of multiple shooting has been used for solving boundary value problems in ODE systems (Stoer and Bulirsch 1993). In the context of

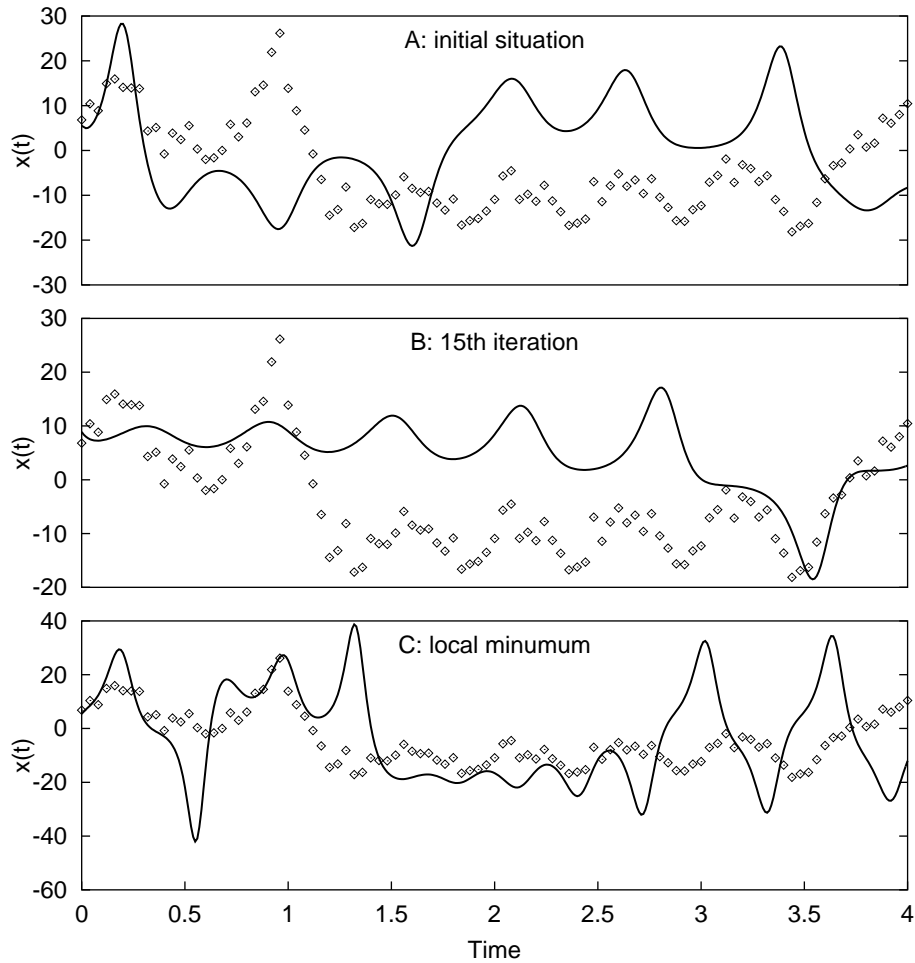


Figure 2.1: Failure of the initial-value approach for the Lorenz system. Right from the start the model trajectory is far away from the data points, rendering them useless for the optimisation process.

parameter estimation, the method was introduced in van Domselaar and Hemker (1975) and, in a much more general context, by Bock (1981, 1983). The motivation to use multiple shooting for parameter estimation is that the initial-value approach effectively neglects information on the dynamics of the system present in the measurements. Even though the time course of at least one component, the observation, is known rather accurately, the initial-value approach does not take advantage of any but the very first observation in the fitting interval. If the parameters are far off from the correct ones, the trial trajectory soon loses contact with the measurements.

For the multiple shooting approach, the fitting interval $[T_0, T_0 + T]$ is

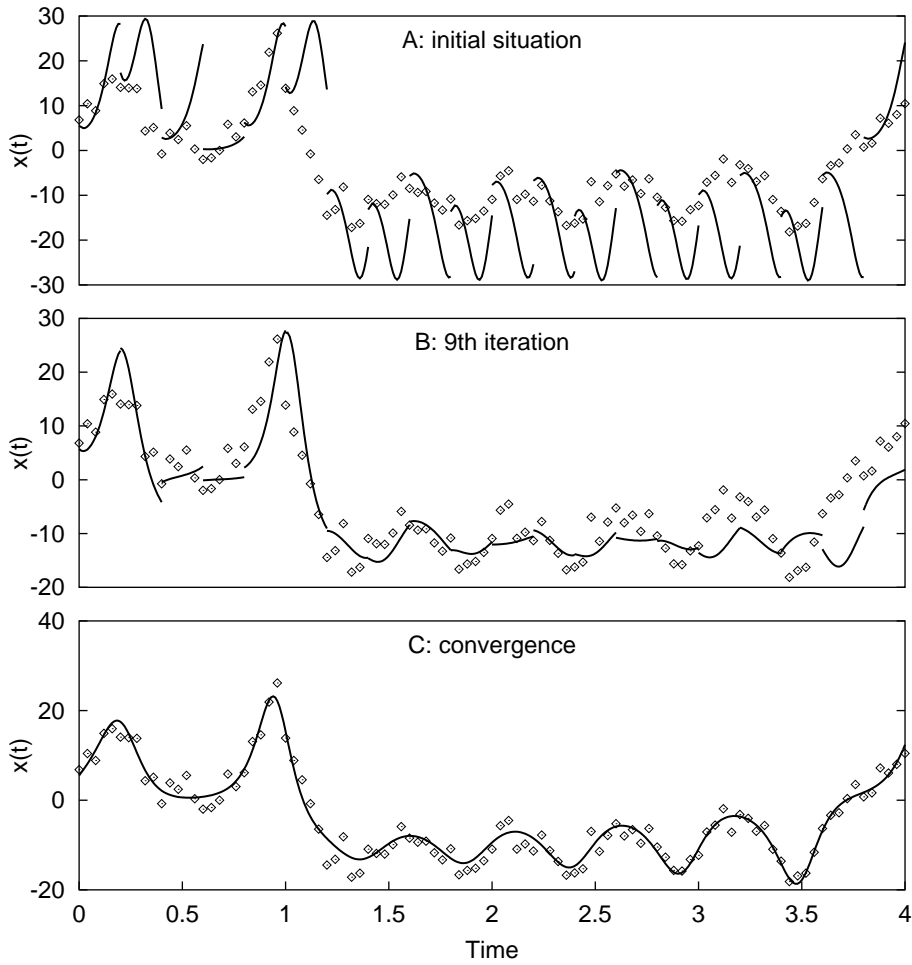


Figure 2.2: Convergence of the multiple shooting algorithm for the Lorenz system. The model trajectory is always near the data points, even when the parameters are rather wrong.

partitioned into M subintervals:

$$T_0 < T_1 < \dots < T_M = T_0 + T.$$

For each subinterval, *local initial values* \mathbf{x}_j^0 are introduced as additional parameters. The total parameter vector is then $\boldsymbol{\theta} = (\mathbf{p}, \mathbf{x}_0^0, \dots, \mathbf{x}_{M-1}^0)$. Let $\mathbf{x}_j(t, \boldsymbol{\theta})$ denote the solution of Eq. (2.6a) within $[T_j, T_{j+1})$ with initial values \mathbf{x}_j^0 . While the dynamic parameters \mathbf{p} are unique over the entire interval, the local initial values are optimised separately in each subinterval. The measurements are used to get starting guesses for them. This approach leads to an initially discontinuous trajectory, which is, however, close to the

measurements. The final trajectory must of course be continuous, i.e. the computed solution at the end of one subinterval must finally equal the local initial values of the next one:

$$\mathbf{x}_{j-1}(T_j, \boldsymbol{\theta}) = \mathbf{x}_j(T_j, \boldsymbol{\theta}) = \mathbf{x}_j^0, \quad j = 1, \dots, M - 1. \quad (2.8)$$

Eq. (2.8) represents a set of equality constraints like those in Eq. (2.4b). The special structure of their linearised form Eq. (2.5b) permits them and the extra variables \mathbf{x}_j^0 , $j = 1, \dots, M - 1$, to be eliminated easily from the resulting large linear system. In this way the dimension of the system of equations to be solved in each iteration is no larger than with the initial-value approach. This procedure is called *condensation* (Bock 1983). It can not be used when the quasi-Newton optimisation method is used for optimisation.

Since only the linearised continuity constraints are imposed on the update step, the iteration is allowed to proceed to the final continuous solution through “forbidden ground”: the iterates will generally be discontinuous trajectories. This freedom allows the method to stay close to the observed data, prevents divergence of the numerical solution and reduces the problem of local minima.

The best way to understand the procedure is to regard an example. The simulation from Fig. 2.1 was repeated, but now with the multiple shooting method (Fig. 2.2). It converges to the correct solution after 26 iterations. More details of the mathematical and implementational aspects of the method are given in Bock (1983, 1987). Other applications can be found in Horbelt et al. (2001b, 1998b); Timmer et al. (2000b); Baake et al. (1992); Baake and Schlöder (1992).

2.2.3 Calculating Sensitivities

The linearisation Eq. (2.5) requires not only the trajectory itself to be computed, but also its first derivatives with respect to the parameters. For ODEs, $\mathbf{x}(t_i, \boldsymbol{\theta})$ is not given explicitly, in contrast to non-linear parameter estimation in regression. The *sensitivity matrix* $\mathbf{S}(t_i, \boldsymbol{\theta}) := \frac{\partial \mathbf{x}(t_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$ can be calculated in different ways:

- External Numerical Differentiation (END)

This method is most simple because no analytical derivatives must be provided. For each parameter θ_j , the model equations are solved one more time with a slightly disturbed parameter value $\theta_j + h$. Then a finite difference yields the sensitivity:

$$\frac{\partial \mathbf{x}(t_i, \boldsymbol{\theta})}{\partial \theta_j} \approx \frac{\mathbf{x}(t_i, \boldsymbol{\theta} + h \mathbf{e}_j) - \mathbf{x}(t_i, \boldsymbol{\theta})}{h}, \quad (2.9)$$

where \mathbf{e}_j is the j -th unit vector. If the integration method uses adaptive stepsizes, the numerical solution of the ODE is discontinuous with

respect to the parameters. Then a small value of h can seriously destabilize the minimisation algorithm. This can be defused by using the same integration steps for the evaluation of $\mathbf{x}(t_i, \boldsymbol{\theta} + h\mathbf{e}_j)$ as for $\mathbf{x}(t_i, \boldsymbol{\theta})$. A further refinement is to use symmetric finite differences:

$$\frac{\partial \mathbf{x}(t_i, \boldsymbol{\theta})}{\partial \theta_j} \approx \frac{\mathbf{x}(t_i, \boldsymbol{\theta} + h\mathbf{e}_j) - \mathbf{x}(t_i, \boldsymbol{\theta} - h\mathbf{e}_j)}{2h}. \quad (2.10)$$

This increases the accuracy from the order of h to the order of h^2 at the cost of only doubling the computation time. Nevertheless, the finite truncation error that always exists with derivative-free methods becomes dominant at the convergence point, where the sensitivity of χ^2 with respect to $\boldsymbol{\theta}$ tends to zero by definition. Therefore the quadratic or at least super-linear convergence behaviour of sophisticated optimisation algorithms is destroyed by external numerical differentiation.

- Internal Numerical Differentiation (IND)

Given a common integration method and fixed or adaptively chosen stepsizes, the solution of an ODE, evaluated at a time t_i is the result of a finite number of arithmetic operations, called the integration scheme. The scheme can be differentiated analytically with respect to the parameters to yield the sensitivities (Bock 1983; Hairer et al. 1987). IND is very efficient, but it must be developed specifically for each integration method.

- Integration of the sensitivity equations

The time evolution of \mathbf{S} is given by

$$\begin{aligned} \dot{\mathbf{S}} &= \frac{d}{dt} \frac{\partial \mathbf{x}}{\partial \boldsymbol{\theta}} = \frac{\partial}{\partial \boldsymbol{\theta}} \dot{\mathbf{x}} = \frac{d}{d\boldsymbol{\theta}} \mathbf{f}(t, \mathbf{x}(t, \boldsymbol{\theta}), \mathbf{p}) \\ &= \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \boldsymbol{\theta}} + \left(\frac{\partial \mathbf{f}}{\partial \mathbf{p}}; \mathbf{0} \right) \\ &= \mathbf{J}\mathbf{S} + (\mathbf{J}_p; \mathbf{0}), \end{aligned} \quad (2.11)$$

where

$$\begin{aligned} \mathbf{J} &= \frac{\partial \mathbf{f}(t, \mathbf{x}, \mathbf{p})}{\partial \mathbf{x}} \\ \mathbf{J}_p &= \frac{\partial \mathbf{f}(t, \mathbf{x}, \mathbf{p})}{\partial \mathbf{p}} \end{aligned}$$

are the *Jacobian matrix* and the *inhomogeneity matrix* of the ODE respectively. $\mathbf{0}$ is the $m \times m$ null matrix expressing that $\frac{\partial \mathbf{f}(t, \mathbf{x}, \mathbf{p})}{\partial \mathbf{x}_j^0}$ is zero.

These ODEs are called *sensitivity equations* or *variational equations*. The fact that they belong to the same class of dynamical systems as the

original equations, is a principle that holds for all classes enumerated in Fig. 1.1. The right-hand side of the sensitivity equations depends on the solution of the original equation (2.6a). Therefore Eqs. (2.6a) and (2.11) are combined to an $m(1+p+m)$ -dimensional ODE that is solved with standard variable step size integration methods. In order to increase the efficiency of the integration, the error control can be modified to regard only the original equation. This is usually sufficient because errors in the sensitivities affect only the rate of convergence and not the accuracy of the result. With this modification the method achieves the same efficiency as IND, while it can be used with virtually every integration method.

The latter method was used in the examples reported here.

2.3 Time-discrete Systems

Since measurements of a dynamical process are usually sampled with equidistant times, it is an obvious step to look at the process also only on a time-discrete level. An advantage of maps is their simplicity in comparison with ODEs. They can be used to test methods for estimating parameters and to investigate fundamental mechanisms.

Time-discrete dynamical systems are mostly autonomous, i.e. do not explicitly depend on time. They are defined by a map function

$$\mathbf{x}_i = \mathbf{f}(\mathbf{x}_{i-1}). \quad (2.12)$$

The right-hand side \mathbf{f} could be defined by the propagation of an ODE from a time t to $t+\Delta t$ with initial values \mathbf{x}_{i-1} , for instance. Another way to convert an ODE into a map is through the technique of *Poincaré section*. For this method the t_i are defined as the times at which the state vector crosses a given $(n-1)$ -dimensional manifold in the state space (e.g. a hyper-plane) in a given direction. This technique is very popular in the study of chaotic systems since it allows the reduction of the state space by one dimension while characteristic properties of the system (e.g. the non-trivial *Lyapunov exponents*) are retained.

Having developed the full apparatus for estimating parameters in ODEs, it is only a short path to transfer this approach to maps. Since maps often exhibit very complex dynamics, the problem of local minima is expected to be relevant in this context. The multiple shooting technique will turn out to be essential for maps. The groups of K points $\{t_{mK}, \dots, t_{mK+K-1}\}$ is combined to the m th “multiple shooting interval”. The extended vector of parameters is $\boldsymbol{\theta} = (\mathbf{p}, \mathbf{x}_0, \mathbf{x}_K, \dots, \mathbf{x}_{(M-1)K})$, where M is the number of multiple shooting intervals. The subsequent states within the m th interval fulfil the dynamical equation

$$\mathbf{x}_i = \mathbf{f}(t_i, \mathbf{x}_{i-1}, \mathbf{p}), \quad i = mK + 1, \dots, mK + K - 1. \quad (2.13)$$

The continuity between the intervals, which is the same equation for the remaining points $i = mK$, $m = 1, \dots, M - 1$, is realized via equality constraints. These may be violated during the iterative process, but they are satisfied at the convergence point. For the solution of the optimisation problem, a quasi-Newton method is used, since the stability of the Gauss-Newton method is too delicate for the extreme nonlinearities existing in many maps. The initial-value approach corresponds to setting $M = 1$ and K equal to the number of data points.

The sensitivities $\mathbf{S}_i(\boldsymbol{\theta}) := \frac{\partial \mathbf{x}_i}{\partial \boldsymbol{\theta}}$ are computed from the map

$$\mathbf{S}_i = \mathbf{J}\mathbf{S} + (\mathbf{J}_p; 0), \quad (2.14)$$

where

$$\mathbf{J} = \frac{\partial \mathbf{f}(t_i, \mathbf{x}_{i-1}, \mathbf{p})}{\partial \mathbf{x}}$$

$$\mathbf{J}_p = \frac{\partial \mathbf{f}(t_i, \mathbf{x}_{i-1}, \mathbf{p})}{\partial \mathbf{p}}$$

analogously to Eq. (2.11) for ODEs.

2.4 Delay Differential Equations

Delay Differential Equations (DDEs) arise from ODEs when the right-hand side depends on the state at one or more *retarded* times. Examples of such systems are used to model processes in physics, biology, medicine and engineering science, among others.

The uniqueness of the solution of a DDE requires the specification of the *initial curve* within an interval as long as the largest delay time involved. In this way an *initial curve problem* is posed, analogously to the initial value problem of an ODE. As the initial curve is a vector in a function space, the state space of a DDE is infinite-dimensional. When it is unknown, the question of identifiability arises. Only a finite number of degrees of freedom can be estimated from measured data.

For ease of notation this presentation is restricted to the most frequent case of a single time lag τ and a single, namely the first component of the state vector entering into the dynamics at the delayed time. Then the initial curve problem reads

$$\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}, x_\tau, \mathbf{p}) \quad \text{for } t > T_0 + \tau \quad (2.15a)$$

$$x_\tau = x_1(t - \tau)$$

$$x_1(t) = h_0(t) \quad \text{for } t \in [T_0; T_0 + \tau]. \quad (2.15b)$$

2.4.1 Direct Problem

The solution of Eq. (2.15) can be carried out in steps of τ (*method of steps*)¹: Let I_k denote the interval $[T_0 + k\tau; T_0 + (k+1)\tau]$. In I_k , $x_1(t - \tau)$ is given by the solution on the preceding interval I_{k-1} . Thus the DDE Eq. (2.15a) can be solved as an ODE with a known explicit time dependence. In general, the first derivative of the solution has a discontinuity at $T_0 + \tau$ which is recurring in the second derivative at $T_0 + 2\tau$ and so on. This should be taken into account in the numerical integration.

ODE integrators with variable step size selection, which are favourable, require the specification of the right-hand side at arbitrary points in time. In the DDE case this means that a continuous solution must be computed for a point in time that is not known beforehand. This is not possible with the usual routines that integrate from point to point and return the trajectory only at the points requested.

The idea of a Runge-Kutta algorithm is to approximate the solution by piecewise polynomials. If the coefficients of these are stored in addition to the function values at each integration step, the solution can be interpolated afterwards. This technique, called *dense output*, was already used by the ODE integrator DOPRI5, in order to circumvent the maximum length of the integration steps being limited by the sampling interval. The code RETARD is an extension of DOPRI5. It supplies a routine computing the lagged variables, that can be called within the calculation of the right-hand side of the DDE. Both codes are made public in Hairer et al. (1987).

2.4.2 Calculating Sensitivities

Analogously to Eq. (2.11) for ODEs, the sensitivities $\mathbf{S} = \frac{\partial \mathbf{x}}{\partial \boldsymbol{\theta}}$ of the state vector with respect to the parameters fulfil the DDE

$$\dot{\mathbf{S}} = \mathbf{J}\mathbf{S} + \mathbf{J}'\mathbf{S}_\tau + (\mathbf{J}_p; \mathbf{0}), \quad (2.16)$$

where

$$\begin{aligned} \mathbf{J} &= \frac{\partial \mathbf{f}}{\partial \mathbf{x}}, & \mathbf{J}' &= \frac{\partial \mathbf{f}}{\partial x_\tau}, \\ \mathbf{J}_p &= \frac{\partial \mathbf{f}}{\partial \mathbf{p}}, & \mathbf{S}_\tau &= \frac{\partial x_1}{\partial \boldsymbol{\theta}}(t - \tau), \end{aligned}$$

An exception is the delay parameter τ , being part of \mathbf{p} when the time lag is unknown. The right-hand side $\mathbf{f}(t, \mathbf{x}(t, \mathbf{p}), x_1(t - \tau, \mathbf{p}), \mathbf{p})$ depends on τ through the retarded time $t - \tau$, in addition to the three dependencies through \mathbf{p} already contributing to Eq. (2.16). Therefore the time evolution

¹ If multiple time lags exist, τ is the smallest of them.

of $\frac{\partial \mathbf{x}}{\partial \tau}$ contains an extra term:

$$\frac{d}{dt} \frac{\partial \mathbf{x}}{\partial \tau} = \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \tau} \Big|_t + \frac{\partial \mathbf{f}}{\partial x_\tau} \left(\frac{\partial x_1}{\partial \tau} \Big|_{t-\tau} - \frac{\partial x_1}{\partial t} \Big|_{t-\tau} \right) + \frac{\partial \mathbf{f}}{\partial \tau}. \quad (2.17)$$

The extra term is given by

$$\frac{\partial x_1}{\partial t}(t - \tau) = \begin{cases} \dot{h}_0(t - \tau) & t \leq T_0 + 2\tau \\ f_1(t - \tau, \mathbf{x}(t - \tau), x_1(t - 2\tau), \mathbf{p}) & t > T_0 + 2\tau. \end{cases} \quad (2.18)$$

2.4.3 Multiple Shooting

Many DDEs exhibit irregular oscillations due to the infinite-dimensional phase space. Therefore a technique similar to the multiple shooting approach is expected to be particularly helpful for the estimation of parameters in DDEs. However, there is no straightforward approach to express the continuity constraints which are an essential element of this technique.

As in the case of ODEs, the fitting interval can be partitioned into subintervals.

Each subinterval j has its own initial curve h_j of length τ , involving infinitely many degrees of freedom. The purpose of the continuity constraints is to ensure that the final trajectory is a solution of the DDE, i.e. the initial curves of subsequent subintervals must each be consistent with the trajectory on the preceding intervals. Therefore the subintervals must have an overlap of length τ with each other. Let g_j define the segment of the $(j - 1)$ th subinterval that overlaps with the j th. Continuity would require g_j and h_j to be exactly equal. This can not be achieved since h_j must somehow be represented by a finite number of parameters p_s and g_j will in general not be a member of the corresponding family of functions. In other words, while g_j is an arbitrary vector in a function space, h_j lies on a p_s -dimensional sub-manifold of that function space. Two methodically different ways to formulate the matching conditions between the two functions are conceivable:

1. $\|g_j - h_j\|$ is minimised, where $\|\cdot\|$ is an appropriately defined norm in the function space.
2. g_j is projected onto the manifold of initial curves and the corresponding p_s parameters are forced to be equal to those of h_j .

The first approach requires two cost functions to be minimised simultaneously, i.e. the mismatch norm must be included as a penalty term in the cost function Eq. (2.4a), scaled with a penalty parameter whose appropriate size is rather difficult to determine. Therefore the second approach is pursued in the following procedure. The parameterisation of the initial curves will be performed via *cubic splines*.




Figure 2.3: Schematic diagram of the multiple shooting intervals and their spline and data segments. **Bold lines:** spline segments. **Thin lines:** data segments. **Open circles:** spline knots.

Estimation procedure

1. Define a grid of points $\{T_0 + kd, k = 0, 1, \dots\}$ with d being chosen appropriately. These points will serve as the knots of cubic splines representing the initial curves. They should be dense enough to allow for a sufficient approximation of the model trajectory. If the measurements are sampled equidistantly, the data points or a subset of them should be used for the knots.
2. Estimate the smallest K such that Kd is equal to or greater than the time lag: $(K - 1)d < \tau \leq Kd$. The splines will be defined on intervals of length Kd . No problem will arise when τ is smaller than Kd (then only that part of the spline will be used as the initial curve) or only a little larger (then the spline will be extrapolated). When τ is much larger than Kd , the spline will give a poor approximation of the initial curve. Therefore, if τ is unknown, it is better to start with a large K .
3. Choose the separation D of the multiple shooting mesh points in multiples of d and define $T_j = T_0 + jD$, $j = 1, \dots, M$. Each subinterval consists of a *spline segment* of length Kd and a *data segment* of length D as depicted in Fig. 2.3. The best choice of D depends on the regularity of the time series and on the quality of the starting guesses for the parameters, as in the case of ODEs. A typical choice is $D = Kd$, in which case the end of each spline segment coincides with the beginning of the following one. $D < Kd$ generates multiply overlapping intervals and will rarely be meaningful.
4. Introduce *spline variables* $\mathbf{s}_j = (\dot{s}_{j0}, s_{j0}, \dots, s_{jK}, \dot{s}_{jK})$ for each spline segment $[T_j; T_j + Kd]$. The cubic spline h_j is defined as a piecewise cubic polynomials function, being continuous up to second order and satisfying

$$h_j(T_j + kd) = s_{jk}, \quad k = 0, \dots, K \quad (2.19)$$

$$\dot{h}_j(T_j + kd) = \dot{s}_{jk}, \quad k = 0, K. \quad (2.20)$$

For a detailed discussion of splines, see Press et al. (1992). They are computed in two steps. In the setup step, $\ddot{s}_{jk} = \ddot{h}_j(T_j + kd)$, $k = 0, \dots, K$ are calculated from \mathbf{s}_j . The \ddot{s}_{jk} are linear in the \mathbf{s}_j :

$$\ddot{s}_{jk} = \mathbf{L}\mathbf{s}_j.$$

In the evaluation step, $h_j(t)$ is evaluated at arbitrary times t , using s_{j0}, \dots, s_{jK} and $\ddot{s}_{j0}, \dots, \ddot{s}_{jL}$. Since the \mathbf{s}_j will be fit variables, sensitivities of the spline functions with respect to their respective spline variables are also needed. For their calculation the explicit knowledge of \mathbf{L} is necessary.

5. The initial curve is the spline function h_j , restricted to the interval $[T_j + Kd - \tau; T_j + Kd]$, where τ is the actual estimate of the delay parameter. Note that the end of the initial curve matches the end of the spline segment, even when τ varies. Eqs. (2.15a) and (2.16) are integrated on the data segment. The residuals between the model trajectory and the data points contribute to the cost function. The first data segment is extended by its respective spline segment in order to include the corresponding data points in the cost function.

For each but the last trajectory, g_{j+1} is defined by the part overlapping with the spline segment of the following interval. It is projected onto the spline manifold $\{\mathbf{s}_{j+1}\} \subset \mathbb{R}^{K+3}$ by simply reading out the function values and derivatives

$$r_{j+1,k} = g_{j+1}(T_{j+1} + kd), \quad k = 0, \dots, K \quad (2.21a)$$

$$\dot{r}_{j+1,k} = \dot{g}_{j+1}(T_{j+1} + kd), \quad k = 0, K. \quad (2.21b)$$

Now the continuity constraints read

$$\mathbf{r}_j = (\dot{r}_{j0}, r_{j0}, \dots, r_{jK}, \dot{r}_{jK}) = \mathbf{s}_j, \quad j = 1, \dots, M. \quad (2.22)$$

As in the case of ODEs, the number of continuity constraints is equal to the number of variables newly introduced in each subinterval. For the lagged variable, the ordinary continuity constraints Eq. (2.8) are disregarded since Eq. (2.19) implies $x_1(T_j + Kd) = s_{jK}$. For the other variables, if there are any, Eq. (2.8) is used in addition to Eq. (2.22).

6. During the first iterations of the optimisation, the spline variables are held fix because they are expected to be estimated well from the data. This is important when the starting guesses for the other fit variables are far from the true values. After the algorithm has converged for the first time, they are released and fitted together with the other variables. (*two-phase procedure*)

Choice of Control Parameters

While Kd is required to approximate the time lag, K itself and D are important control parameters for the fit. Their optimal choice must be made with regard to the problem at hand. When long time series with many irregular oscillations are examined, many subintervals are needed in order to achieve global convergence, especially if good starting guesses of the parameters are not available. After convergence the algorithm can be restarted with improved starting guesses and without multiple shooting in order to be perfectly sure that the projection Eq. (2.21) does not distort the analysis.

K controls the accuracy of the splines. If it is large, the question of over-fitting must be considered. For the second and subsequent subintervals, this is not an issue since the additional degrees of freedom introduced by the spline variables are counterbalanced by an equal number of continuity constraints.

For the initial curve it is an issue. However, the influence of the initial curve on the dynamics is that of an external input. High frequency portions are damped by the integration. Thus, the DDE behaves like an intrinsic low-pass filter. Therefore chaotic delay systems have often low-dimensional attractors that are attained after a short transient period. All but a few degrees of freedom of the initial curve decay rapidly. As a consequence the largest part of the trajectory is not sensitive to these degrees of freedom. In the corresponding directions in the space of fit variables, $\chi^2(\boldsymbol{\theta})$ is a slowly varying function. These directions are linear combinations of the variables of the first spline segment, so these spline variables are delicate to noise and their estimates will have large confidence intervals. This could be termed *over-fitting* if the construction of the beginning of the model trajectory was a major aim. Yet the main interest will generally be an accurate estimate of the model parameters and this aim is achieved by means of the remaining part of the trajectory.

2.5 Other classes

A few remarks shall be made concerning those classes that can not be treated in detail here.

Differential-algebraic equations (DAEs) are ODEs in which one or more differential equations are replaced by algebraic relations between state variables. Such systems occur, for instance, in the modelling of chemical reaction systems and in Hamiltonian systems, in which constants of motion are conserved.

DAEs with simple algebraic portions can be translated into pure ODEs by eliminating individual variables with the help of the algebraic equations. More complex DAEs can be treated in a similar way but the resulting ODEs

may become excessively complicated. In this case it would be preferable to use integration methods that are explicitly developed for DAEs (Brenan et al. 1989).

Integro-differential equations (IDEs) are a generalisation of DDEs in which the delay time is not a fixed value. It rather follows some statistical distribution. For biological systems this scenario is often more realistic than the assumption of fixed delays. Mathematically it is formulated by means of an integral term in the evolution equations:

$$\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}, \mathbf{x}_\tau, \boldsymbol{\theta}), \quad (2.23)$$

with

$$\mathbf{x}_\tau(t) = \int_{-\infty}^t \mathbf{g}(t, t', \mathbf{x}(t')) dt'. \quad (2.24)$$

If the time delay has a lower bound $\tau > 0$, i.e. , $\mathbf{g}(t, t', \mathbf{x}) = 0$ for $t-t' < \tau$, the methods of Section 2.4 can be applied, where for each evaluation of \mathbf{f} , an integral and its sensitivities have to be computed numerically. Since IDEs are promising candidates for models in many practical applications, future work will be devoted to this technique.

Finally, methods for estimating parameters in **partial differential equations** (PDEs) can be found in Schittkowski (1999); Bär et al. (1999); Kelpin et al. (2000); Coca and Billings (2000).

Chapter 3

Time-discrete Systems

In this chapter, the maximum likelihood methods described in chapter 2 will be tested on the most prominent deterministic time-discrete system, the logistic map. Some interesting properties and subtleties of the parameter estimation task can be studied more easily on such a simple system. For instance, the positive Lyapunov exponent of the map will play an interesting role in the following. A three-dimensional ODE would be required to study equivalent properties for a continuous systems.

3.1 Introduction

The map function is written in the form

$$x_i = f_a(x_{i-1}) = 1 - ax_{i-1}^2, \quad i = 1, \dots, N \quad (3.1)$$

with the parameter $a \in (0; 2]$ and the state $x_i \in [1 - a; 1]$. The observation equation reads

$$y_i = x_i + \eta_i, \quad i = 0, \dots, N, \quad (3.2)$$

where η_i denotes Gaussian white noise with variance σ^2 . Readers preferring the equivalent formulation $z_i = rz_{i-1}(1 - z_{i-1})$, are referred to Appendix A for a description of the transformation. In analogy to ODEs, a sequence of states (x_0, x_1, \dots) satisfying Eq. (3.1) shall be termed a trajectory.

A simple method for estimating a is to regard the measured values y_i as functions of their precedents y_{i-1} and to minimise the *one-step prediction error*

$$\chi^2(a) = \sum_{i=1}^N \frac{(y_i - f(y_{i-1}))^2}{\sigma^2}. \quad (3.3)$$

This cost function ignores three important facts:

- The probability of a system state to take a certain value x is not the same for all x in the range of the map. The x_i are rather distributed ac-

ording to the *invariant measure* of the map. This probability density depends on the parameter.

- The y_{i-1} in the numerator are treated as if they were noiseless “independent variables”. Actually they are as noisy as the “dependent variables” y_i . This is called the *errors-in-variables problem* (Carroll et al. 1995; Kantz and Jaeger 1997). It can be accounted for by minimising the *orthogonal distances* between pairs of data points (y_{i-1}, y_i) and the graph of the map function. This is called the *total least squares method*.
- Two pairs (y_{i-1}, y_i) and (y_{j-1}, y_j) are not independent of each other, they are related through the dynamical equation (3.1).

The first two points lead to biased estimates \hat{a} , even in the limit of infinitely many data points. When $a=a_{\text{true}}=2$, the expectation value of \hat{a} evaluates analytically to

$$\langle \hat{a} \rangle = \frac{4\sigma^2 + 3}{24\sigma^4 + 24\sigma^2 + 3} a_{\text{true}},$$

which is only 1.1 for a signal-to-noise ratio of 2.

McSharry and Smith (1999) showed that the total least squares method is not sufficient to remove the bias. They developed a method that yields unbiased estimates of a by taking into account the first two points listed above. The attribute “maximum likelihood” they attached to their method is misleading since a real maximum likelihood method must not ignore that one true trajectory is underlying the measurements. The third point means that information about the data is ignored. As a consequence the variance of \hat{a} can not attain its lowest possible value, the *Cramér-Rao bound* (Cox and Hinkley 1994).

Jaeger and Kantz (1996) proposed *orthogonal distances* to correct for the errors-in-variables and the *n-step prediction error* to take into account the dynamic nature of the data, at least over a few iterations of the map function. The n-step prediction error is based on short trajectories that are locally fitted to the measurements. Due to numerical difficulties, they could not extend their methods beyond $n=4$ iterations though they recognised that this would be desirable.

The maximum likelihood method used here corresponds to choosing $n=N$ in the context of the *n-step-prediction error*: it finds the trajectory satisfying the dynamical equations over the entire interval and coming closest to the data. This is sometimes termed *solving the shadowing problem*. Having solved this problem, the only independent variable left is the very first one, the initial value x_0 . In principle the first point in the list above applies to this initial value. However, this is expected to have only a boundary effect that has a negligible impact on the parameter estimate.

Since much more information is taken into account in this way, the estimate is expected to have a much lower variance. The difficulty is that the parameter space is cluttered with local minima corresponding to suboptimal trajectories. The multiple shooting approach is a technique to find the best fit trajectory without getting trapped too easily in local minima. In the following sections the method is applied to simulated data of the logistic map.

3.2 Logistic Map without Noise

At first noiseless data are used in order to separate difficulties that are due to noise from those stemming from bad starting guesses. They were simulated according to Eq. (3.1) with $a=1.85$, $x_0=0.8$ and $N=10$. The parameter a was estimated with the initial-value approach and the multiple shooting approach (Fig. 3.1). For the latter, the number K of data points per

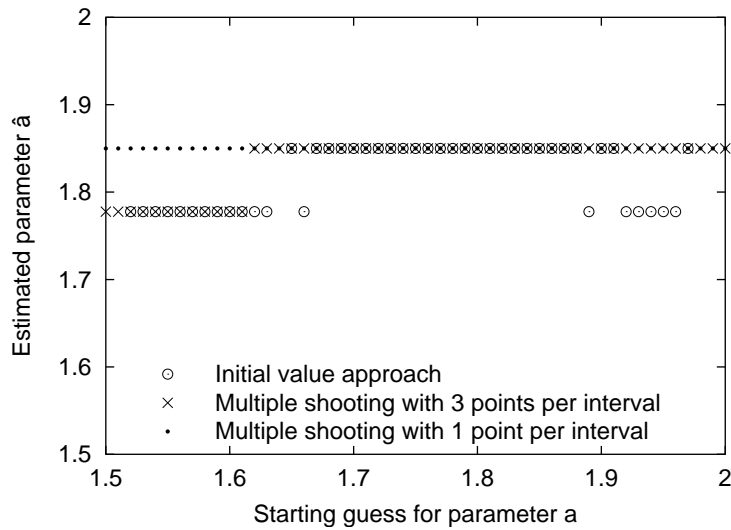


Figure 3.1: Estimating the parameter a in the logistic map from noiseless data with the starting guess varying from 1.5 to 2 in steps of 0.01. True parameter $a=1.85$, initial value $x_0=0.8$, number of data points $N + 1=11$. **Circles:** results for the initial-value approach. The true parameter is found only when the starting guess is already near the true value. All false solutions are located on a single line with $a=1.778$, indicating a clear local minimum. Some circles are missing: for these starting guesses the algorithm did not converge at all. **Crosses:** results for the multiple shooting approach with $K=3$ data points per multiple shooting interval. The minimisation is less susceptible, yet not immune against local minima. **Points:** results for the multiple shooting approach with $K=1$. The parameter is estimated correctly within the entire interval.

multiple shooting interval (see page 14) was set to 1 and to 3 respectively.

The starting guess a_0 for the unknown parameter a was varied from 1.5 to 2 in steps of 0.01. The starting guesses for the initial values were taken from the data, i.e. they were the true values. While the multiple shooting method with $K=1$ converged to the correct value for all starting guesses, the initial-value approach stopped in a local minimum when a_0 was below 1.6 or above 1.9. With $K=3$ the result was better than with the initial-value approach but it was worse than with $K=1$.

The simulation was repeated with $N=20$. Again the multiple shooting algorithm supplied accurate estimates. This time the initial-value approach failed when the deviation of the starting guess was larger than $2 \cdot 10^{-4}$. It is a general rule that the unreliability of the initial value approach increases with the length of the trajectory.

This comparison clearly demonstrates the necessity of the multiple shooting approach for estimating parameters in chaotic maps. All following simulations in this chapter were done with one data point per multiple shooting interval. Nevertheless the multiple shooting method is not necessarily a guarantee against effects of the starting guess. In fact local minima were




Figure 3.2: Objective function for noiseless data of different lengths N . Initial value $x_0=0.8$, true parameter $a=1.85$. Curves for higher N are higher than those for smaller N since the summands of Eq. (2.3) are non-negative. The global minimum at $a=1.85$ becomes successively sharper but the number of local minima increases.




Figure 3.3: Same as Fig. 3.2 for fixed $a=1.85$ and varying x_0 .

encountered when the starting guess was below 1.15.

Fig. 3.1 shows that the reason for the lack of reliability must be sought in local minima of the χ^2 landscape. When a false solution was found, it was always $(a, x_0) = (1.778, 0.820)$. To get an impression of the mechanism, the objective function Eq. (2.3) is regarded. $\chi^2(a, x_0)$ was computed for noiseless data of different lengths N as a function of a for a fixed initial value $x_0=0.8$ (Fig. 3.2), and also as a function of x_0 for fixed $a=1.85$ (Fig. 3.3).

With each additional data point included in the cost function, the global minimum at $(a, x_0) = (1.85, 0.8)$ becomes sharper, so the unknown variables can be estimated more accurately. But at the same time, the number of local minima in which the iterative process could stop unintentionally, increases. This phenomenon is most distinct when χ^2 is plotted against x_0 .

From Figs. (3.2) and (3.3) it follows that additional noise has a severe impact on the parameter estimates. Firstly, the starting guesses for the states, which are taken from the data, are contaminated and can lead into a false minimum. Secondly, the shape of the objective function is deformed so that a false local minimum can become the global minimum. For instance, in Fig. 3.2, there is a local minimum at $a=1.716$ for $N=10$. It is so low that moderate noise on the first few data points could lift the global minimum above it, having the maximum likelihood estimate jump to 1.716. Before

studying the results for the logistic map with noisy observations, a digression to the *shift map* is made, in order to understand the causalities better.

3.3 Shift Map

The shift map is related to the tent map which is turn similar to the logistic map for $a=2$ (see Appendix A). The definition of the shift map reads:

$$x_i = f(x_{i-1}) = \begin{cases} 2x_{i-1} & 0 \leq x_{i-1} < 0.5 \\ 2x_{i-1} - 1 & 0.5 \leq x_{i-1} < 1. \end{cases} \quad (3.4)$$

This simplest example of a map having chaotic properties shows the essential ingredients of a chaotic system: stretching (slope=2 at each point) and folding: $f(x + 0.5) = f(x)$. With regard to the parameter estimation problem, the initial value x_0 is the only unknown parameter to be estimated from noisy measurements.

The n th iterate of the shift map can be written

$$f^n(x_0) = x_n = 2^n x_0 \pmod{1},$$

where $x \pmod{1}$ denotes the fractional part of x . That means that f^n is a heavily oscillating function for large n and this also shows up in the objective function. The name shift map originates from its binary representation. When the system state is written in the form

$$x_i = \sum_{j=1}^{\infty} a_{ij} 2^{-j}, \quad a_{ij} \in \{0, 1\},$$

the map function for the a_{ij} reads:

$$a_{i,j} = a_{i-1,j+1},$$

i.e. , the sequence of bits is shifted one bit to the left by the map. This instructive representation leads immediately to the following conclusions:

1. An infinitesimal separation ϵ between two initial states x_0 and x'_0 amplifies to $2^n \epsilon$ after n iterations, thus the Lyapunov exponent of the shift map is

$$\lambda = \log(2)$$

2. When the map is iterated using the usual double precision arithmetic with a 52 bit mantissa, every state beyond the 52nd iteration is zero or pure artifact. This limit shall be called the *numerical horizon*.

3. If the noise on each data point is small ($\sigma \ll 0.5$), subsequent measurements can be used to improve the accuracy of the result by a factor of 2 each, thus the standard deviation of the estimated initial value decreases exponentially with the number of observed states:

$$\sigma_{x_0} = 2^{-N} \sigma = e^{-\lambda N} \sigma. \quad (3.5)$$

On the contrary, estimation methods that neglect the existence of an underlying trajectory, yield estimates whose standard deviation scales as $1/\sqrt{N}$ according to the law of large numbers.

The exponential scaling is based on the exponential divergence of trajectories which is characterised by the positive Lyapunov exponent. It is therefore expected to be inherent to all chaotic systems. In particular, for the logistic map, Figs. (3.2) and (3.3) show that the minima of χ^2 become exponentially sharper with increasing N .

4. If, on the other hand, σ is comparable with 0.5, the most significant bit of x_0 , a_{01} , can not be estimated reliably because y_0 is noisy and y_1, y_2, \dots do not depend on this bit. That means that the accuracy of the result is limited by a lower bound that is independent of N . To quantify this bound, assume that $x_0 < 0.5$, i.e. $a_{01} = 0$ without loss of generality and that all data points but the very first are measured with negligible noise: $\sigma_0 = \sigma$, $\sigma_1 = \sigma_2 = \dots = 0$. The subsequent points y_1, y_2, \dots determine \hat{x}_0 to be either x_0 or $x_0 + 0.5$. When $y_0 < x_0 + 0.25$, the correct value is estimated, otherwise $x_0 + 0.5$ is estimated and the error $\Delta = \hat{x}_0 - x_0$ is 0.5. Thus the mean squared error of the estimate is

$$\langle \Delta^2 \rangle = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{1}{2} \frac{\eta^2}{\sigma^2}\right) \Delta^2 d\eta \quad (3.6)$$

$$= \frac{1}{4} \int_{0.25/\sigma}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} z^2\right) dz. \quad (3.7)$$

The exponential law Eq. (3.5) breaks down when $\sigma_{x_0}^2$ arrives at this bound.

3.4 Logistic Map with Noise

The analytic results for the shift map raise the hope that the high sensitivity to initial values, while causing subtleties in the task of convergence to the global minimum, can also be helpful with regard to the accuracy of the estimate. The Lyapunov exponent of the logistic map for $a=1.85$ is $\lambda=0.503$. The corresponding numerical horizon is at the 71st iterate. Since it is not clear whether results for such long time series can be taken seriously, the following simulations were made with a maximum length $N=60$.




Figure 3.4: Parameter a estimated from noisy time series of length $N=20$ as a function of the noise level. For each noise level 100 realizations were made. **Plus signs and error bars:** ensemble averages and standard deviations of the mean, calculated from 100 realizations. The standard deviation of the estimate is ten-times higher than the standard error of the mean. **Broken line:** true parameter. For high noise levels, the estimate has large bias and variance. This is a finite sampling effect, i.e. the bias vanishes with increasing N .




Figure 3.5: Standard deviation of \hat{a} as a function of the noise level in double-logarithmic plot. **Crosses:** standard deviation σ_a of the estimated parameter \hat{a} . **Solid line:** Linear function of σ for comparison of the slopes. σ_a is proportional to σ for moderate noise levels.

Data were simulated with $a=1.85$, $x_0=0.8$ and $N=20$ and observation noise with *noise levels* between 0.1% and 100% was added. The noise level is defined as the ratio between the standard deviations of the noise and the true signal. For each noise level, 100 realizations of the time series were generated and the parameters were estimated from them. To suppress any effects of the starting guess, the true values for all fit variables were supplied as starting guesses.

Fig. 3.4 shows mean and standard deviation of the estimated parameter \hat{a} . The true parameter always lies within the confidence intervals calculated from 100 realizations. However, the standard error of the mean is ten-times smaller than the confidence intervals. It indicates a significant bias for high noise levels that is due to the finite sampling and decreases with increasing N . In Fig. 3.5 the standard deviation σ_a of \hat{a} is plotted double-logarithmically against the noise level. As expected, σ_a is proportional to σ for moderate noise levels.

Next the validity of Eq. (3.5) is examined for the case of the logistic map, which has a Lyapunov exponent $\lambda=0.503$ for $a=1.85$. For this purpose the same simulations as above were done, but now with fixed noise levels of 0.1%, 1% and 10% and N varying from 3 to 60. For each combination, 1000 trajectories were simulated. In order to avoid artifacts in the N -dependence, the initial value x_0 was drawn randomly according to the invariant measure

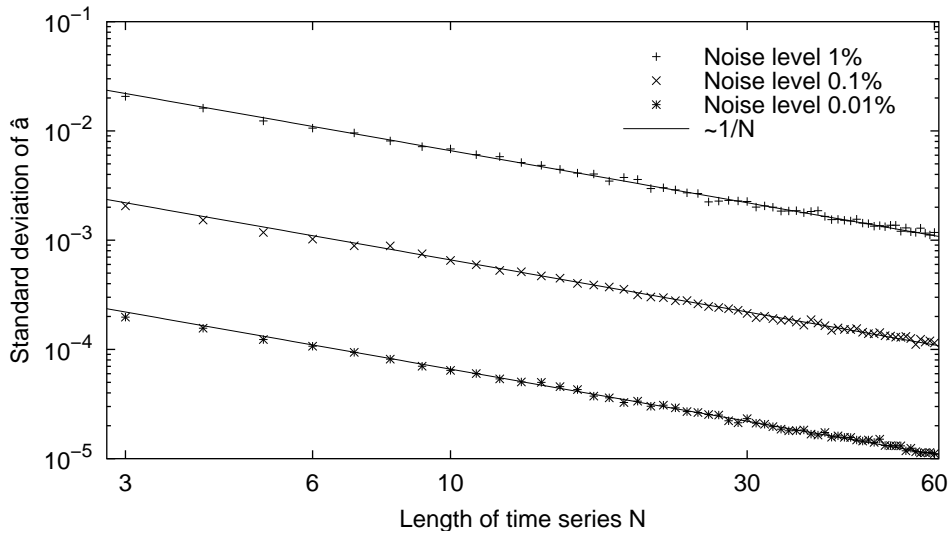


Figure 3.6: Standard deviation of the estimate as a function of the length N in double logarithmic plot. **Symbols:** standard deviation of \hat{a} , calculated from 1000 realizations each, for the indicated noise level. **Solid lines:** reciprocal of N with appropriate factors. The exponential law Eq. (3.5) with $\lambda=0.503$ is not fulfilled. Instead the standard deviation is inversely proportional to N .




Figure 3.7: Same as Fig. 3.6, but with one parameter being fixed to the true value. Noise level: 10%. **Plus signs:** Standard deviation of estimated a when x_0 was fixed. **Crosses:** Standard deviation of estimated x_0 when a was fixed. **Solid line:** Expected exponential function $e^{-\lambda N}$, where $\lambda=0.503$ is the Lyapunov exponent of the logistic map for $a=1.85$.

of the map. Then the parameters were estimated from each trajectory and the standard deviation was calculated (Fig. 3.6). Again the algorithm was supplied with optimal starting guesses. Nevertheless, no indication of an exponential scaling law can be found. Instead the standard deviation seems to scale perfectly as $1/N$.

Anomalous scaling behaviour has been reported for dimension estimates (Theiler 1990) and for Lyapunov exponent estimates (Theiler and Smith 1995), as exceptions from the ordinary $1/\sqrt{N}$ law. In the present case it is not a surprise that the precision scales better than $1/\sqrt{N}$ since the method takes into account more information than methods that treat subsequent pairs of data as independent.

The reason for the non-exponential scaling in contrary to the shift map is, that two variables were estimated simultaneously while the shift map has only one parameter at all. In order to verify this, the same simulations were made with a or x_0 being fixed to their true values respectively (Fig. 3.7). The true value of x_0 was 0.8. This time the exponential law is clearly visible and covers 12 orders of magnitude.

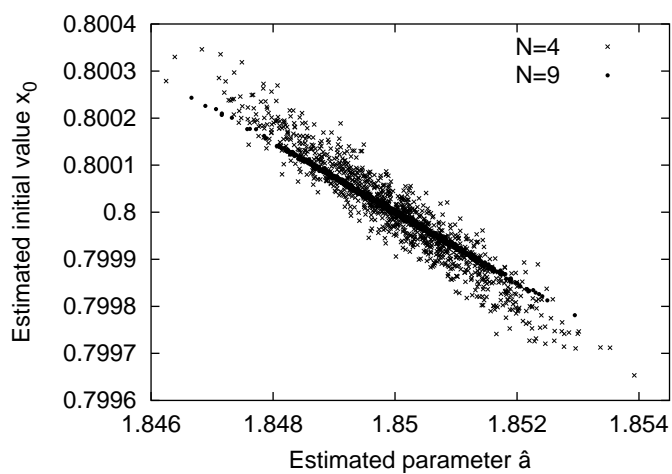


Figure 3.8: Correlation between the estimated parameters \hat{a} and \hat{x}_0 for two different lengths N . With increasing N the correlation becomes stronger.




Figure 3.9: Grey-scale plot of $\chi^2(x_0, a)$. Dark areas correspond to high values of χ^2 . The global minimum at $(0.8, 1.85)$ is indicated by a cross. The objective function is heavily jagged, especially in the direction of x_0 . The long narrow valleys are responsible for the strong correlation between a and x_0 .




Figure 3.10: Example of a fit that yielded a rather wrong parameter estimate ($\hat{a}=1.896$) but nevertheless a very good correspondence between model and true trajectory. It shows that the large error of the parameter estimate is not due to a local minimum. The noise level was 10%.

3.5 Correlation between the Parameter Estimates

The non-exponential behaviour in Fig. 3.6 shall be examined in more detail now.

The mechanism leading to Eq. (3.6) can not be the cause of the large variances for long time series. While it can explain how the global minimum becomes higher than another local minimum, one would expect that the true parameters remain at least near a local minimum of χ^2 . Since the true variables were supplied as starting guesses, the estimate can not converge to a false minimum.

Plotting the estimated parameters against each other shows a strong correlation between \hat{a} and \hat{x}_0 (Fig. 3.8). Therefore χ^2 was examined as a function of both parameters (Fig. 3.9). The picture is characterised by long narrow valleys in which the parameter vector can move nearly freely without much change in the objective function.

Fig. 3.10 shows one of the fits that contributed to Fig. 3.6 and yielded a rather wrong parameter estimate. N was 60 and the noise level was 10%. The true initial value was $x_0=-0.8075$. The estimated parameters are $\hat{a}=1.896$ and $\hat{x}_0=-0.8070$. Astonishingly the trajectory is similar to the true one all along the 60 points although the parameter deviates considerably from the true value. Thus this is not a problem of local minima. The suspicion arises that for each parameter a' , there is a transformation h of




Figure 3.11: Empirical representation of the self-similar function h . The figure contains 10^4 points altogether. A small portion of them stems from unsuccessful fits that do not reflect the transformation h . They can be identified as isolated points. Most of them are located near the diagonal, since the true states were supplied as starting guesses. The other points form bold lines, representing the sought function. h is monotonic, but it has many discontinuities.

the state variable x to a new variable x' that fulfils the logistic equation for the parameter a' :

$$h(f_a(x)) = f_{a'}(h(x))$$

The transformation is depicted in a commutative diagram:

$$\begin{array}{ccccc}
 \dots x_{i-1} & \xrightarrow{f_a} & x_i & \xrightarrow{f_a} & x_{i+1} \dots \\
 \downarrow h & & \downarrow h & & \downarrow h \\
 \dots x'_{i-1} & \xrightarrow{f_{a'}} & x'_i & \xrightarrow{f_{a'}} & x'_{i+1} \dots
 \end{array}$$

Since the initial value is also affected by the transformation, deviations of x_0 are related to deviations of a , explaining the correlation between the estimates of the two variables.

Power expansion approaches showed that there is at least no continuous transformation of this kind. In order to conceive what is causing the wrong estimates, noiseless data of length $N=100$ were fitted with the parameter a




Figure 3.12: Same as Fig. 3.11, after subtracting the line of identity.

being fixed to the wrong value $a'=1.95$, so that only the initial value was a free parameter. This was done 100 times, with the true initial value being scanned from -0.85 to 1 . In this way, true trajectories and fit trajectories with a total of 10^4 points each were generated. They were plotted against each other in Fig. 3.11. Despite isolated spurious points that are due to unsuccessful fits, a clear image of the transformation function is discernible. h turns out to be monotonic, but highly discontinuous. Near $x = x' = 0$ the range contains a gap of size 0.257 . Similar, but successively smaller discontinuities occur at all inverse iterates of 0 , $\{x|f^n(x) = 0, n \in \mathbb{N}\}$, leading to a self-similar structure. The deformations become more distinct when $x' - x$ is plotted against x (Fig. 3.12). The transformation is close to the identity, though a' deviates considerably from a . That means that a large error in the parameter causes small deviations in the states. Reading it the other way around, the algorithm compensates part of the noise in the data by adjusting the parameter so that the transformation h imitates the noise.

Having revealed the strong correlation between \hat{a} and \hat{x}_0 , the first point in the list on page 24 has to be reconsidered. If \hat{a} is highly correlated with \hat{x}_0 , its a priori distribution could have an effect on the estimated parameter.

The mechanism causing the correlation is not a unique property of the logistic map. Similar transformations were found in the Moran-Ricker map

$f(x) = x e^{a(1-x)}$. Moreover, it is not restricted to maps. Systems with continuous time and state space could also possess such correlations, but the corresponding transformations must be continuous since discontinuities would cut the state space into non-contiguous parts.

3.6 Conclusion

This chapter examined to what extent the superior properties of the maximum likelihood estimator can be used to obtain accurate estimates of the parameters of a chaotic time-discrete system. The logistic map was chosen as a paradigmatic representative. Two problems turned out to aggravate the task:

- The unknown initial value x_0 is strongly correlated with the dynamic parameter a . While a single parameter can be estimated with an exponentially decreasing standard deviation, a simultaneous fit of both variables improves only like $1/N$, which is still much better than the $1/\sqrt{N}$ law achieved with other methods.
- The objective function is cluttered with local minima. Increasing N does not increase their level relative to the global minimum, but makes them sharper and adds even more of them. That means that errors caused by noise on the first data points can not be compensated by subsequent data points.

The latter point is a disadvantage of the maximum likelihood method. Its superior asymptotic properties rely on the assumption that the minimum of the objective function is actually found. If this can not be ensured, other methods might be preferable.

Chapter 4

Application to the Q-switched CO₂ Laser

The maximum likelihood approach for estimating parameters in dynamical systems is feasible if a parametric model of the right-hand side is on-hand. Regarding the CO₂ laser, the so called four-level model (4LM) is well established (Meyer-Bourbonneux et al. 1976; Ciofini and Meucci 1995). It is a nonlinear, five-variable ODE with six unknown parameters that shall be estimated from univariate time series of the intensity of a Q-switched CO₂ laser. The experiments were made at the Istituto Nazionale di Ottica Applicata in Firenze.

The application of an analysis method to measured time series is a much greater challenge than its test on simulated data. Because of partially unknown experimental conditions, and due to the preference for simple models, the dynamic equations are only an approximation of the actual dynamics. The same is true for the observation equation. The application reported in this chapter showed that the observation function is an essential part of the model and that the absence of a proper description can easily lead to wrong results.

For this reason, after the description of the experimental setup in the first section, the complete preprocessing procedure is included in Section 4.2. Section 4.3 presents the dynamic model and derives theoretical relations between the dynamic parameters and the initial values of the state vector, that will be used to reduce the number of degrees of freedom. After testing the procedure in Section 4.4, it is applied to ten records of measured data, that have been recorded under different experimental conditions. Section 4.6 concludes with a summary.




Figure 4.1: Experimental set-up. G: grating; LT: laser tube; PS: current-stabilised power supply; EOM: electro-optic modulator; HVA: high-voltage amplifier; M: spherical mirror; PZT: piezoelectric translator; D: detector; P: preamplifier; DO: digital oscilloscope; WS: workstation. Quantities: i : discharge current; n : photon density; I : laser intensity; V : detector voltage; y : recorded signal; TTL: EOM driver signal.

4.1 Experimental Setup

The experimental setup (Fig. 4.1) consists of a single-mode CO₂ laser with an intracavity electro-optic modulator (EOM) (Meucci et al. 1996). The optical cavity, 1.35 m long, is defined by an 80% reflectivity spherical mirror (radius of curvature 3.0 m) and by a grating blazed for 10.6 μm . The mirror, acting as the cavity outcoupler, is mounted on a piezoelectric translator in order to control the tuning between the cavity mode and the centre of the molecular line. The active medium, a gas mixture of 14% CO₂, 14% N₂, 2% H₂ and 70% He at an average pressure of 21 Torr, is contained in a pyrex tube terminated by Brewster windows. The medium is excited by a high-voltage DC discharge current. The power supply is highly stabilised against current fluctuations such as those occurring during the Q-switch. The current is stabilised to better than 0.01 mA.

At 6 ms intervals the EOM driver and the high-voltage amplifier provide 600 V pulses of 3 ms duration with a rise time shorter than 50 ns. At the switch-on at time $t = 0$ a fast jump of the cavity loss rate $K(t)$ from a higher value K_1 to a lower value K_2 is induced. $K_2 = 1.37$ MHz was assumed in all cases. This kind of steep external modulation induces the dynamic variables to explore a large region of their phase space. In this way the measured time series is supposed to contain a considerable amount of information about the laser.

The intensity is detected with a Hg-Cd-Te photodiode and amplified. The bandwidth of the detecting device is approximately 400 kHz. A time series of length 64 ms, containing 10-11 Q-switch events, is recorded with a digital oscilloscope having a resolution of 12 bit. The sampling frequency is 1 MHz. The TTL signal from the EOM driver is also recorded. Fig. 4.2 shows both time series. The vertical axis of the detector signal is inverted




Figure 4.2: Measured time series of the detector signal and the TTL signal as recorded by the digital oscilloscope. Discharge current $i=5.22$ mA. The detector signal axis is inverted because low values correspond to high laser intensities. The TTL signal is switched on at 6 ms intervals for 3 ms duration each. The cavity loss rate is low when the TTL signal is high and vice versa.

because low values correspond to high laser intensities. The TTL signal is high when the cavity losses $K(t)$ are low.

After the Q-switch (*on-state*) the intensity describes a large spike, followed by relaxation oscillations. At the end of the EOM pulse, $K(t)$ jumps back to K_1 for another 3 ms (*off-state*) before the next Q-switch is induced. For the estimation of parameters, only the first 500 μs after each Q-switch will be used. The most important portions of the time series are the first few relaxation oscillations after each Q-switch. They last for some 100 μs and are not resolved in Fig. 4.2.

As the peak intensity depends on the value of the population inversion which is proportional to the discharge current, the discharge current i is chosen as the control parameter, repeating the procedure for 10 different current values between 4.20 and 5.22 mA. The records were divided into three groups with respect to their dynamical behaviour. The following table gives an overview.

Regime	Low currents	Medium currents	High currents
Current i	4.20	4.40	4.80
$[mA]$	4.30	4.50	4.90
		4.60	5.05
		4.70	5.22

With increasing current a growth of the peak height is observed, accom-

panied by a reduction of the delay time between the EOM pulse and the spike (Meucci et al. 1992).

4.2 Preprocessing

This section describes all steps necessary to relate the dynamic variables of the model to the recorded time series.

4.2.1 Baseline Correction

Using the TTL signal, the individual Q-switch events are identified and cut out from the continuous time series. They are also called pulses in the following. 500 points before each Q-switch are used to determine the baseline for this pulse. From the 10-11 baselines of each record the ensemble average $B(i)$ and its standard deviation are computed and plotted in Fig. 4.3. For the six low and medium currents, $B(i)$ is approximately constant, indicating zero intensity. Its average

$$\bar{B} = \frac{1}{6} \sum B(i)$$

is given by a dashed line in Fig. 4.3. The variability between the records is much larger than the error bars, indicating that it is mainly due to long term drifts with correlation times that are longer than the record length (64 ms), but shorter than the breaks between the measurement of the records. In order to suppress these drifts as much as possible, each pulse is adjusted so that its baseline vanishes.




Figure 4.3: Baselines estimated from raw data for all currents. For high currents the baseline is non-vanishing, indicating that the laser does not switch off completely.

Beginning with $i=4.80$ mA the baseline increases with the current, indicating a non-zero steady level in the off-state. Now the true baseline, corresponding to zero intensity, is not known anymore. As a substitute, the total average baseline \bar{B} of the six low and medium currents is used and each pulse within a record for a high current i is adjusted such that its baseline is $B(i) - \bar{B}$. For high currents the error bars are larger because the steady-level in the off-state depends on experimental parameters which are subject to fluctuations in the discharge current. Finally all signals are inverted so that higher values correspond to higher intensities.

4.2.2 Pulse Merging

Since the sampling frequency is limited to 1 MHz, the first peak after the Q-switch contains only 10 data points for high currents. On the other hand, the different pulses of a record have a very similar shape for medium and high currents. It is therefore to be expected that they can be overlaid to a single pulse containing all data points, thereby compensating for the low sampling frequency.

Simple Alignment

While the baselines have already been adjusted as described earlier, the time axes remain to be aligned among the pulses. Fig. 4.4 shows the result when this is done by using the TTL signal as a trigger. The time at which it crosses a threshold value, which is the mean between the on- and the off-level, is taken as the origin of the time axis. Since the TTL signal has the same sampling frequency and switches from off to on within much less




Figure 4.4: Alignment of the pulses using the TTL signal as a trigger. The aligned times are still integral multiples of the sampling interval. Discharge current $i=5.22$ mA.




Figure 4.5: Result of the internal alignment procedure. All data points lie almost perfectly on a unique curve. The arrow illustrates the procedure taking the second maximum as a threshold being crossed in the first decline. Discharge current $i=5.22$ mA.

than $1 \mu\text{s}$, the switch time can be detected only with an uncertainty of $1 \mu\text{s}$. Consequently the new time values are still integral multiples of the sampling interval. So the TTL signal is an insufficient aid for this purpose. Instead the following procedure is used to align the pulses.

Internal Alignment

- Determine the height h of the second maximum by quadratic interpolation.
- Determine the trigger time t_0 at which the decline after the first peak crosses the threshold h , using linear interpolation.
- Define t_0 as a temporary new origin of the time axis for this pulse.

The result is shown in Fig. 4.5 for the highest current. All data points lie almost perfectly on a unique curve. The method works well for all high and medium currents though the resulting curves are not as smooth as in the example shown. They can be seen in Figs. 4.20 and 4.21 on page 64 and 65.

External Alignment

Regarding now the TTL signals in the new time coordinates, the time of the Q-switch relative to the trigger time can be determined with an uncertainty of about $\pm 0.2 \mu\text{s}$. It is used as the final origin of the time axis for this record. This external alignment between the records is less important than




Figure 4.6: Merged pulses including off-state phase for medium and high currents. The currents are 4.40, 4.50, 4.60, 4.70, 4.80, 4.90, 5.05 and 5.22 mA from bottom to top. The vertical lines at $t=0$ and $t=3039 \mu\text{s}$ indicate the switch-on and switch-off times respectively. The spikes shortly after the switch-on exceed the upper limit of the vertical axis. The first relaxation oscillations are not visible in this figure.

the internal alignment between the pulses because the dynamical models used are invariant under global time shifts.

For the low currents, the pulses are so variable that it does not make sense to align them with each other. They are not merged into a single time series but their time values are normalised by the simple alignment procedure described above, using the TTL signal as a trigger. Since these records exhibit slower dynamics than those for medium and high currents, the spikes are represented by more points and thus the low sampling frequency is a minor problem.

Fig. 4.6 shows the complete merged time series for all medium and high discharge currents. The low currents are a special case that will be treated separately. The switch-on and switch-off times are indicated by vertical lines. The prominent first peaks after the Q-switch are cut off by the vertical range

in order to emphasise the other features in this picture. After a few hundred μs the intensity becomes constant, despite a slow increase that is ascribed to temperature effects.

Following the switch-off, the measured intensities rise to a higher level before they decrease to their respective baselines within 1 or 2 ms. This slow switch-off behaviour reflects the properties of the EOM which is designed to be very fast at the switch-on (time constant smaller than 50 ns) while the switch-off characteristic is less important.

The last 1-2 ms before the following switch-on are characterised by a slow decrease that is again ascribed to temperature effects. The off-phase clarifies the distinction of medium and high discharge currents. For the medium currents (four lowest curves), zero intensity is attained, while for the high currents the baseline progressively increases.

In Meucci et al. (1991) the time interval to reach thermal equilibrium was directly measured to be of the order of 3 ms for a voltage-stabilised power supply. Fig. 4.6 suggests that in the case studied here, there are still slight drifts after 3 ms. However, for the analysis of the dynamical behaviour of the laser in the remainder of this chapter, only the first 500 μs after the Q-switch will be utilised. The remaining part of the on-phase, representing a steady level, does not contain any additional information. The off-phase is useless due to the unknown characteristic of the EOM. Within this short fitting interval, the temperature drifts are negligible. Therefore one can assume that 3 ms are a sufficient time interval to reach an acceptable equilibrium and then to switch on the laser again.

In Fig. 4.7 the integral of the laser intensity is plotted as a function of the discharge current for medium and high currents. While the solid curve




Figure 4.7: Total integrated output power as a function of the discharge current. **Solid line:** integral of the laser intensity over the complete pulse. **Dashed line:** integral from 0 to 4 ms.

was obtained by integrating over the complete pulse, from one Q-switch to the following, for the dashed curve the integration was restricted to the interval between 0 and 4 ms which includes the on-state and the relaxation to the off-state, but not the stationary part of the off-state. The figure shows that the amount of energy radiated due to the non-zero baseline for high currents is rather small. Thus the increase of the total power with the current is mainly due to the increase of the on-state level.

4.2.3 Calibration Procedure and Observation Equation

The measurement process includes two further steps distorting the measured laser intensity I . They must be taken into account either by preprocessing or by means of the observation equation (2.2).

The detector voltage V is assumed to depend on I according to the weakly nonlinear function

$$V = f(I) := \frac{I}{1 + \alpha I + \beta I^2} \quad (4.1)$$

with unknown constants α and β . This is a refinement of the nonlinearity used in Varone et al. (1995) that is able to describe the observed behaviour more precisely than the original formula. In the range of interest there exists an inverse of f which is denoted f^{-1} . Eq. (4.1) needs not to be understood in terms of physical processes. The purpose of the following preprocessing is to free the data from as many systematic errors as possible before analysing them with the methods described in Chapter 2.

For low and medium currents the impact of the detector nonlinearity is slight and therefore not identifiable from the data. However, since the constants α and β do not depend on the current, they can be determined from measurements with the highest current and then be used for all currents. For this purpose, two time series y and y' were recorded under the same experimental conditions except that for the second measurement the detector was slightly moved aside such that only a fraction a of the light was recorded. Striped quantities refer to this second data set in the following. Combining Eq. (4.1) and its inverse yields

$$V' = f(I') = f(aI) = f(af^{-1}(V)), \quad (4.2)$$

i.e. V' is a function of V , parametrised by a , α and β .

Fig. 4.8 shows both signals after the external alignment procedure described in Section 4.2.2. The upper curve is the same as in Fig. 4.5, despite the external alignment. The lower one corresponds to the second measurement with the displaced detector.

The internal alignment procedure was defined such that the trigger time is invariant under the application of the detector nonlinearity, so that if the




Figure 4.8: Comparison of time series for the highest current $i=5.22$ mA. All pulses in a record were merged according to Section 4.2.2. Arrows indicate the trigger times and the points in time where the signals take their maximum.

entire curve were subjected to a transformation similar to Eq. (4.1), the trigger time would not change. Another point in time having this property is the time where the signal takes its maximum. These times are indicated by arrows in the figure.

The uncertainty left by the external alignment is $\pm 0.05 \mu\text{s}$ and $\pm 0.2 \mu\text{s}$ for the upper and lower curve respectively. However, these uncertainties concern only a global shift between the two records and not the distance between the time of the maximum and the trigger time. From Fig. 4.8 it can be clearly seen that this distance differs between the full measurement and the down-sized version. The difference is more than $1 \mu\text{s}$. The reason is that the measured value y is not proportional to V in Eq. (4.1). The detector circuit has a low pass characteristic that can be expressed as a linear ODE:

$$\dot{y} = \tau^{-1}(V - y), \quad (4.3)$$

where y is the recorded signal and τ is the time constant of the detecting device. The following paragraphs will show how this time constant can be determined from the two measurements explored above.

If τ is known, Eq. (4.3) can be solved for V :

$$V = y + \tau \dot{y} \quad (4.4)$$

From the recorded time series y and y' the first derivatives \dot{y} and \dot{y}' were calculated by means of a kernel estimator (Gasser et al. 1985) and then V and V' were reconstructed using Eq. (4.4) for various values of τ .

As for the case $\tau = 0$, $V' = y'$ is plotted versus $V = y$ in Fig. 4.9A. There is no unambiguous functional relationship. On the other hand, with




Figure 4.9: Plot of V' versus V ; **A:** $\tau = 0$, no unambiguous functional relationship is present; **B:** $\tau = 2.35 \mu\text{s}$, all points lie approximately on a single curve; dashed line: best fit of Eq. (4.2) to these points.

$\tau = 2.35 \mu\text{s}$, which was numerically found to be the optimal value, the points $(V; V')$ are well described by a single curve (Fig. 4.9B, solid line). This value is in good agreement with the one obtained when measuring the frequency cut-off with standard methods.

Now, $\alpha = 4.55 \cdot 10^{-4}$, $\beta = -9.03 \cdot 10^{-9}$ and $a = 0.081$ were estimated by a nonlinear fit of Eq. (4.2) to these points (Fig. 4.9B, dashed line). Moderate precision was sufficient for this calculation because these constants represent only small corrections to the measurement curve. The constant τ determined in this way was used to reconstruct $I(t)$ for all measurements in the remainder of this chapter. The nonlinearity function f in Eq. (4.1) was directly taken into account as the observation function g in Eq. (2.2).

4.3 Theoretical Investigation

Following the observation model, this section describes all aspects of the dynamic model.

4.3.1 Dynamical Model

Lasers are distinguished with respect to the number of differential equations ruling their fundamental dynamics. According to this classification the CO_2 laser is a class-B laser, being ruled by two equations for the photon density and the population inversion. The corresponding model is the *two-level model* (Arecchi et al. 1982). In the past it became evident that upper and




Figure 4.10: Schematic diagram of the four-level model after Meucci et al. (1992).

lower levels have different relaxation times and that there is considerable interaction between the lasing levels and the other levels of the rotational manifold (Dupré et al. 1975; Arecchi et al. 1988; Ciofini and Meucci 1995). The resulting five-dimensional four-level model (4LM), as described in Meucci et al. (1992), is used here as a model for the dynamics of the CO₂-laser. A schematic diagram of the model is shown in Fig. 4.10. It includes two resonant levels whose populations are denoted by N_1 and N_2 , respectively. There are Z additional rotational levels in each vibrational band with total population densities M_1 and M_2 .

The populations decay reversibly into levels of the same band with the rate constant γ'_R , and irreversibly into other levels with rate constants γ_1 and γ_2 respectively. The flux from a lasing level N_i to its other rotational levels is $Z\gamma'_R N_i =: \gamma_R N_i$. The population inversion $\delta = N_2 - N_1$ is achieved by a pump rate $\gamma_2 P$ acting on each upper level.

The measured laser intensity is proportional to the photon number $n = |E|^2$. The complex electric field amplitude E is amplified by stimulated emission with a field-matter-coupling constant G and damped with a cavity loss parameter $K(t)$ which is switched from a higher value K_1 to a lower value K_2 at time zero. The time constant of this Q-switch is smaller than

50 ns and can be neglected in the model. The 4LM reads

$$\dot{E} = (-K(t) + \frac{1}{2}G\delta)E + \xi(t) \quad (4.5a)$$

$$\dot{N}_1 = -\gamma_R N_1 + \gamma'_R M_1 + G\delta n - \gamma_1 N_1 \quad (4.5b)$$

$$\dot{N}_2 = -\gamma_R N_2 + \gamma'_R M_2 - G\delta n - \gamma_2 N_2 + \gamma_2 P \quad (4.5c)$$

$$\dot{M}_1 = +\gamma_R N_1 - \gamma'_R M_1 - \gamma_1 M_1 \quad (4.5d)$$

$$\dot{M}_2 = +\gamma_R N_2 - \gamma'_R M_2 - \gamma_2 M_2 + \gamma_2 ZP \quad (4.5e)$$

with

$$\begin{aligned} \delta &= N_2 - N_1 \\ \gamma_R &= Z\gamma'_R \\ n &= |E|^2 \end{aligned}$$

Eq. (4.5a) is a complex Ornstein-Uhlenbeck process. $\xi(t)$ models the spontaneous emission fluctuations. It is a complex stochastic function whose real and imaginary parts are independent white Gaussian noise with zero mean and variance s^2 per unit time interval:

$$\begin{aligned} \xi(t) &= \xi_1(t) + i\xi_2(t) \quad (4.6) \\ \langle \xi_i(t) \rangle &= 0 \\ \langle \xi_i(t)\xi_j(t') \rangle &= 0 \quad \text{for } t \neq t' \quad \text{or } i \neq j \\ \text{var} \left(\int_t^{t+dt} \xi_i(t') dt' \right) &= s^2 dt. \end{aligned}$$

where $\langle . \rangle$ denotes the expectation value and $\text{var}(\cdot)$ the variance with respect to the probability distribution of ξ . s^2 is related to the population density of the upper lasing level through $s^2 = GN_2$.

In the following, stars will denote complex conjugation. The Langevin equation for E , Eq. (4.5a) is transformed into one for n :

$$\begin{aligned} \dot{n} &= \dot{E}E^* + E\dot{E}^* \\ &= (AE + \xi)E^* + E(AE^* + \xi^*) \quad \text{with } A = -K + \frac{1}{2}G\delta \\ &= 2AEE^* + 2\text{Re}\{\xi^* E\} \\ &= 2An + 2\sqrt{n}\text{Re}\{\xi^* e^{i\phi}\}, \end{aligned}$$

where ϕ is the phase of E . Since $\xi(t)$ has an isotropic two-dimensional Gaussian distribution in the complex plane, the same is true for $\xi^*(t)$ and for $\xi^*(t)e^{i\phi}$. Therefore $\tilde{\xi}(t) := \text{Re}\{\xi^*(t)e^{i\phi}\}$ is real-valued Gaussian white noise with variance s^2 . It follows that

$$\dot{n} = 2An + 2\sqrt{n}\tilde{\xi} \quad \text{with} \quad \langle \tilde{\xi}(t)\tilde{\xi}(t') \rangle = s^2\delta(t-t'). \quad (4.7)$$

$\tilde{\xi}(t)$ is not correlated with $n(t)$, which is determined by the past history of $\xi(t)$. Therefore the expectation value of $2\sqrt{n}\tilde{\xi}$ vanishes and the deterministic part of the ODE reads

$$\frac{d}{dt}\langle n \rangle = 2A\langle n \rangle. \quad (4.8)$$

It will be shown that in the region of interest, this equation is sufficient for the modelling of measured time series. The stochastic noise can have a macroscopic effect but it causes only a time shift of the model trajectory. Therefore it will be neglected.

As a further change, the average population densities of the non-lasing rotational levels $m_i = M_i/Z$ are used as dynamical variables instead of M_i , because they have the same order of magnitude as N_i , thus rendering the ODE integration and the optimisation numerically more efficient. Now the four-level model reads:

$$\dot{n} = (-2K(t) + G\delta)n \quad (4.9a)$$

$$\dot{N}_1 = -\gamma_R N_1 + \gamma_R m_1 + G\delta n - \gamma_1 N_1 \quad (4.9b)$$

$$\dot{N}_2 = -\gamma_R N_2 + \gamma_R m_2 - G\delta n - \gamma_2 N_2 + \gamma_2 P \quad (4.9c)$$

$$\dot{m}_1 = +\gamma'_R N_1 - \gamma'_R m_1 - \gamma_1 m_1 \quad (4.9d)$$

$$\dot{m}_2 = +\gamma'_R N_2 - \gamma'_R m_2 - \gamma_2 m_2 + \gamma_2 P. \quad (4.9e)$$

with

$$\delta = N_2 - N_1$$

$$\gamma_R = Z\gamma'_R.$$

It is important to observe that there is a symmetry between γ_1 and γ_2 in the model. When a solution of Eq. (2.6a) with certain parameters is given, the transformation

$$\tilde{n} = n \quad (4.10a)$$

$$(\tilde{N}_1, \tilde{N}_2) = (P - N_2, P - N_1) \quad (4.10b)$$

$$(\tilde{m}_1, \tilde{m}_2) = (P - m_2, P - m_1) \quad (4.10c)$$

$$(\tilde{\gamma}_1, \tilde{\gamma}_2) = (\gamma_2, \gamma_1) \quad (4.10d)$$

leads to another solution of Eq. (4.9) with the same observation. Thus it cannot be decided from a measurement of n , whether the first or the second solution is true. This is accounted for by restricting to solutions with $\gamma_1 \geq \gamma_2$.

The laser intensity I measured by the detector is proportional to the photon density n inside the cavity:

$$I = k n. \quad (4.11)$$

The unknown coefficient k is eliminated by transforming n, N_i, M_i, P and G^{-1} to new states and parameters by multiplication with k . The ODE system is invariant under this gauge transformation but one has to keep in mind that the transformed state variables and parameters contain an unknown factor. In particular, G and P are not comparable with any theoretically computed number. The qualitative behaviour of the model is visualised in Fig. 4.14 on page 59.

4.3.2 Laser Startup Process

In the following the dynamics of the state vector from the Q-switch to the first peak will be shown to go through a stochastic, a linear and a nonlinear regime. The neglect of the stochastic term in the ODE and relations for the initial values of the state vector will follow from these considerations.

Eq. (4.9) has the following fixed point:

$$n = 0 \quad (4.12a)$$

$$N_1 = m_1 = 0 \quad (4.12b)$$

$$\delta = N_2 = m_2 = P. \quad (4.12c)$$

As long as the *laser net gain* $-2K + G\delta$ is negative, the fixed point is dynamically stable and is called the *off-state*. The stochastic part of Eq. (4.7) causes small fluctuations around the fixed point. $n = |E|^2$ has a χ^2 distribution with two degrees of freedom and mean

$$\langle n \rangle = \frac{2s^2}{2K_1 - GP}. \quad (4.13)$$

When $K(t)$ is switched to its lower level, the gain becomes positive and the fixed point becomes unstable. Now the stochastic fluctuations start up an exponential increase of the photon number:

$$\langle n(t) \rangle \sim e^{(-2K_2 + GP)t}. \quad (4.14)$$

The stochastic regime ends, when the photon density has grown much larger than s . Then further stochastic influence on the dynamics can be neglected. At this time, n, N_1 and m_1 are still small and N_2, m_2 and δ can be approximated by their steady-state values. Thus the ODE is reduced to the linear system

$$\frac{d}{dt} \begin{pmatrix} n \\ N_1 \\ m_1 \end{pmatrix} = \mathbf{L} \begin{pmatrix} n \\ N_1 \\ m_1 \end{pmatrix} \quad (4.15)$$

with

$$\mathbf{L} := \begin{pmatrix} -2K_2 + GP & 0 & 0 \\ GP & -\gamma_R - \gamma_1 & \gamma_R \\ 0 & \gamma'_R & -\gamma'_R - \gamma_1 \end{pmatrix}. \quad (4.16)$$

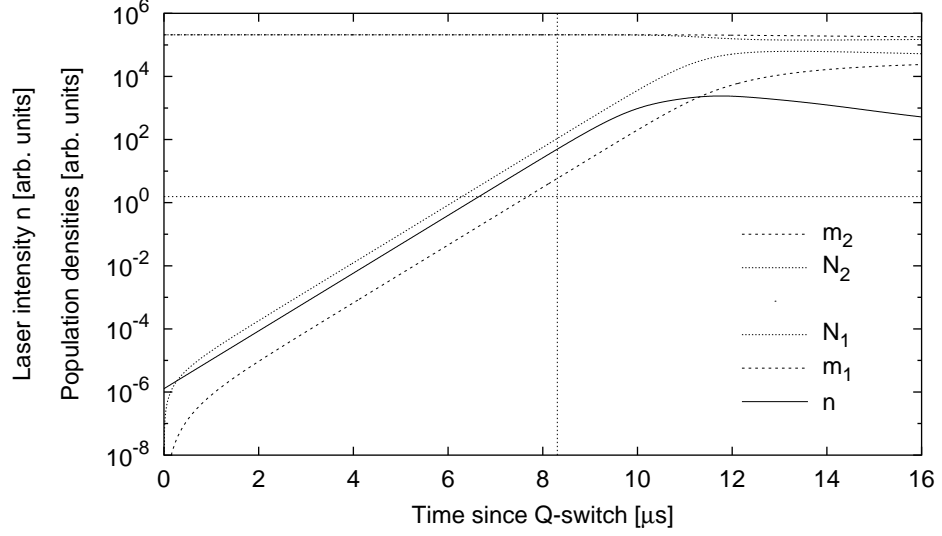


Figure 4.11: Theoretical startup behaviour. The model was simulated with the parameters that were estimated in Section 4.5.2 from measured data for $i=4.60$ mA. The initial values were taken from Eq. (4.12). Then all components of the state vector are plotted logarithmically. The legend lists the components in the same order as they appear at the right end of the graph. **Horizontal dotted line:** noise level apparent in the measured data set. **Vertical dotted line:** beginning of fit interval (see Section 4.3.3).

The region of validity of Eq. (4.15) is called the linear regime.

The only positive eigenvalue of \mathbf{L} is the initial laser net gain $\lambda = -2K_2 + GP$. The amplitude of the spike and the time at which it occurs are strongly dependent on λ . The state vector increases exponentially along the direction of the corresponding eigenvector \mathbf{v} , defined through

$$\mathbf{L} \cdot \mathbf{v} = \lambda \cdot \mathbf{v}, \quad (4.17)$$

until the effects of the nonlinear term $G\delta n$ in Eq. (4.9) become significant and the nonlinear regime begins.

The exponential startup of the trajectories is visualised in Fig. 4.11 by a simulation with the parameters that were estimated in Section 4.5.2 for $i=4.60$ mA. Within $1 \mu\text{s}$ the populations N_1 and m_1 have taken their proportions relative to n that are defined by Eq. (4.17). A long linear regime follows in which these proportions are maintained and the populations of the upper levels N_2 and m_2 are constant. At $7 \mu\text{s}$ the observable reaches a value that is equivalent to the noise level of the measured data. After $9 \mu\text{s}$ the nonlinear regime begins and the state vector deviates from the direction of \mathbf{v} .

The corresponding picture with measured data for medium and low cur-




Figure 4.12: Startup for medium and low currents in semilogarithmic representation. The leftmost curves correspond to the currents $i=4.70$, 4.60 , 4.50 and 4.40 mA, from left to right. Following are the pulses of the record with $i=4.30$ mA and, in the right half of the figure, the pulses for $i=4.20$ mA. The dynamical range of the signal is 2-3 orders of magnitude.

rents is shown in Fig. 4.12. Due to the observation noise, only a small portion of the linear regime is visible. By extrapolating the straight lines representing the linear regimes until they meet the vertical axis, one can estimate the mean initial photon number in Eq. (4.13) and therefore the level of the quantum noise s^2 . This will be done in Section 4.5.3. From the visual inspection one can deduce that $\langle n(t=0) \rangle$ is much smaller than the observation noise.

Since the direction of \mathbf{v} is independent of $\xi(t)$, only the length of the state vector is affected by the noise term. However, a scaling of an exponential increase is equivalent to a shift of the time axis, i.e. the trajectories of the individual pulses are displaced against each other. This phenomenon can be seen as *pulse time jitter* in the experiment. It disappears when the current exceeds the threshold for laser activity in the off-phase. In this way the boundary between medium and high currents can easily be detected in the experiment. The turn-on statistics have been investigated in detail by many authors (Arecchi et al. 1989; Balestri et al. 1991; Balle et al. 1991, 1994; Ciofini et al. 1990; Grassi et al. 1994).

For the further analysis, time shifts are not considered since they are removed by the internal alignment procedure described on page 43. Therefore the deterministic ODE Eq. (4.9) will be taken as the model ruling the dynamical behaviour of the laser after the Q-switch.

4.3.3 Consistency of the Initial State Vector

In addition to the general methods outlined in Chapter 2 some strategies for parameter estimation are specific to the considered laser model. A crucial point is the beginning of the fit interval. When the initial values of all state variables were treated as free parameters, the algorithm signalled rank deficiency and the estimated unobserved components were unreasonable, i.e. it was not possible to estimate all parameters and initial values simultaneously from the given data. This problem can be solved by ensuring the initial state to be consistent with the past history of the data, i.e. by incorporating additional physical knowledge about the laser buildup. This will be done in the following paragraphs. The theory of the laser startup process outlined in Section 4.3.2 is important in this context.

a) Low and Medium Excitation Currents

For currents which are not too high, the system is in the off-state Eq. (4.12) before the Q-switch. If the fit starts at $t = 0$, the steady-state can in principle supply the initial values for the unobserved components. Unfortunately, the trial trajectory would then be arbitrarily sensitive to numerical fluctuations of the initial value n_0 , bringing instabilities into the iterative optimisation process. Therefore the fit was started at the time at which the intensity reaches 2% of its peak value. At this time the system is still in the linear regime.

Accordingly, the proportions of n , N_1 and m_1 are given by the eigenvector defined in Eq. (4.17). Eqs. (4.12c) and (4.17) are used to relate the initial values of the unobserved components at the beginning of the fit interval to n_0 by means of equality constraints as described in Section 2.1.1.

b) High Excitation Currents

For high currents and resultingly large pump parameters, the situation is different. The laser gain is then always positive, even during the breaks, when the cavity losses are at their higher value K_1 . Consequently the laser is not switched off and on, it is rather switched between two different on-states. The data points prior to $t = 0$ correspond to a steady-state with non-vanishing intensity. For these currents the fit start was placed at $t = 0$ and the steady-state equations that result when the right-hand sides of Eqs. (4.9b-e) are set to zero, were used as constraints to the initial values. In the same way the constraint $\dot{n} = 0$ could be demanded at $t = 0$, but it

would not provide any additional information about the parameters because Eq. (4.9a) contains the unknown constant K_1 . Conversely this will be used to estimate the value of K_1 in Section 4.5.3.

4.4 Simulation Study

The outcome of a fitting procedure is not always guaranteed to be meaningful. There could be a transformation like the one in Eq. (4.10) that changes the state vector and the parameters but leaves the observation unaffected. It is also thinkable that there exists a continuous family of such transformations so that the parameter vector may be varied continuously without a change of the objective function. In this way a continuous ambiguity is created. Such redundancies are not always easily seen.

Therefore, in this section, the outlined methods are tested under conditions that are as close as possible to realistic situations. A time series of the laser intensity was simulated over $T=500 \mu\text{s}$ with a sampling frequency of 11 MHz using Eqs. (4.9) and (4.1) with the parameters that were estimated in Section 4.5 for the excitation current $i = 4.80 \text{ mA}$. Similar results were obtained for the other currents. Noise with the same properties as in the measured data was added. Then the parameters $G, \gamma_1, \gamma_2, \gamma_R, \gamma'_R$ and P were estimated as described in Chapter 2. Their starting guesses have been set to twice the true values. The initial value n_0 was another free parameter while the other initial values were related to n_0 according to Section 4.3.3b. The constraints from Section 4.3.3a could be used as well because the baseline $n(t<0)$ is still very low for this current. Indeed the estimated parameters were the same in this case.

Fig. 4.13 shows three stages of the iterative process. 40 subintervals were used for the multiple shooting method. The density of the multiple shooting mesh was adapted to the problem. The initial dynamics including the spike show a fast change of the state variables over several orders of magnitude. In contrast the relaxation oscillations evolve slowly to a steady-state. These different conditions were accounted for by using many short subintervals for the spike phase and rather long subintervals in the relaxation phase.

After 23 iterations the fit converged. The estimated parameters deviated by at most 4% from the true values. Fig. 4.14 shows all five components of the model trajectory. The ability to construct estimates of the hidden variables of the physical system is one of the advantages of this modelling procedure compared to delay embedding techniques.

4.5 Experimental Results

Each record contains 10 or 11 Q-switch pulses, sampled with a resolution of 1 MHz. For each pulse the baseline was estimated by the last 500 points

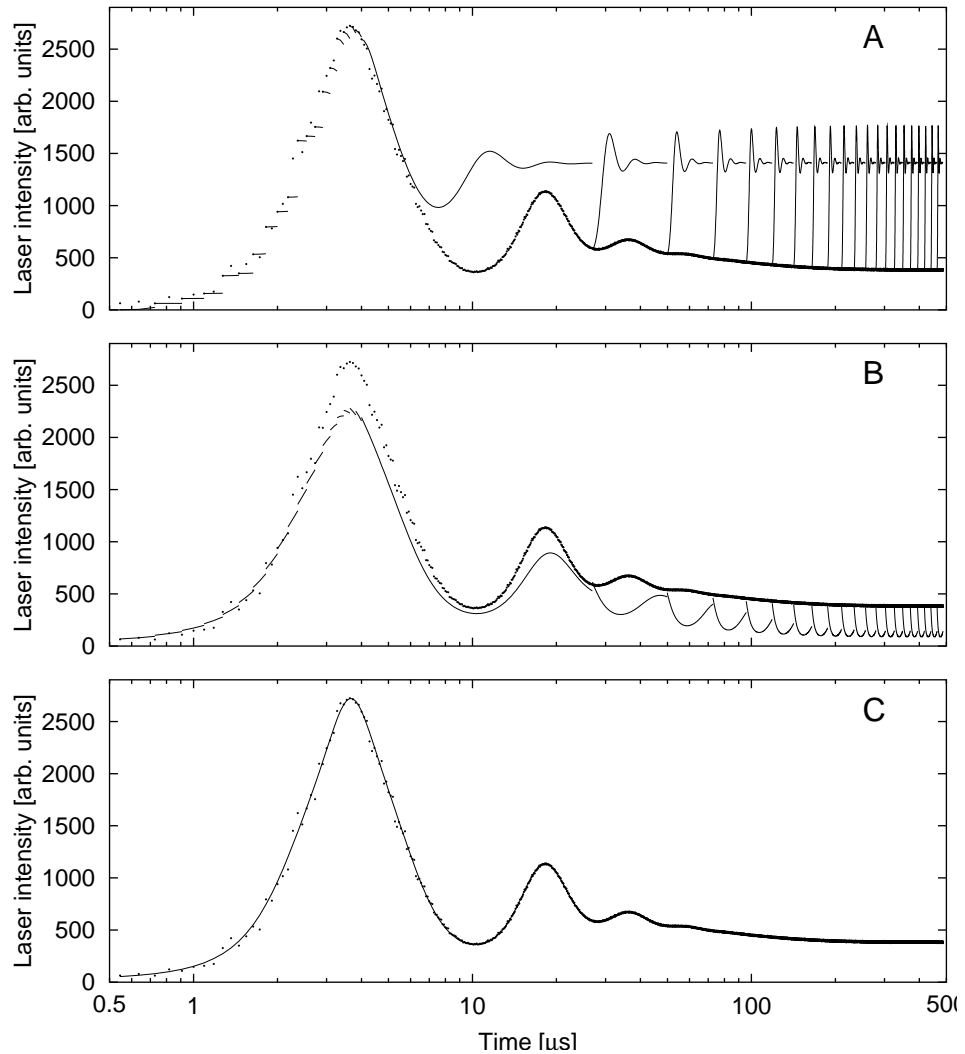


Figure 4.13: Iterative fit process for a simulated time series. **Dots:** data. **Lines:** model trajectory. The time axis is logarithmic in order to make the first peak more distinct. **A:** initial guess. The density of the multiple shooting mesh is changed to a lower value at $t = 4 \mu\text{s}$. In the rear part the initial values “do not yet know” that the Q-switch is already past. **B:** after 15 iterations. The trajectory is still discontinuous. **C:** convergence after 23 iterations.




Figure 4.14: Constructed time series of observed and unobserved components of the state vector for the fit shown in Fig. 4.13. The laser intensity n exhibits a large spike followed by relaxation oscillations to a steady level. The oscillations are reflected in the population densities of the lasing levels N_i while they are damped in the other levels of the rotational manifold.

before the Q-switch and subtracted from the signal. For high currents the laser intensity is non-zero at $t = 0$. Therefore the baselines determined for low and medium currents had to be averaged and substituted for the high currents. The impact on the estimated parameters was never larger than 1 %.

α , β and τ were determined as described in Section 4.2.3. Then Eq. (4.4) was used to correct all records for the low-pass characteristics. Eq. (4.1) was directly taken into account as the observation equation. The preprocessed data are given as points in the following figures.

4.5.1 Low Currents

For the lowest current, the dynamics differ considerably between the individual pulses as seen in Fig. 4.15. To understand this phenomenon, the initial laser net gain $\lambda = -2K_2 + GP$ is regarded. As mentioned in Section 4.3.2, the amplitude of the spike is strongly influenced by λ . The situation is depicted in Fig. 4.16. P and therefore GP increase with the pump current, as indicated by horizontal lines. $2K(t)$ is switched from K_1 to the lower level K_2 . For currents below $i_0 = 4.20$ mA, no laser action can be observed at all, implying that λ crosses zero below this current. For the lowest current $i = i_0$, GP is only slightly above $2K_2$. Small fluctuations in the pump current then cause large relative variations in the laser net gain which is represented by the hatched area between the two curves. As a consequence, the pulses have




Figure 4.15: Measured data for the lowest pump current $i=4.20$ mA. The points of each pulse are connected with lines.




Figure 4.16: Depiction of the mechanism generating the large pulse variations for low currents. **Horizontal lines:** increasing level of GP . **Thick line:** time course of $2K(t)$ which is switched from K_1 to the lower level K_2 . For the lowest current, $i=4.20$ mA, the net gain $-2K_2 + GP$ (**hatched**) is only slightly positive.

different heights and shapes.

This qualitative effect has already been understood theoretically and confirmed experimentally (Balle et al. 1994; Grassi et al. 1994); however no quantitative comparison between theory and experiment was made in these studies. The approach here is able to model the measured data in full detail. In a multiexperiment analysis all pulses were utilised simultaneously. Only the pump parameter P was allowed to attain an individual value for each pulse, while $G, \gamma_1, \gamma_2, \gamma_R$ and γ'_R were forced to be the same for all pulses. The end of the fit interval was set to $300 \mu\text{s}$ after the Q-switch. Beyond this point the signal is expected not to contain substantial information due to the low signal-to-noise ratio.

As Fig. 4.17 shows, both the amplitudes and shapes of all pulses are reproduced well. The estimated parameters are: $G = 26.7 \cdot 10^{-12} \text{s}^{-1}$, $\gamma_1 = 25.7 \text{ kHz}$, $\gamma_2 = 4.79 \text{ kHz}$, $\gamma_R = 42.0 \text{ kHz}$, $\gamma'_R = 24.9 \text{ kHz}$, $P = 109 \cdot 10^{15}$.

In Fig. 4.18 the estimated value of P and the resulting net gain λ are plotted versus the experimental spike height. A clear monotonic functional relationship can be seen. The empirical relative standard deviation of P is only 1.4%, i.e. P does not make large use of its freedom to vary between the records. On the other hand, λ has a relative standard deviation of 22%, being responsible for the broad spectrum of spike shapes. If P were fixed to be the same for all pulses, the model would only be able to produce 11 spikes with equal height and shape.




Figure 4.17: Data (**points**) and model trajectories (**lines**) for $i=4.20 \text{ mA}$. For the sake of clarity only 6 of 11 pulses are shown. The large variability of the curves was reproduced very well although only P and n_0 were varied independently for each pulse. The estimated parameters are given in the text.




Figure 4.18: Relation between the experimental spike heights (abscissa) and the estimated values of P (left axis) and $\lambda = -2K_2 + GP$ (right axis) for $i=4.20$ mA. Note the ranges of the two axes for P and λ . There is a factor 2 between the minimum and the maximum of λ while the P values cover a range that is only the 20th part of their mean.




Figure 4.19: As Fig. 4.18, but for $i=4.30$ mA. The relative variation of P is the same as above while that of λ is less.

The good correspondence between measurements and model can be taken as a strong argument for the validity of the four-level model for this experimental condition. For $i=4.30$ mA the situation is similar, but less distinct (see Fig. 4.19). In this case the standard deviation of P is again 1.4% while that of λ is 6.5%.

A comparison between the estimated values of the pump parameter for the two currents shows that the relative sensitivity of P with respect to the discharge current $\frac{1}{P} \frac{\partial P}{\partial i}$ is of the order of some mA^{-1} . Thus a discharge current uncertainty of 0.01 mA would suffice to induce a relative uncertainty of P of several percent, so the estimated variance of P is in agreement with the experimental conditions.

4.5.2 Medium and High Currents

For currents $i=4.40$ mA and higher, the variance in the spike height is strongly reduced since the variance of the net gain is small compared with its absolute value. Moreover the spike becomes narrower with increasing current, thus being sampled by successively fewer data points. In order to increase the sampling frequency, all pulses of a record were merged into a single time series as described in Section 4.2.2. The data points were gathered together with time values shifted such that each pulse had time zero at a trigger point. The trigger point was the time at which the first decline after the spike fell below the second maximum. For medium currents the pulse shapes are not yet perfectly equal, thus the merged time series seem to contain some noise as can be seen in Fig. 4.20.

The procedure for estimating the laser parameters was applied to each record to estimate the parameters $G, \gamma_1, \gamma_2, \gamma_R, \gamma'_R$ and P . Data and best fit trajectories are shown in Figs. 4.20 and 4.21. In order to emphasise the spike region, a logarithmic abscissa is used. No rank deficiency was encountered by the algorithm. The model trajectories follow the experimental data well. Especially the spikes are reproduced well. For high currents the second peak is slightly overestimated and the relaxation oscillations that are coherent over the entire time interval, are not imitated in detail by the model trajectories.

Fig. 4.22 summarises the estimated parameters for medium and high currents. The statistical errors shown in the figure are calculated as described in Section 2.1.2. They are based on the assumption that the observational noise is uncorrelated. However, due to the low pass characteristic of the detector and the pulse merging procedure, the noise has got a more complicated correlation structure. Moreover systematic errors introduced by the preprocessing are not reflected in the estimated confidence intervals. Therefore the displayed error bars should not be over-interpreted.

For high currents the laser is continuously above the threshold as mentioned before. This qualitative difference between medium and high currents




Figure 4.20: Medium currents: measured data (**points**) and best fit trajectories (**lines**). The currents are 4.70, 4.60, 4.50 and 4.40 mA from left to right. The time axis is logarithmic in order to make the first peaks more distinct. There is some jitter in the data because the individual pulses of a record do not have exactly the same shape.

was emphasised by drawing lines between points within, but not between these two groups.

The rate constants $\gamma_1, \gamma_2, \gamma_R$ and γ'_R increase with the current, which reveals their temperature dependence. γ_1 was estimated to be much larger than γ_2 in accordance with Ciofini and Meucci (1995) and Zehnlé et al. (1992). Thus, based on dynamically modelled measured time series, the assumption $\gamma_1 = \gamma_2$ made in Arecchi et al. (1989) and Bromley et al. (1993) must be rejected. In this context the application of the simpler two-level model (Arecchi et al. 1988; Ciofini et al. 1990) turned out to be insufficient to describe the given data adequately.

The field matter coupling constant G is inversely proportional to the collisional broadening γ_\perp (Ciofini et al. 1990). G decreases with increasing current. This can be understood as a temperature-induced increase of γ_\perp . Therefore $1/G$ was plotted together with the rate constants and indeed it shows a similar behaviour.

The pump rate $\gamma_2 P$ increases monotonically with the pump current as expected. The effective number of rotational levels is not estimated directly,




Figure 4.21: High currents: measured data (**points**) and best fit trajectories (**lines**). The currents are 4.80, 4.90, 5.05 and 5.22 mA from bottom to top. The baseline increases with increasing current.

it is calculated through $Z = \gamma_R/\gamma'_R$. It increases with the current and has values between 1 and 2, in contrast to thermodynamic considerations suggesting higher values. Also, γ'_R and γ_R are considerably smaller than the values used in the literature (Christensen et al. 1969; Meyer-Bourbonneux et al. 1976; Meucci et al. 1992; Ciofini and Meucci 1995). As the gas mixture, its total pressure, the geometry of the discharge tube (internal diameter, electrode shapes, etc.) and the excitation currents are different, it is difficult to compare estimated parameters with values obtained by spectroscopic measurements. However, the low values estimated for Z are an indication of an intrinsic weakness of the 4LM which is too simplified a model to represent the complex molecular dynamics in a CO_2 laser. Taking into account the four unobserved components and the coupling between their initial values and the parameters (Section 4.3.3), it is difficult to assess the impact of model mis-specifications on the estimated parameters. Future work must be devoted to an evaluation of refined models.

The comparison between the experimental results and the outcomes of the fitting procedure remains essential. In this framework, the numerical estimations confirm the essential role played by the incoherent processes ruled by the collisional rates γ_R and γ'_R .




Figure 4.22: Estimated parameters with error bars as a function of the excitation current for medium and high currents. From top to bottom: **Top left axis:** $1/G$. **Top right axis:** rate constants γ_R , γ'_R , γ_1 and γ_2 . **Bottom left axis:** number of rotational levels Z . **Bottom right axis:** pump rate $\gamma_2 P$. $1/G$ and P contain an unknown scaling factor that is the same for all records. The results and error bars are discussed in the text.

4.5.3 Estimation of n_ξ and K_1

With the parameters of the model estimated, the laser startup process can be analysed quantitatively. The initial laser net gain $\lambda = -2K_2 + GP$ defines the slope of the straight lines that characterise the rise of the photon density in the semilogarithmic representation Fig. 4.12 on page 55. The mean initial photon number is estimated by

$$n_\xi = \langle n(0) \rangle \approx n(t)e^{-\lambda t}, \quad (4.18)$$

where $(t, n(t))$ is a point on the trajectory in the linear regime. n_ξ is smaller than 10^{-5} for all medium and low currents, except for $i=4.70$ mA where it is $5 \cdot 10^{-5}$. Thus n_ξ is eight orders of magnitudes lower than the spike intensity and the linear regime covers at least six decades of the photon density.

As a final result, as mentioned in Section 4.3.3, the steady-state condition $\dot{n} = 0$ is used to estimate the value K_1 of the cavity losses before the Q-switch for high currents. The result was between 2.42 and 2.54 kHz for all four records, i.e. approximately the same value was estimated independently in all these records.

4.6 Conclusion

In this chapter measured time series from a Q-switched CO₂ laser were modelled on the base of the four-level model (4LM). The unobserved dynamical variables of a five-dimensional differential equation scheme were constructed and the internal parameters were estimated. The method was tested on simulated data under realistic conditions and it was applied to ten records of measured data with differing excitation currents.

For low pump currents a large variability of the spike shapes was explained quantitatively through rather small variations in the pump parameter. For a wide range of higher pump currents the 4LM is able to reproduce well the large spike as well as the long tails of the relaxation phase.

The study also confirms in an unequivocal and independent way the important role of the rotational manifolds in the dynamics of a CO₂ laser. For the first time a quantitative correspondence between measured time series from a CO₂ laser and model trajectories from the four-level model was obtained.

Chapter 5

Delay Differential Equations

Delay systems are of practical importance in fields as different as engineering science and biology. Arecchi et al. (1991) and Heil et al. (2000, 2001) investigated lasers with delayed feedback. Cooke and van den Driessche (1996) treated infectious diseases models. De Gaetano and Arino (2000) modelled the human glucose metabolism and Baker et al. (1998) examined cell growth patterns, to name only some examples.

The estimation of parameters in delay differential equations (DDEs) faces the same difficulty as for ODEs, the problem of local minima. In Section 2.4.3 a multiple shooting method was developed for DDEs to solve this problem. It will be applied to two examples in Sections 5.2 and 5.3. But first, an alternative approach is examined critically.

5.1 The Delay Chain Approach

A widely used method for the approximation of a DDE is to introduce a compartment z that has an input $x_1(t)$ and decays with a rate constant $k=\tau^{-1}$ (Harrison 1995). In Fig. 5.1A the setting is drawn schematically.

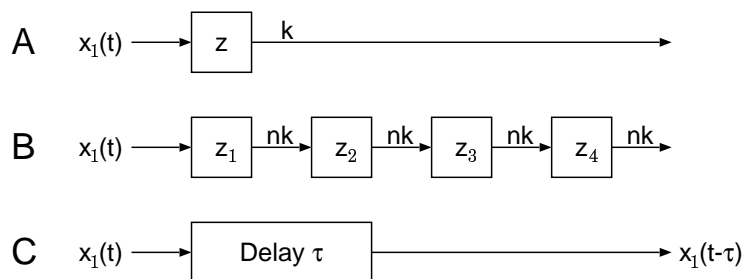


Figure 5.1: Three ways to model a delay. **A:** single compartment. **B:** delay chain, here with length $n=4$. **C:** real delay.

An example will be given in Section 5.1.3. A particle (an atom, molecule, individual, etc.) entering into z at time zero will dwell there for a time that is exponentially distributed. The mean dwell time is thus the same as in a DDE with time lag τ (Fig. 5.1C).

The method can be refined by using a *delay chain* of n compartments, each one coupled to its successor with a rate constant $k=n\tau^{-1}$ (Fig. 5.1B). The resulting overall dwell time is the sum of n independent identically distributed random variables. It has again the mean τ , but for $n \rightarrow \infty$ the distribution becomes a Gaussian with the variance scaling as n^{-1} according to the central limit theorem. Nevertheless, an essential difference from real DDEs remains. A particle can dwell in the delay chain for an arbitrary time, both for a negligible time and for a very long time. In nonlinear systems these extremes can lead to a dynamic behaviour that is not found in the corresponding DDE.

The following simulation study investigates critically, to what extent a system whose true dynamics is of DDE type, can be modelled adequately by means of a delay chain.

5.1.1 An Example from Infectious Disease Modelling

As an example an infectious disease model is chosen that was proposed to describe propagation of phocine distemper virus in North Sea seal populations during an epidemic in 1980 (Swinton et al. 1998). The model is schematically drawn in Fig. 5.2. Models of this kind are commonly referred to as *SEIR models*. The dynamic variables S, E, I, R, X represent the populations of susceptible, exposed, infectious, recovered and dead seals respectively. Exposed seals are those that are infected but not yet infectious.




Figure 5.2: Schematic diagram of the infectious disease model for seal populations. S, E, I, R, X represent the populations of susceptible, exposed, infectious, recovered and dead seals respectively. The number of infectious seals couples back into the infection flux F_1 . F_2 is a linear coupling, F_3 is a linear coupling or a delay term for the ODE model and the DDE model respectively. For the delay chain approach, I is replaced by a chain of n compartments analogously to the z compartment in Fig. 5.1.

The corresponding ODEs read

$$\dot{S} = -F_1 \quad \text{with} \quad F_1 = \beta SI/N \quad (5.1a)$$

$$\dot{E} = F_1 - F_2 \quad F_2 = \gamma E \quad (5.1b)$$

$$\dot{I} = F_2 - F_3 \quad F_3 = \nu I \quad (5.1c)$$

$$\dot{R} = F_3(1 - m) \quad (5.1d)$$

$$X = \frac{m}{1 - m} R \quad N = S + E + I + R. \quad (5.1e)$$

F_1 describes the rate of seals newly infected to the virus. It is proportional to the portion of infectious animals I/N , where N is the total living population.

The external influence leading to the outbreak of the disease is assumed to be limited to a short time interval before the first data point, being reflected in a small non-zero initial value of E . I , R and X can be assumed to be zero at $t=0$, while S contains initially the largest portion of individuals.

$$E = E_0 \quad (5.2a)$$

$$S = S_0 \quad (5.2b)$$

$$I = R = X = 0. \quad (5.2c)$$

F_2 describes the transition from exposed to infectious. Finally, F_3 describes the end of the infectious period due to recovery or death. This flux is shared out between X and R in fixed portions m and $1 - m$. The variable X can be computed statically from R . X has no influence on the system, but it is the only observable. The growth of the population is not included in the model. It can be neglected on the time scales considered.

In order to examine whether the ODE model is a reasonable description of the delay dynamics, both F_2 and F_3 can be considered to be replaced with delay terms:

$$F_2(t) = F_1(t - \tau_{\text{lat}}) \quad (5.3)$$

$$F_3(t) = F_2(t - \tau), \quad (5.4)$$

where τ_{lat} and τ are the latent and the infectious period respectively. Since the former is usually much shorter than the latter, Eq. (5.3) is expected to have a smaller impact on the overall dynamics than Eq. (5.4). For this reason and in order to simplify the following considerations, only the latter is used to replace F_3 in Eq. (5.1c) and to convert the ODE into a DDE. This model is far from being a complete and accurate description of the real situation. However, it is good enough for examining the adequacy of ODE techniques to emulate the dynamic behaviour.

5.1.2 Modelling with ODEs

A time series of length $T=150$ d (days) was simulated, using Eqs. (5.1) and (5.4) and the parameters $\beta=0.179$ d⁻¹, $\gamma=0.1333$ d⁻¹, $m=0.7$, $\tau=10$ d, $S_0=100$, $E_0=10$. The parameters were taken from a fit of the ODE model to a time series observed during the 1980s epidemic. No noise was added and the sampling interval (3 d) was chosen to be shorter than in the original data in order to underline the impact of the model mis-specification.

Then the X component of the data was analysed using the ODE model. The quasi-Newton method was used for the minimisation. Owing to large residuals, the Gauss-Newton method failed to converge frequently.

E_0 , β , γ and m were fixed to the values used for generating the data. The initial values of I and R were fixed to zero according to Eq. (5.2c). The parameter ν and the initial value of S were free parameters. They were estimated as $S_0=80.8$ and $\nu=0.0901$ d⁻¹. The resulting model trajectories are compared with the simulated curves in Fig. 5.3.

The X component of the model trajectory resembles the simulation. However, considering the fact that only F_3 is mis-specified, the correspondence is poor. The initial value of the S component was underestimated by 20%. The final value $S(T)$ is only $\frac{1}{6}$ of the true value. The other components E and I have similar amplitudes to, but different shapes than the true curves. For all components it is noticeable that in the fitted model, the maximum of the epidemic occurs considerably later than in the simulation. This shows that essential properties of the true dynamics are distorted by the approximation. One could object that the outcome of the comparison could possibly be more favourable if more than two parameters had been optimised. However, this option would not remove, but only hide the discrepancies.

5.1.3 Modelling with the Delay Chain Model

After the simple ODE model turned out to be an insufficient approximation, the delay chain approach presented above was applied to the I compartment of the infectious disease model. The dynamic variable I was replaced by a chain of n variables I_1, \dots, I_n and Eqs. (5.1c and d) were replaced by the following ODEs:

$$\dot{I}_1 = F_2 - G_1 \quad \text{with } G_i = n\nu I_i \quad (5.5a)$$

$$\dot{I}_2 = G_1 - G_2 \quad (5.5b)$$

$$\vdots$$

$$\dot{I}_n = G_{n-1} - G_n \quad (5.5c)$$

$$\dot{R} = G_n(1 - m). \quad (5.5d)$$

The number of infectious seals is given by $I = I_1 + \dots + I_n$.




Figure 5.3: Fit of the ODE model to data simulated with the DDE model. The parameters are given in the text. For the fit only the X component was used. At the beginning the curves lack smoothness because of the low sampling frequency. The R components of the simulation and the fit are not shown since R is proportional to X . Result: S is considerably underestimated and the maximum of the epidemic occurs later than in the simulation.




Figure 5.4: Same simulation as in Fig. 5.3, fitted with the delay chain model with $n=4$. Result: though the observable is well reproduced, S is still clearly underestimated. The E and I components correspond much better to the simulated curves than for $n=1$.




Figure 5.5: Final value of the S compartment as a function of the delay chain length. **Open circles:** estimated $S(T)$. **Broken line:** true value. It is slowly approached. **Solid line:** phenomenological formula Eq. (5.6).

For $n=1$ the model is identical to the simple ODE model used above. The analysis described in the preceding section was repeated with this model for various lengths of the delay chain. With increasing n the correspondence between simulation and fit trajectories becomes progressively better. As an example the outcome for $n=4$ is shown in Fig. 5.4. Though the observable is well reproduced, S is still significantly underestimated.

One of the details of the true dynamics that is not reproduced well, is the number of susceptible seals at the end of the epidemic, $S(T)$. An interesting question is how long the chain must be in order to reproduce all such details with satisfying accuracy. In Fig. 5.5, $S(T)$ is shown as a function of the chain length. As expected, it approaches the true value $S_{\text{true}}(T)$ with increasing n . However, a significant discrepancy exists even for $n=25$, which is far beyond the limits of practical usefulness.

A phenomenological analysis shows that the error $S_{\text{true}}(T) - S(T)$ can be described almost perfectly as the reciprocal of a linear function of n :

$$S(T) \approx S_{\text{true}} - \frac{32.52}{n + 0.6652} \quad (5.6)$$

5.1.4 Discussion

This section demonstrated that the delay chain technique can be used to approximate a DDE containing a moderate nonlinearity by an ODE model. However, the model mis-specification leads to systematic errors that do not

diminish rapidly with increasing length of the chain. A final evaluation of the usefulness of this method depends on the specific system one wants to examine. For highly nonlinear systems it is presumably not suited. On the other hand, biological systems with time lags like the one used as an example here, could be candidates. Typically they are not pure delay systems, since the time lags vary between individual subjects. The delay element is rather of integro-differential type which is something in between ODEs and DDEs from the viewpoint of dwell time distributions. Therefore the delay chain approach is possibly less harmful in some of these systems. Nevertheless the impact of such approximations should be monitored carefully.

5.2 Estimating Parameters of the Mackey Glass System

In this section, the algorithm for estimating parameters in DDEs developed in Section 2.4 is applied to the well-known Mackey-Glass equation

$$\dot{x} = \frac{ax_\tau}{1 + x_\tau^c} - bx. \quad (5.7)$$

It was originally introduced as a model of blood generation for patients with leukemia (Mackey and Glass 1977). Since then it has become a standard system for the exploration of high-dimensional chaos and for the evaluation of methods of chaotic time-series analysis (Kittel et al. 1995; Farmer 1982).

Stationary time series were simulated using Eq. (5.7) with the standard parameters $a=0.2$, $b=0.1$, $c=10$, $\tau=30$ and the sampling interval $\Delta t=1$. The dynamics was initialised with a constant initial curve $h_0(t)=0.8$. A transient period of length 10000 at the beginning was skipped. Fig. 5.6 shows a noise-free time series of length 5000.




Figure 5.6: Time series of the Mackey-Glass system. Parameters $a=0.2$, $b=0.1$, $c=10$, time lag $\tau=30$.

White Gaussian noise with standard deviation $\sigma=0.14$ was added to the data, corresponding to a noise level of 50%. At this noise level, negative data points occur spuriously and cause instabilities when they are used to get starting guesses for the initial curves. Therefore the state vector was forced to be positive by means of inequality constraints as described on page 6.

5.2.1 Known Time Lag

The time lag is more difficult to estimate than the other parameters since it has a complex influence on the dynamics. Therefore this challenge is deferred to Section 5.2.2. For now, τ is assumed to be known.

The parameters a , b and c were estimated from a segment of 1000 data points. The multiple shooting method was applied with 33 subintervals, each having a spline segment of length 30, containing 11 spline knots, placed at every third data point. The data segments had the same length. So, the constants defined in Section 2.4.3 are $d=3$, $K=10$, $D=Kd=30$ and $M=33$. The starting guesses for the parameters were set to ten-times the true values.

Fig. 5.7 shows three stages of the iterative optimisation procedure. The initial situation is plotted in panel A. The trajectories of the individual multiple shooting intervals are simply overlapping, so that at each point except those at the boundaries, two trajectories can be seen. The cubic splines describing the initial curves were set up to go through the data points at the knots. Then the DDE was integrated over each data segment. In the first phase, the spline variables were held fixed and the parameters were optimised. Panel B shows the situation when the algorithm had converged for the first time after 7 iterations. The true trajectory is already imitated well. At this stage the estimated parameters were $a=0.189$, $b=0.0910$, $c=8.97$. In the second phase the spline variables were fitted together with the parameters. After 11 additional iterations the algorithm converged finally (panel C). Now the model trajectories seem to be perfectly continuous. That means that the spline knots are dense enough to give a reasonable representation of the rear segments of the preceding trajectories. The final estimates of the parameters were $a=0.199$, $b=0.100$ and $c=9.93$, i.e. , no parameter deviated by more than 1% from its true value.

In Fig. 5.8A the fit trajectory is compared with the true trajectory. They are virtually identical over the entire interval, except for the beginning. In order to inspect the discrepancies in more detail, the difference between the two curves is plotted in panel B. The errors are largest in the very first spline segment. The spline variables of the first segment can not be estimated with high accuracy because the subsequent dynamics is not sensitive to all linear combinations of them, as mentioned at the end of Section 2.4.3.




Figure 5.7: Estimation of the parameters a , b and c from a noisy time series of the Mackey-Glass system for known time lag. **Points:** simulated data. True parameters: $a=0.2$, $b=0.1$, $c=10$, time lag $\tau=30$, noise level: 50%, sampling interval: $\Delta t=1$. **Lines:** model trajectories. At each time two trajectories are overlapping. **A:** initial situation. Starting guesses of the parameters: $a=2$, $b=1$, $c=100$. **B:** end of the first phase, in which the spline variables were fixed, after 7 iterations. Parameters at this time: $a=0.189$, $b=0.0910$, $c=8.97$. **C:** final solution after 11 additional iterations, virtually continuous. Estimated parameters: $a=0.199$, $b=0.100$ and $c=9.93$.




Figure 5.8: **A:** comparison of the estimated trajectory shown in Fig. 5.7C (**bold broken line**) with the true trajectory (**thin solid line**). **B:** difference between the two trajectories. The largest errors occur in the very first spline segment.

5.2.2 Unknown Time Lag

Up to now the time lag τ was assumed to be known. This is not unreasonable since various concepts for the estimation of the delay time have been reported (Kember and Fowler 1993; Bünner et al. 1996; Ellner et al. 1997; Voss and Kurths 1997, 1999). Nevertheless it is desirable to estimate τ in the same way as the other parameters. Mathematically this is not a problem; the sensitivity equations have only to be extended by the extra term in Eq. (2.17). However, the dynamics is much more sensitive to τ than to the other parameters. The intrinsic oscillations of the time series translate into fluctuations of the objective function with respect to τ . When the starting guess of τ deviates from the true time lag by an amount that is comparable to typical time scales of the system, then the parameter vector is likely to be separated from the true vector by a barrier of the cost function. The sensitivity equations are themselves a DDE with the time lag τ . When this time lag is mis-specified, they provide misleading information for the minimisation process. The situation is similar to that of the initial-value approach. There the multiple shooting method solved the problem by dividing the fit

interval and thereby shortening the length of continuous trajectories. In the case of DDEs this would not help, since it does not make sense to make the trajectories shorter than the time lag.

Despite these problems, the estimation of τ works remarkably well. The simulation made in the preceding section was repeated under the same conditions except that τ was set to $30 - \frac{\pi}{4} = 29.2146$, in order to demonstrate that time lags that are not multiples of the sampling interval can be estimated as well. Again 50% measurement noise was added and the four parameters were estimated by the multiple shooting method. The starting guesses for the parameters a , b and c were again set to ten-times the true values, while τ was set to 60. Accordingly the parameters for the setup of the multiple shooting intervals were $d=3$, $K=20$, $D=Kd=60$ and $M=16$. The convergence in phase one was reached after 29 iterations, the final convergence after 14 additional iterations. The estimated parameters were $a=0.195$, $b=0.0985$, $c=9.97$ and $\tau=29.20$. The fact that τ is a delicate parameter is reflected in its extraordinarily high accuracy. Its relative error is only 0.05%. True tra-




Figure 5.9: Estimation of all four parameters from noisy data. **A:** simulated data (points), true trajectory (thin solid line) and model trajectory (bold broken line). Only slight discrepancies are visible in the first few oscillations. **B:** difference between the two trajectories.

jectory, data and model trajectory can be seen in Fig. 5.9. Again a very good correspondence between model and true trajectory is achieved.

5.2.3 Systematic Test with Random Starting Guesses

The results presented in the preceding sections are only examples. They are not intended to claim that the method would succeed in estimating the parameters with high accuracy in all instances. To evaluate the performance and the reliability of the method, a systematic simulation study is made in this section.

One thousand different noisy time series of length 500¹ were generated with the parameters $a_0=0.2$, $b_0=0.1$, $c_0=10$, and a time lag $\tau_0=29.2146$. For the simulation, a constant initial curve $h_0(t)=x_0$ was used and a transient period of length 10000 was skipped as before, but now x_0 was a random number, uniformly distributed in $[0:1]$. Thereby not only the noisy time series but also the true trajectory was different in each pass. After the simulation the algorithm for estimating the parameters was applied as demonstrated above. The starting guesses for the parameters were drawn randomly from the following probability densities:

$$\rho_a(a) = \frac{1}{a_0} \exp\left(-\frac{a}{a_0}\right) \quad \text{for } a > 0 \quad (5.8a)$$

$$\rho_b(b) = \frac{1}{b_0} \exp\left(-\frac{b}{b_0}\right) \quad \text{for } b > 0 \quad (5.8b)$$

$$\rho_c(c) = \frac{1}{c_0 - 1} \exp\left(-\frac{c - 1}{c_0 - 1}\right) \quad \text{for } c > 1 \quad (5.8c)$$

$$\rho_\tau(\tau) = 0.1 \quad \text{for } 25 \leq \tau < 35. \quad (5.8d)$$

The exponential distributions of a and b allowed large deviations from the true values. The parameter c must not be greater than one, otherwise the sensitivity equations would be divergent. Therefore the exponential distribution was modified to satisfy this constraint. The expectation values for these three distributions are the respective true parameters. The starting guess for τ was uniformly distributed in the interval $I_\tau = [25; 35]$.

The method succeeded in estimating the parameters within 50 iterations in 852 out of 1000 trials. For the other realizations it failed to converge within this limit. Considering the rather adverse conditions this is a good result. Fig. 5.10 shows all starting guesses that were used and indicates which of them led to a successful fit and which not. There is no indication of a correlation between the starting guess and the success of the fit.

The estimated parameters from the successful trials are summarised in Fig. 5.11. Only a single real outlier can be seen at index 624, i.e. the

¹ The length was reduced to save computation time. The entire simulation study took about 10 hours on a 650 MHz Pentium III.




Figure 5.10: Visualisation of the radius of convergence of the method. Each point in these figures displays the starting guesses for one of one thousand realizations. The true parameter values are indicated with lines. **Upper panel:** parameters a and b . **Lower panel:** parameters c and τ . Due to the exponential distributions, very large parameter values occurred. However, they seem not to have led to significantly more unsuccessful fits.

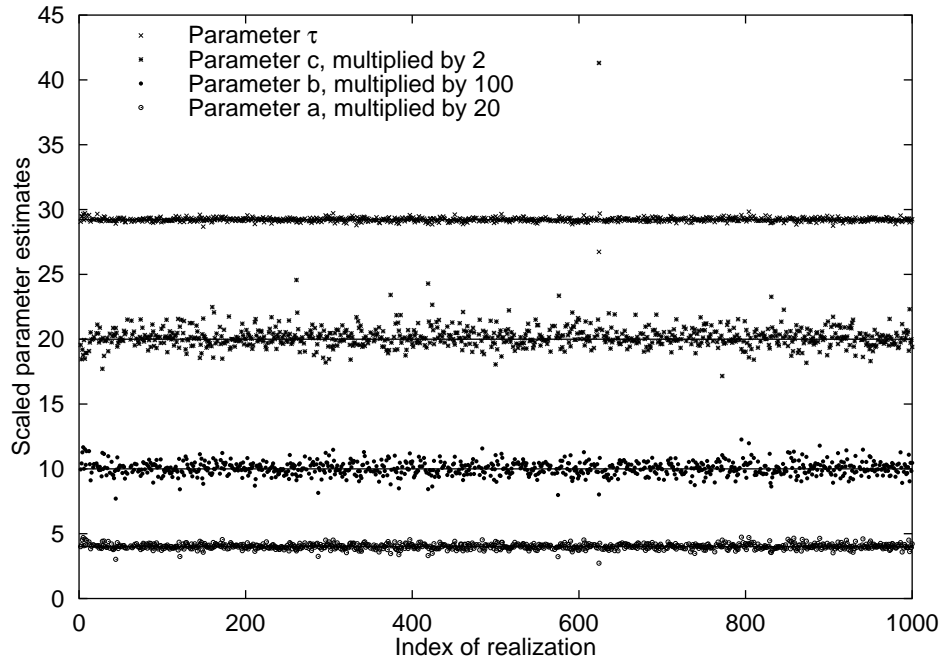


Figure 5.11: Summary of estimated parameters. **Symbols:** estimated parameters as indicated in the legend. **Solid lines:** true values.

algorithm rarely signals convergence when it is actually trapped in a clear local minimum. This is an important advantage of the method. A missing result is much better than a misleading result. It offers the opportunity to try different starting guesses or to examine another segment of the time series, if available. Table 5.1 shows the mean estimated parameters and their standard deviation. All parameters were estimated with a relative accuracy of 5%, except τ which was determined with a ten-times better precision. The reason is the high sensitivity of the dynamics to changes in τ . So the investment made in an accurate starting guess for τ is paid back in the form of an extraordinarily high precision of the estimate.

Parameter	True value	Mean \pm Standard deviation	rel.error
a	0.2	0.2009 ± 0.010	5%
b	0.1	0.1004 ± 0.0054	5%
c	10	10.09 ± 0.52	5%
τ	29.2146	29.22 ± 0.16	0.5%

Table 5.1: Mean estimated parameters and standard deviations, calculated from 852 successful applications of the method.

Parameter	True value	Mean	± Standard deviation	rel.error
a	0.2	0.2001	± 0.002	1%
b	0.1	0.1001	± 0.001	1%
c	10	10.005	± 0.07	0.7%
τ	29.2146	29.2146	± 0.025	0.09%

Table 5.2: Same as Table 5.1 for a noise level of 10%. Taking all realizations together, τ could be determined with a relative accuracy of 0.003%.

The simulation study was repeated with a noise level of 10% (see Table 5.2). In this case only for three of one thousand realizations the procedure failed. The precision of the estimated parameters is even higher, as expected.

When the range I_τ of starting guesses for τ was chosen wider, the method failed to converge for significantly more realizations. If the time lag had a large uncertainty in a realistic setting, one would apply the method for several different starting guesses. The true time lag has a rather large “basin of attraction” in which the starting guess may be located in order to converge to the true value. Therefore a rather small number of different time lags must be tried. Most other methods have to scan the entire range of possible values.

5.2.4 Summary

In this section the multiple shooting procedure developed in Section 2.4 was applied to the Mackey-Glass system with time lag 30. For rather high levels of observation noise it yields very accurate estimates of all parameters and the underlying true trajectory is well reconstructed. The method rarely signals convergence when the result is clearly wrong, and it converges to the true parameters from rather inaccurate starting guesses. In particular, the starting guesses for a , b and even c had often even a different order of magnitude from the respective true parameters. The time lag was estimated with an accuracy of 0.5% from 500 data points with 50% measurement noise. It does not need to be a multiple of the sampling interval and the starting guess may deviate considerably from the true value.

5.3 Application to Measured Data

As a final example the method is applied to a time series from a chaotic electronic oscillator. The experimental setup is drawn in Fig. 5.12. It is part of an experiment realizing a method to anticipate future states of nonlinear time-delayed feedback systems (Voss 2001). A transistor and an amplifier are used to construct a nonlinear response function f . The low-pass filtered output (x) is fed into a delay element coupled back into the




Figure 5.12: Schematic diagram of the electronic circuit used by Voss (2001). The low-pass filter (R_1 and C) and the delay line are inspired by Kittel et al. (1998). The nonlinearity is built up of the transistor T_1 , the adjustable amplifier OA_1 and the resistors R_2 - R_7 . **Electronic components:** Delay line: bucket brigade line MN 3011 with 3328 stages, triggered by MN 3101 (both National Panasonic); OA_1 : LM 324N; $C=660$ pF; $R_1=470$ k Ω , $R_2=100$ k Ω lin., $R_3=22$ k Ω , $R_4=4.7$ k Ω , $R_5=10$ k Ω , $R_6=1$ k Ω , $R_7=47$ k Ω , T_1 : BC 238C.

input (x_τ). The dynamical model reads

$$\dot{x} = -\alpha x + f(x_\tau). \quad (5.9)$$

The nonlinear response function f is parameterised by the third order polynomial

$$f(x) = a_0 + a_1 x + a_2 x^2 + a_3 x^3 \quad (5.10)$$

A time series of length 100 ms was recorded with a sampling interval of 0.01 ms. A glance at the time series shows that it is highly oversampled. The sampling interval pretends to provide information that actually does not exist in the time series due to low-pass filtering. Therefore only every 10th data point was used for the analysis in the following. Fig. 5.13 shows that this is sufficient to reproduce all relevant features of the original time series. The parameter estimates did not change by more than 0.03% due to the reduction.

The dynamical parameters are the time lag τ , the damping factor α and the coefficients a_i . No a priori information about the parameters was utilised. The most difficult task is the coarse estimation of the delay time. The time series shows oscillations with a period of about 4 ms, indicating a time lag in the millisecond range. Therefore the delay parameter was scanned from 1 ms to 20 ms in steps of 1 ms. Each value was used as a starting guess τ_0 and the estimation procedure was carried out as described earlier. The starting guesses for all other parameters were simply set to




Figure 5.13: Clipping of the time series showing that it is oversampled. **Dots:** original data. **Circles:** data points used for fitting.

zero. The length of the spline segments was chosen 5 ms longer than τ_0 in order to allow for some variation of τ .

For $\tau_0=12$ and 13 ms the method converged within 20 iterations to the values $\tau=13.3$ ms, $\alpha=3.7$ kHz, $a_0=0.8864$ V kHz, $a_1=-3.982$ kHz, $a_2=-15.26$ V⁻¹kHz and $a_3=-11.54$ V⁻²kHz. For other starting guesses no convergence was achieved or the final value of the objective function was more than 100 times higher than for the successful fits.

Having found good estimates for all parameters, the procedure was started again without multiple shooting and with the length of the spline segment adjusted to τ . The final estimates are $\tau=13.28$ ms, $\alpha=3.443$ kHz, $a_0=0.8332$ VkHz, $a_1=-3.775$ kHz, $a_2=-14.36$ V⁻¹kHz and $a_3=-10.67$ V⁻² kHz. The model trajectory is shown together with the time series in Fig. 5.14. Hardly any difference can be seen in the entire interval. In order to facilitate a comparison, the attractor of the system is reconstructed by a delay embedding in Fig. 5.15, both for the measured time series and for the model trajectory. The model attractor resembles well the attractor reconstructed from the data, except for some features that are visible in the right-hand image, but hidden by noise on the left. Thus the modelling procedure achieves even more than a perfect imitation of the measured data: it performs noise reduction.




Figure 5.14: Comparison of best fit trajectory (solid line) and observed time series (points).




Figure 5.15: Phase space reconstruction of the attractor of the system. **A:** from observed time series. **B:** from model trajectory.

Chapter 6

Summary

This thesis examined methods for estimating parameters in deterministic dynamical systems. The maximum likelihood approach that was chosen in view of its superior asymptotic properties, facilitated a common description of the parameter estimation problem for some rather different classes of systems.

First of all time-discrete systems were regarded in order to plumb the strengths and limits of the maximum likelihood approach. The logistic map was used as a well studied canonical example. The essential point in the maximum likelihood approach is that it takes into account the information that there is a true trajectory underlying the measurements. The difficulty is to find it without being trapped in a local minimum of the objective function in the parameter space. The multiple shooting technique, developed by Bock for the class of ordinary differential equations (ODEs), was transferred to time-discrete systems in order to circumvent this problem.

This technique turned out to be of particular importance for the logistic map since the conventional initial-value approach failed to converge to the global minimum when the starting guess for the unknown parameter deviated slightly from its true value. This property of the logistic map could be understood as a consequence of the heavily jagged objective function. The expectation that this characteristic of chaotic maps could be useful to obtain accurate estimates of the parameters, was confirmed in the case of a single variable being estimated from noisy time series. The standard deviation of the estimate scaled as $e^{-\lambda N}$, where λ is the Lyapunov exponent of the map and N is the length of the time series. However, when both the initial value of the trajectory and the dynamic parameter were estimated, which is the realistic setting, then the exponential law changed to a $1/N$ -law. This is still better than the $1/\sqrt{N}$ scaling achieved with conventional methods that ignore the underlying dynamic nature of the data.

Finally the strong correlation between the initial value and the parameter of the logistic map could be understood in terms of a highly discontinuous

self-similar function that transforms the states of the system so that they fulfil the same dynamic equation for a different parameter.

In contrast to the pedagogical example of the logistic map, Chapter 4 described a practical application of the parameter estimation method. Measured data from a CO₂ laser were modelled by a nonlinear system of five differential equations derived from first principles. It contains six unknown parameters and four unobserved components. The analysis of measured data involves difficulties and subtleties that are not found with simulated time series. Among them are non-stationarities, limits of the detector system and model mis-specifications. The analysis underlined that the observation equation is as important a part of the model as the dynamical equations. If the observation process is modelled incorrectly, unpredictable errors in the fitting procedure can result.

A meticulous investigation of the time series recorded under two different conditions revealed that the detector circuit used in the experiment had a significant low-pass characteristic in addition to its nonlinearity. These properties could be determined in a calibration procedure. They were used to choose a proper pre-processing procedure and to define the observation equation.

Another important element in the modelling procedure was the reduction of the number of fit variables by means of equality constraints that were derived from an analytical description of the laser startup process. These constraints that represent physical knowledge about the laser, related the initial values of the unobserved components to the initial photon number and the dynamic parameters. In this way the unobserved dynamical variables were constructed and the internal parameters of the laser were estimated. The method was tested on simulated data under realistic conditions and applied to ten records of measured data with differing excitation currents.

For low pump currents, a large variability of the pulse shapes could be explained through rather small variations in the pump parameter. For a wide range of higher pump currents the four-level model was confirmed as an appropriate model to describe the given data. The study confirmed the important role of the rotational manifolds in the dynamics of a CO₂ laser. For the first time a quantitative correspondence between measured time series from a CO₂ laser and model trajectories from the four-level model was obtained.

The last chapter was devoted to delay differential equations (DDEs). This class of dynamical systems is used to describe a large variety of applications in physics, biology, medicine and engineering science, among others. The delay chain approach to approximate a DDE with an ODE was shown to be practical to some extent but to lead to systematic errors in the result that vanish only as n^{-1} with increasing length of the delay chain. Therefore it is recommended to use a DDE approach to model delay systems.

Next the estimation of parameters of a canonical DDE example, the

Mackey-Glass equation, was investigated. For this purpose a multiple shooting technique for DDEs was developed. Due to the theoretically infinite phase space of a DDE, there is no straightforward approach to transfer the concept of the continuity constraints to DDEs. The proposed procedure parameterises the initial curve of each subinterval by cubic splines and forces the projection of the solution of the preceding interval onto the space of spline functions to be identical to them.

Using this procedure, all parameters of the Mackey-Glass system with time lag 30 could be estimated for rather high levels of observation noise. Moreover, the method has a large radius of convergence, i.e. the starting guesses for the parameters were allowed to deviate from the true values by a factor of ten, except for the delay parameter that was allowed to vary between 25 and 35.

The time lag does not need to be a multiple of the sampling interval, as in many other methods. Since Chapter 3 showed that a disadvantage of the maximum likelihood approach is the extremely complex structure of the objective function, a simulation study was performed evaluating the actual reliability of the method. It turned out to yield reliable and precise estimates of all parameters for fairly high noise levels.

Finally, the method was applied to measured data from an electronic circuit, demonstrating that it is well suited to yield accurate estimates of the parameters of an experimental delayed-feedback system.

The examples presented in this thesis have demonstrated the power of the maximum likelihood method to extract as much information as possible from measurements of a dynamical system. Parameters are estimated with the highest precision and physical properties of the underlying system are identified. This approach of investigating dynamical systems by modelling “in vivo” experiments is expected to enable detailed insight into numerous challenging complex phenomena.

Appendix A

Equivalence of Logistic Map, Tent Map and Shift Map

This appendix provides explicit transformations between two versions f_x and f_z of the logistic map (with special parameters), the tent map f_t and the shift map f_s . All maps, their variables and the relations between them are depicted in the following commutative diagram:

$$\begin{array}{ccccccc}
 x & \xleftarrow{h_x} & z & \xleftarrow{h_z} & t & \xrightarrow{h_s} & s \\
 \downarrow f_x & & \downarrow f_z & & \downarrow f_t & & \downarrow f_s \\
 x' & \xleftarrow{h_x} & z' & \xleftarrow{h_z} & t' & \xrightarrow{h_s} & s'
 \end{array}$$

h_x , h_z and h_s are invertible transformations. Define the logistic map in the form used in Chapter 3

$$x' = f_x(x) = 1 - ax^2, \quad a \in [1.5; 2], \quad (\text{A.1})$$

the logistic map in its standard form

$$z' = f_z(z) = rz(1 - z), \quad r \in [1; 4], \quad (\text{A.2})$$

the tent map

$$t' = f_t(t) = \begin{cases} 2t & 0 \leq t < 0.5 \\ 2 - 2t & 0.5 \leq t < 1, \end{cases} \quad (\text{A.3})$$

and the shift map

$$s' = f_s(s) = \begin{cases} 2s & 0 \leq s < 0.5 \\ 2s - 1 & 0.5 \leq s < 1. \end{cases} \quad (\text{A.4})$$

The two versions of the logistic map are related through the transformations

$$\begin{aligned} a = \frac{r}{4}(r-2) \quad & \text{and} \quad x = h_x(z) = \frac{4z-2}{r-2} \\ & \text{or} \\ r = 1 + \sqrt{1+4a} \quad & \text{and} \quad z = h_x^{-1}(x) = \frac{-1 + \sqrt{1+4a}}{4}x + \frac{1}{2}, \end{aligned}$$

respectively.

The similarity to the tent map holds only for $r=4$ or $a=2$ respectively. The transformation is (Ulam and von Neumann 1947):

$$z = h_z(t) = \sin^2\left(\frac{\pi}{2}t\right).$$

Finally, the tent map is related to the shift map through the operation

$$s = t \text{ xor } \frac{t}{2},$$

where “xor” means the exclusive or of the binary representations of the operands, i.e. , when

$$s = \sum_{j=1}^{\infty} s_j 2^{-j}, \quad s_j \in \{0, 1\} \quad \text{and} \quad (\text{A.5})$$

$$t = \sum_{j=1}^{\infty} t_j 2^{-j}, \quad t_j \in \{0, 1\} \quad \text{then} \quad (\text{A.6})$$

$$s_j = \begin{cases} t_1 & \text{if } j = 1 \\ 0 & \text{if } j > 1 \text{ and } t_j + t_{j-1} \text{ is even} \\ 1 & \text{if } j > 1 \text{ and } t_j + t_{j-1} \text{ is odd.} \end{cases} \quad (\text{A.7})$$

Since this transformation is somehow pathological, it should not be taken too seriously. However, it shows that the properties of the shift map described in Section 3.3 have some relevance in the context of the logistic map.

Acknowledgements

I would like to thank

- Prof. Dr. J. Honerkamp for giving an experimentalist the opportunity to graduate in the field of data analysis.
- PD Dr. Jens Timmer for the orientation towards practical applications as opposed to toy models and for the lessons on the secondary virtues of a scientist. *whoa*
- Henning Voss for helpful hints regarding parameter estimation in maps and for the electronic-circuit data presented in Section 5.3.
- Chris Bertram, for proof-reading my disastrous “sentences” (not this page), thereby giving me lessons in good English.
- Thomas Müller and Wolfgang Huber for reminding me of the aim.
- Thorsten Müller and Bernd Kappler for the long easy time in a nice room on the 9th floor.
- Henning Rust, Kim Boström, Douglas Maraun, Daniel Faller, and all the colleagues of the Centre for Data-Analysis and Modelling and the Scientific Consulting Group for a really good time.
- the developers of Linux, L^AT_EX, gnuplot and tgif for having created a productive environment that did not crash in the decisive moment.

Bibliography

- Arecchi, F. T., Gadomski, W., Meucci, R., and Roversi, J. A. (1988). Swept dynamics of a CO₂ laser near threshold: two- versus four-level model. *Opt. Commun.*, 65:47–51.
- Arecchi, F. T., Gadomski, W., Meucci, R., and Roversi, J. A. (1989). Dynamics of laser buildup from quantum noise. *Phys. Rev. A*, 39:4004–4015.
- Arecchi, F. T., Giacomelli, G., Lapucci, A., and Meucci, R. (1991). Dynamics of a CO₂ laser with delayed feedback: The short-delay regime. *Phys. Rev. A*, 43:4997–5004.
- Arecchi, F. T., Meucci, R., Puccioni, G., and Tredicce, J. (1982). Experimental evidence of subharmonic bifurcations, multistability, and turbulence in a Q-switched gas laser. *Phys. Rev. Lett.*, 49:1217.
- Baake, E., Baake, M., Bock, H. G., and Briggs, K. M. (1992). Fitting ordinary differential equations to chaotic data. *Phys. Rev. A*, 45(8):5524–5529.
- Baake, E. and Schlöder, J. (1992). Modeling the fast fluorescence rise of photosynthesis. *Bull. Math. Biol.*, 54:999–1021.
- Baker, C. T. H., Bocharov, G. A., Paul, C. A. H., and Rihan, F. A. (1998). Modelling and analysis of time-lags in some basic patterns of cell proliferation. *J. Math. Biol.*, 37:341–371.
- Balestri, S., Ciofini, M., Meucci, R., and Arecchi, F. T. (1991). CO₂ laser with swept pump parameter: The nonlinear regime. *Phys. Rev. A*, 44:5894–5897.
- Balle, S., Colet, P., and San Miguel, M. (1991). Statistics for the transient response of single-mode semiconductor laser gain switching. *Phys. Rev. A*, 43:498–506.
- Balle, S., San Miguel, M., Abraham, N. B., Tredicce, J. R., Alvarez, R., D’Angelo, E. J., Gambhir, A., Thornburg, K. S., and Roy, R. (1994). Transient in multivariable dynamical systems depend on which parameter is switched as illustrated in lasers. *Phys. Rev. Lett.*, 72:3510–3513.

- Bär, M., Hegger, R., and Kantz, H. (1999). Fitting partial differential equations to space-time dynamics. *Phys. Rev. E*, 59:337–342.
- Bock, H. G. (1981). Numerical treatment of inverse problems in chemical reaction kinetics. In Ebert, K., Deuffhard, P., and Jäger, W., editors, *Modelling of Chemical Reaction Systems*, pages 102–125. Springer, Heidelberg and New York.
- Bock, H. G. (1983). Recent advances in parameter identification techniques for O.D.E. In Deuffhard, P. and Hairer, E., editors, *Numerical Treatment of Inverse Problems in Differential and Integral Equations*, pages 95–121. Birkhäuser, Basel.
- Bock, H. G. (1987). *Randwertproblemmethoden zur Parameteridentifizierung in Systemen nichtlinearer Differentialgleichungen*. PhD thesis, Universität Bonn. Bonner Mathematische Schriften Nr. 183.
- Brenan, K. E., Campbell, S. L., and Petzold, L. R. (1989). *The Numerical Solution of Initial Value Problems in Differential-Algebraic Equations*. North-Holland, New York.
- Bromley, D., D’Angelo, E. J., Grassi, H., Mathis, C., and Tredicce, J. R. (1993). Anticipation of the switch-off and delay of the switch-on of a laser with a swept parameter. *Opt. Commun.*, 99:65–70.
- Bünner, M., Popp, M., Meyer, T., Kittel, A., Rau, U., and Parisi, J. (1996). Recovery of scalar time-delay systems from time series. *Phys. Lett. A*, 211:345–349.
- Carroll, R. J., Ruppert, D., and Stefanski, L. A. (1995). *Measurement Error in Nonlinear Models*. Chapman and Hall, London.
- Casdagli, M. (1991). Chaos and deterministic versus stochastic nonlinear modeling. *J. Roy. Stat. Soc. B*, 54:303–328.
- Christensen, C. P., Freed, C., and Haus, H. A. (1969). Gain saturation and diffusion in CO₂ lasers. *IEEE J. Quantum Electron.*, 5:276–283.
- Ciofini, M. and Meucci, R. (1995). Determination of the effective number of rotational levels affecting the dynamics of CO₂ lasers. *IEEE J. Quantum Electron.*, 31:886–893.
- Ciofini, M., Meucci, R., and Arcchi, F. T. (1990). Transient statistics in a CO₂ laser with a slowly swept pump. *Phys. Rev. A*, 42:482–486.
- Ciofini, M., Politi, A., and Meucci, R. (1993). Effective two-dimensional model for CO₂ lasers. *Phys. Rev. A*, 48:605–610.

- Coca, D. and Billings, S. A. (2000). Direct parameter identification of distributed parameter systems. *Int. J. Syst. Sci.*, 31:11–17.
- Cooke, K. L. and van den Driessche, P. (1996). Analysis of an SEIRS epidemic model with two delays. *J. Math. Biol.*, 35:240–260.
- Corrêa, M. V., Aguirre, A., and Mendes, E. M. A. M. (2000). Modeling chaotic dynamics with discrete nonlinear rational models. *Int. J. Bifurcat. Chaos*, 10:1019–1032.
- Cox, D. and Hinkley, D. (1994). *Theoretical Statistics*. Chapman & Hall, London.
- Cremers, J. and Hübler, A. (1987). Construction of differential equations from experimental data. *Z. Naturforsch.*, 42a:797–802.
- Crutchfield, J. and McNamara, B. (1987). Equations of motion from data series. *Complex Systems*, 1:417–452.
- De Gaetano, A. and Arino, O. (2000). Mathematical modelling of the intravenous glucose tolerance test. *J. Math. Biol.*, 40:136–168.
- Dupré, J., Meyer, F., and Meyer, C. (1975). Influence des phénomènes de relaxation sur la forme des impulsions fournies par un laser CO₂ déclenché par un absorbant saturable. *Rev. Phys. Appl.*, 10:285–293.
- Ellner, S. P., Kendall, B. E., Wood, S. N., McCauley, E., and Briggs, C. J. (1997). Inferring mechanism from time-series data: Delay-differential equations. *Phys. D*, 110:182–194.
- Farmer, J. D. (1982). Chaotic attractors of an infinite-dimensional dynamical system. *Phys. 4D*, pages 366–393.
- Gasser, T., Müller, H.-G., and Mammitzsch, V. (1985). Kernels for non-parametric curve estimation. *J. Roy. Stat. Soc. B*, 47:238–252.
- Gill, P. E., Murray, W., and Wright, M. H. (1981). *Practical Optimization*. London: Academic Press.
- Giona, M., Lentini, F., and Cimagalli, V. (1991). Functional reconstruction and local prediction of chaotic time-series. *Phys. Rev. A*, 44:3496 – 3502.
- Gombert, A., Glaubitt, W., Rose, K., Dreiholz, J., Zanke, C., Blasi, B., Heinzl, A., Horbelt, W., Sporn, D., Doll, W., Wittwer, V., and Luther, J. (1998a). Glazing with very high solar transmittance. *Solar Energy*, 62:177–188.
- Gombert, A., Rose, K., Heinzl, A., Horbelt, W., Zanke, C., Blasi, B., and Wittwer, V. (1998b). Antireflective submicrometer surface-relief gratings for solar applications. *Solar Energy Materials and Solar Cells*, 54:333–342.

- Gouesbet, G., le Sceller, L., Letellier, C., Brown, R., Buchler, J. R., and Kolláth, Z. (1996). Reconstructing a dynamics from a scalar time series. *Ann. N.Y. Acad. Sci.*, 808:25–50.
- Grassi, H., Green, C., Hoffer, L. M., Lippi, G. L., and Tredicce, J. R. (1994). Experimental study of the turn-on statistics of class-B lasers under the influence of additive noise. *Phys. Rev. A*, 50:805–812.
- Hairer, E., Nørsett, S. P., and Wanner, G. (1987). *Solving Ordinary Differential Equations I*. Springer, Berlin.
- Hanson, R. J. and Haskell, K. H. (1982). Two algorithms for the linearly constrained least squares problem. *ACM Trans. Math. Software*, 8:323–333.
- Harrison, G. W. (1995). Comparing predator-prey models to Luckinbill’s experiment with *Didinium* and *Paramecium*. *Ecology*, 76:357–374.
- Hegger, R., Kantz, H., Schmüser, F., Diestelhorst, M., Kapsch, R., and Beige, H. (1998). Dynamical properties of a ferroelectric capacitor observed through nonlinear time series analysis. *Chaos*, 8:727.
- Heil, T., Fischer, I., and Elsässer, W. (2000). Stabilization of feedback-induced instabilities in semiconductor lasers. *J. Opt. B*, 2:413–420.
- Heil, T., Fischer, I., Elsässer, W., Mulet, J., and Mirasso, C. R. (2001). Chaos synchronization and spontaneous symmetry-breaking in symmetrically delay-coupled semiconductor lasers. *Phys. Rev. Lett.*, 86:795–798.
- Honerkamp, J. (1994). *Stochastic Dynamical Systems*. VCH, Weinheim.
- Horbelt, W., Müller, T., Timmer, J., Melzer, W., and Winkler, K. (2000). Analysis of nonlinear differential equations: parameter estimation and model selection. In Brause, R. and Hanisch, E., editors, *Medical Data Analysis*, volume 1933 of *Lecture Notes in Mathematics*, page 152. Springer, Berlin.
- Horbelt, W., Timmer, J., Bünner, M., and Honerkamp, J. (1998a). Estimating parameters in nonlinear differential equations with application to Q-switched lasers. In *Technische Anwendungen von Erkenntnissen der Nichtlinearen Dynamik*, pages 311–314. VDI, Frankfurt.
- Horbelt, W., Timmer, J., Bünner, M., Meucci, R., and Ciofini, M. (2001a). Dynamical modelling of measured time series from a Q-switched CO₂ laser. *Chaos, Solitons & Fractals*, submitted.
- Horbelt, W., Timmer, J., Bünner, M., Meucci, R., and Ciofini, M. (2001b). Identifying physical properties of CO₂ laser by dynamical modeling of measured time series. *To be published in Phys. Rev. E*, 64.

- Horbelt, W., Timmer, J., and Melzer, W. (1998b). Estimating parameters in differential equations with application to physiological data. In Osipenko, G., editor, *Differential equations and applications*, pages 23–33. State Technical University, Saint Petersburg.
- Jaeger, L. and Kantz, H. (1996). Unbiased reconstruction of the dynamics underlying a noisy chaotic time series. *Chaos*, 6:440–450.
- Kaneko, K. (1993). *Theory and applications of coupled map lattices*. Wiley, Chichester.
- Kantz, H. and Jaeger, L. (1997). Improved cost functions for modelling noisy chaotic time series. *Phys. D*, 109:59–69.
- Kelpin, F. D. L., Kirkilionis, M. A., and Kooi, B. W. (2000). Numerical methods and parameter estimation of a structured population model with discrete events in the life history. *J. Theor. Biol.*, 207:217–230.
- Kember, G. and Fowler, A. C. (1993). A correlation-function for choosing time delays in phase portrait reconstructions. *Phys. Lett. A*, 179:72–80.
- Kittel, A., Parisi, A., and Pyragas, K. (1995). Delayed feedback-control of chaos by self-adapted delay-time. *Phys. Lett. A*, 198:433–436.
- Kittel, A., Parisi, J., and Pyragas, K. (1998). Generalized synchronization of chaos in electronic circuit experiments. *Phys. D*, 112:459.
- Lorenz, E. N. (1963). Deterministic aperiodic flow. *J. Atmos. Sci.*, 20:130.
- Mackey, M. C. and Glass, L. (1977). Oscillation and chaos in physiological control systems. *Science*, 197:287–289.
- McSharry, P. E. and Smith, L. A. (1999). Better nonlinear models from noisy data: Attractors with maximum likelihood. *Phys. Rev. Lett.*, 83:4285–4288.
- Meucci, R., Ciofini, M., and Abbate, R. (1996). Suppressing chaos in lasers by negative feedback. *Phys. Rev. E*, 53:R5537–R5540.
- Meucci, R., Ciofini, M., and Wang, P. (1992). Analysis of the dynamical behavior of a Q-switched CO₂ laser: the linear and the nonlinear regime. *Opt. Commun.*, 91:444–452.
- Meucci, R., Wang, P., and Lapucci, A. (1991). Dynamic behavior of the optogalvanic effect in a CO₂-laser. *Opt. Lett.*, 16:1040–1042.
- Meyer-Bourbonneux, F., Dupré, J., and Meyer, C. (1976). Influence des la relaxation rotationnelle dans le déclenchement passif d'un laser CO₂. application aux molécules C₃H₄ et CH₃Cl. *Can. J. Phys.*, 54:205–210.

- NAG (1997). *NAG Fortran Library Mark 18*. The Numerical Algorithms Group Ltd, Oxford, UK.
- Press, W. H., Teukolsky, S. A., Vetter, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C*. Cambridge University Press, Cambridge, 2nd edition.
- Rust, H., Timmer, J., Horbelt, W., and Voss, H. U. (2000). Parametric, non-parametric and parametric modelling of a chaotic circuit time series. In Oishi, S. and Schwarz, W., editors, *Proceedings of the 2000 International Symposium of Nonlinear Theory and its Applications, NOLTA 2000*, volume 2, pages 713–716, Dresden. w.e.b.-Univ.-Verl.
- Schittkowski, K. (1999). PDEFIT: A fortran code for data fitting in partial differential equations. *Optim. Method. Softw.*, 10:539–582.
- Stoer, J. and Bulirsch, R. (1993). *Introduction to Numerical Analysis*. Springer, Heidelberg and New York, 2nd edition.
- Swinton, J., Harwood, J., Grenfell, B. T., and Gilligan, C. A. (1998). Persistence thresholds for phocine distemper virus infection in harbour seal *phoca vitulina* metapopulations. *J. Anim. Ecol.*, 67:54–68.
- Theiler, J. (1990). Statistical precision of dimension estimators. *Phys. Rev. A*, 41:3038–3051.
- Theiler, J. and Smith, L. A. (1995). Anomalous convergence of lyapunov exponent estimates. *Phys. Rev. E*, 51:3738–3741.
- Timmer, J., Horbelt, W., Bünner, M., Ciofini, M., and Meucci, R. (2000a). Estimating parameters in differential equations with application to laser data. In D.S. Broomhead, E.A. Luchinskaya, P. M. and Mullin, T., editors, *Stochaos: Stochastic and Chaotic Dynamics in the Lakes*, pages 617–623, Melville, NY, USA. American Institute of Physics.
- Timmer, J., Müller, T., and Melzer, W. (1998). Numerical methods to determine calcium flux from calcium transients in muscle cells. *Biophys. J.*, 74:1694–1707.
- Timmer, J., Rust, H., Horbelt, W., and Voss, H. (2000b). Parametric, nonparametric and parametric modelling of a chaotic circuit time series. *Phys. Lett. A*, 274:123–134.
- Ulam, S. and von Neumann, J. (1947). *AMS Bull.*, 53:1120.
- van Domselaar, B. and Hemker, P. W. (1975). Nonlinear parameter estimation in initial value problems. Technical Report NW 18/75, Math. Centre Amsterdam.

- Varone, A., Politi, A., and Ciofini, M. (1995). CO₂ laser dynamics with feedback. *Phys. Rev. A*, 52:3176–3182.
- Voss, H. and Kurths, J. (1997). Reconstruction of nonlinear time delay models from data by the use of optimal transformations. *Phys. Lett. A*, 234:336–344.
- Voss, H. and Kurths, J. (1999). Reconstruction of nonlinear time delay models from optical data. *Chaos, Solitons & Fractals*, 10:805–809.
- Voss, H. U. (2001). Real-time anticipation of chaotic states of an electronic circuit. *preprint*.
- Weigend, A. S., Huberman, B. ., and Rumelhart, D. E. (1990). Predicting the future: A connectionist approach. *Int. J. Neur. Comp.*, 1:193–209.
- Weiss, C. O., Vilaseca, R., Abraham, N. B., Corbalán, R., Roldán, E., de Valcárcel, G., Pujol, J., Hübner, U., and Tang, D. Y. (1995). Models, predictions, and experimental measurements of far-infrared NH₃-laser dynamics and comparisons with the Lorenz-Haken model. *Appl. Physics B*, 61:223–242.
- Zehnlé, V., Dangoisse, D., and Glorieux, P. (1992). Behavior of a CO₂ laser under loss modulation: critical analysis of different theoretical models. *Opt. Commun.*, 90:99–105.