A microscopic view of numerous phytoplankton cells, likely diatoms, showing their characteristic circular or oval shapes and intricate internal structures. The cells are densely packed and appear to be in motion, creating a complex, textured pattern. The lighting is bright, highlighting the details of the cell walls and internal organelles.

A system Identification View
on two aquatic topics

Phytoplankton dynamics & water mass mixing

Anouk Debrauwere

120004

**A system identification view
on two aquatic topics:
*phytoplankton dynamics and
water mass mixing.***

Anouk Debrauwere

Promotor: **Prof. Willy Baeyens**

Co-promotor: **Prof. Johan Schoukens**

Advisors: **Dr. Fjo De Ridder**

Dr. Marc Elskens

Proefschrift ingediend tot het behalen van de academische graad van
doctor in de Wetenschappen

Vrije Universiteit Brussel, Faculty of Science
Analytical and Environmental Chemistry

April 2007



Cover & content © Anouk Debrauwere

Members of the jury

- Prof. Robert Finsy (president)
– Vrije Universiteit Brussel
- Prof. Willy Baeyens (promotor)
– Vrije Universiteit Brussel
- Prof. Johan Schoukens (co-promotor)
– Vrije Universiteit Brussel
- Prof. Frank Dehairs
– Vrije Universiteit Brussel
- Prof. Rik Pintelon
– Vrije Universiteit Brussel
- Prof. An Smeyers-Verbeke
– Vrije Universiteit Brussel
- Prof. Joos Vandewalle
– Katholieke Universiteit Brussel
- Prof. Eric Deleersnijder
– Université Catholique de Louvain
- Prof. Karline Soetaert
– NIOO – KNAW
- Prof. Johannes Karstensen
– University of Kiel

Thanks to all of you.

Contents

Contents.....	vii
Notation	xiii
List of symbols.....	xiii
List of abbreviations.....	xv
Samenvatting.....	xvii
Summary	xix
1. General introduction	1
Structure	5
List of own publications.....	7
2. Basic Concepts	9
2.1. Introduction.....	10
2.2. The system	11
2.3. The measurements	11
2.4. The model.....	14
2.5. The parameter estimator	16
2.5.1. Least squares estimator.....	17
2.5.2. Weighted Least Squares estimator	18
2.5.3. Output Weighted Least Squares	20
2.5.4. Input Output Weighted Least Squares estimator	20
2.6. A note on the residual cost function value	23
2.7. The parameter uncertainty	24
2.7.1. Linearised Covariance Matrix (LCM)	24
2.7.2. Monte-Carlo and other resampling methods	26
 <i>Part A: Compartmental models</i>	
3. Tracer experiments and models.....	31
3.1. Introduction.....	32
3.2. Tracer experiments	32

3.2.1. Necessity	32
3.2.2. Experimental design	33
3.2.3. Elements	34
3.3. Models for tracer experiments.....	37
3.3.1. The elemental process.....	38
3.3.2. Assumptions and elemental equations	38
3.3.3. Generalised <i>Elskens et al.</i> [2002] model.....	40
3.4. Conclusions.....	43
4. Estimation of flux rates and their uncertainty.....	45
4.1. Introduction.....	46
4.2. Problem statement.....	46
4.2.1. The system	46
4.2.2. The measurements	46
4.2.3. The model.....	47
4.2.4. Noise assumptions.....	50
4.2.5. Aims	50
4.3. Results.....	51
4.3.1. IOWLS parameter estimator consistency.....	51
4.3.2. Residual cost function distribution	54
4.3.3. Parameter uncertainty estimation using the LCM.....	55
4.4. Summary	60
5. Identifiability of two-compartments models	63
5.1. Introduction.....	64
5.2. Methods.....	66
5.2.1. Testing identifiability	66
5.2.2. Investigated models and experimental setups	68
5.2.3. Procedure.....	71
5.3. Results.....	73
5.4. Conclusions.....	81
6. Model selection.....	83
6.1. Introduction.....	84
6.2. Existing methods	86
6.2.1. AIC and MDL	86
6.2.2. Other model selection criteria.....	88

6.2.3. Cross-validation	90
6.2.4. Note on the choice of the model set	90
6.3. Model selection based on the statistical interpretation of the minimal WLS cost function	91
6.3.1. No model errors	91
6.3.2. Model errors present	94
6.4. Simulation tests	96
6.4.1. FIR model.....	96
6.4.2. Compartmental model.....	99
6.5. Results on field data	105
6.6. Conclusions.....	109
 <i>Part B: Ocean mixing model</i>	
7. Optimum Multiparameter analysis	115
7.1. Introduction.....	116
7.2. Method	116
7.3. Imperfections	118
7.4. Concluding remarks	120
 8. Parametric Optimum Multiparameter analysis	
8.1. Introduction.....	124
8.2. Method	124
8.2.1. Parameterisation to increase precision	124
8.2.2. Model selection to ensure accuracy	130
8.2.3. Weighting scheme	132
8.2.4. Constrained optimisation scheme.....	134
8.2.5. POMP algorithm	135
8.3. Simulation tests	137
8.3.1. Synthetic dataset	137
8.3.2. Error free situation	138
8.3.3. Noiseless situation with wrong splines	138
8.3.4. Noisy data and no systematic errors	139
8.3.5. Conclusions	140
8.4. Results on field data	140
8.4.1. Dataset and hydrological setting.....	141
8.4.2. Input into POMP.....	143
8.4.3. Results from POMP analysis: mixing fractions.....	148

8.4.4. Uncertainties.....	151
8.4.5. Residuals.....	151
8.4.6. Comparison with classical OMP analysis.....	154
8.5. Conclusions.....	158
9. Estimation of heteroscedastic noise variances.....	161
9.1. Introduction.....	162
9.2. Methods.....	162
9.2.1. Literature.....	162
9.2.2. No model errors present.....	164
9.2.3. Model errors present.....	169
9.2.4. Practical comments and algorithm.....	171
9.3. Simulation tests.....	174
9.3.1. Heterogeneous distribution of degrees of freedom.....	174
9.3.2. Performance of model error correction.....	177
9.4. Results on field data.....	179
9.4.1. CLIVAR-SR3 2001 data.....	179
9.4.2. Estimated noise variances.....	179
9.4.3. POMP analysis with new noise variance estimates.....	181
9.5. Conclusions.....	185
10. Identifiability of the source water matrix.....	189
10.1. Introduction.....	190
10.2. Derivation.....	190
10.2.1. Method 1: Jacobian rank.....	190
10.2.2. Method 2.....	192
10.2.3. Example.....	194
10.3. Discussion and conclusions.....	195
11. General conclusions.....	197
12. Perspectives for future research.....	201
12.1. Introduction.....	202
12.2. Future developments in system identification?.....	202
12.3. Future developments in modelling tracer experiments?.....	203
12.4. Future developments of POMP analysis?.....	204

12.5. Future application of POMP: detecting biogeochemical processes.....	207
12.5.1. Introduction	207
12.5.2. Method.....	208
12.5.3. Results and possible future applications.....	209
12.6. Future application of POMP: inclusion of more tracers	212
Appendix.....	213
Reference List	215

Notation

The following general notation is used: scalars are represented by italic symbols, vectors by bold lowercase symbols and matrices by bold uppercase symbols. Only recurring symbols are included in this list.

List of symbols

- \otimes Kronecker product, see Appendix
- $\mathbf{0}(\dots)$ matrix composed by zeros, with dimensions specified between brackets
- $\mathbf{1}(\dots)$ matrix composed by ones, with dimensions specified between brackets
- θ model parameter vector, length n_θ
- $\hat{\theta}$ estimated value of model parameter vector, length n_θ
- Θ potential temperature, i.e. temperature corrected for density effects, °C
- $\hat{\sigma}_h^2$ estimated noise variance of subset h
- $\hat{\sigma}^2$ estimated noise variance vector
- æ atom% excess, expressing the excess of heavy isotope in a sample (Eq. 3.1)
- \mathbf{B} spline matrix in water mixing application, size $N \times n_B$
- \mathbf{C} spline coefficients matrix, size $n_S \times n_B$
- D dissolved phase or concentration of this phase (simplified from $[D]$)
- \mathbf{e} general noise vector, length n_d
- h index representing subset number, $h \in \{1, \dots, n_h\}$ (Chapter 9)
- \mathbf{I}_a $a \times a$ identity matrix
- \mathbf{J} Jacobian matrix (Eq. 2.16), size $n_d \times n_\theta$
- \mathbf{J}_c Jacobian matrix corrected for the presence of equality constraints, as defined in Eq. (9.17)
- \mathbf{J}_h Jacobian matrix of subset h , size $N_h \times n_\theta$ (univariate case) or $N_h n_v \times n_\theta$ (multivariate case)

\mathbf{J}_w	Jacobian matrix weighted with the measurement uncertainties, as defined in Eq. (9.15)
\mathbf{J}_{wc}	weighted Jacobian matrix corrected for the presence of equality constraints
n_θ	number of unknown parameters θ
n_B	number of splines in the water mixing application, $n_B = n_l n_z$
n_d	number of observations, equals N in the univariate case, otherwise equals Nn_v
n_h	number of subsets (Chapter 9)
n_l	number of B-splines in the latitude direction
n_S	number of sources in the water mixing application
n_v	number of variables
n_z	number of B-splines in the depth direction
N	number of samples
N_h	number of samples in subset h
$\mathcal{N}(\mu, \sigma)$	normal distribution with mean μ and standard deviation σ
P	particulate phase or concentration of this phase (short for $[P]$)
q_i	unlabelled species concentration in tracer experiment (natural analogue of the tracer), measured in compartment i
Q_i	labelled species (tracer) concentration in tracer experiment, measured in compartment i
\mathbf{r}	residuals vector, defined as $\mathbf{y} - f_{model}(\mathbf{u}, \theta)$, length n_d
\mathbf{r}_h	residuals of subset h , vector of length N_h (univariate case) or $N_h n_v$ (multivariate case)
R	regeneration rate (to dissolved phase)
\mathbf{S}	sources matrix in water mixing application, size $n_v \times n_S$
\mathbf{u}	input variable vector
U	uptake rate (from dissolved to particulate phase)
\mathbf{v}	vector containing the residual subset variances, length n_h (univariate case) or $n_h n_v$ (multivariate case)
$x_i(k)$	mixing fraction of source i at position k
\mathbf{X}	mixing fractions matrix in the water mixing application, size $n_S \times N$
\mathbf{y}	output data vector, length n_d
\mathbf{y}_h	output data in subset h , vector of length N_h (univariate case) or $N_h n_v$ (multivariate case)

\mathbf{y}_{-h}	all data in \mathbf{y} which are not element of subset h , vector of length $N - N_h$ (univariate case) or $(N - N_h)n_v$ (multivariate case)
\mathbf{Y}	multivariate measurement matrix, size $n_v \times N$

List of abbreviations

AAIW	Antarctic Intermediate Water
AIC	Akaike Information Criterion
CSW	Circumpolar Surface Water
CTD	Conductivity – Temperature – Depth
IOWLS	Input Output Weighted Least Squares
LCM	Linearised Covariance Matrix
LS	Least Squares
MDL	Minimum Description Length
OMP	Optimum MultiParameter
OWLS	Output Weighted Least Squares
POMP	Parametric Optimum MultiParameter
SAMW	Subantarctic Mode Water
STSW	Subtropical Surface Water
SWT	Source Water Type
UCDW	Upper Circumpolar Deep Water
WLS	Weighted Least Squares

Samenvatting

De algemene doelstelling van deze studie is het aanpassen en verbeteren van bestaande modelleringsprocedures, met de bedoeling om meer en meer betrouwbare informatie te extraheren uit de beschikbare waarnemingen. Concreet betekent dit dat veel aandacht besteed is aan het kwantificeren van onzekerheden, gezien deze nodig zijn om een lijn te trekken tussen wat significant is in de metingen en het deel dat niet moet geïnterpreteerd worden. Twee toepassingen werden nader bekeken: tracerexperimenten gemodelleerd door compartimentele modellen (Deel A) en een multivariaat model voor het beschrijven van mengende watermassa's (Deel B). De nadruk van deze studie ligt echter op de gebruikte en ontwikkelde methoden. Zijn bijdrage is niet zo zeer om uitwisselingsnelheden of mengingsfracties te hebben geschat, maar eerder om hun *accurate* schatting te hebben verwezenlijkt, tegelijk met een schatting van de gekoppelde onzekerheid. De voornaamste verwezenlijkingen van deze doctoraatsstudie kunnen in het heel kort samengevat worden als volgt:

- (i) In rekening brengen van inputonzekerheden bij het schatten van modelparameters en hun onzekerheden.
- (ii) Modelselectie gebaseerd op de statistische interpretatie van de residuele Gewogen Kleinste Kwadraten kostfunctie.
- (iii) Verfijning van de Optimum Multiparameter analysis voor de grootschalige reconstructie van spatiale verdeling van mengende watermassa's.
- (iv) Opstellen van algoritme om heteroscedastische ruisvarianties te schatten uit de residuen, met een correctie voor eventuele modelfouten.
- (v) Nagaan van identificeerbaarheid van gegeven experiment-model combinaties.

Summary

The general aim of this study is to modify or improve existing modelling procedures in order to extract more or more reliable information from observations at hand. In practice, this means that much attention is directed to quantifying uncertainties, since these enable to draw a line between what is significant and what should not be interpreted. Two applications were considered: tracer experiments modelled by compartmental models (Part A) and a multivariate water mass mixing model (Part B). However, the focus lies on the methods used and developed to improve the models and their inferences. Indeed, the merit of this work is not to have enabled the estimation of flux rates and mixing fractions, but to have enabled their accurate estimation, together with an estimate of the associated uncertainty. Very briefly, these are the main achievements:

- (i) Inclusion of the input uncertainties in the estimation of model parameters and their uncertainties.
- (ii) Model selection method based on the statistical interpretation of the residual Weighted Least Squares cost function.
- (iii) Improvement of Optimum Multiparameter analysis for large-scale reconstruction of mixing water mass distributions.
- (iv) Construction of an algorithm to estimate heteroscedastic noise variances, from residuals but corrected for model errors.
- (v) Verification of identifiability of given experiment-model combinations.

Chapter 1:

General introduction

“It is probably true quite generally that in the history of human thinking the most fruitful developments take place at those points where two different lines of thought meet. These lines may have roots in quite different parts of human culture, in different times or different cultural environments or different religious traditions: hence if they actually meet, that is, if they are at least so much related to each other that a real interaction can take place, then one may hope that new and interesting developments may follow.”

Werner Heisenberg, in: "Physics and Philosophy", Allen & Unwin, 1963.

Studying natural systems often involves the determination of physical parameters which are not directly experimentally accessible, e.g. rate constants, fractionation factors, characteristic scales of certain processes, etc. In order to quantify these parameters, generally some other, more easily reachable system variables are measured instead. Then, using a chosen formal relation between these measurable quantities and the unknown parameters (a so-called model), the latter can finally be extracted from the measurements. System identification offers a systematic and formal framework for these steps. The basic steps of this system identification ladder are introduced in more detail in Chapter 2, and will return throughout this thesis. In brief they consist of

- (1) collecting data,
- (2) selecting a model or model class to describe these data,
- (3) match the model to the data,
- (4) evaluate the quality of the results.

However, these well-established steps may not always be sufficient. Depending on the specific question, refinements or additional steps may be necessary. Most chapters of this thesis consist of a methodological development that can be seen as one of these additional steps.

The study of natural systems is impeded by a number of aspects. Firstly, these systems are in general vastly more complex than the models that are built to describe them. In other words, a perfect model is an illusion and therefore it is necessary to assess the performance of each model in describing those processes which are of interest. The only information possibly useful for this purpose are the data at hand. But these data are not perfect either: they are associated with an uncertainty which is more complicated than the mere analytical precision. Field data are not only subject to variations due to sampling, handling and analysis, but also to all variations acting on the system, i.e. environmental variability. These variations are not straightforward to quantify, especially because field measurements are often too difficult or expensive to repeat. Each environmental modelling exercise has to take into account these specific difficulties. This often comes down to

making a compromise between a number of aspects like optimal accuracy, optimal precision and lowest cost. With the help of techniques from the system identification theory these criteria can be quantified and a trade-off can be made in an objective way.

This brings us to a formulation of the objective of this study. The general aim is to modify, improve or extend existing modelling procedures in order to extract more or more reliable information from the observations at hand. In practice, this means that much attention is directed to quantifying uncertainties and model errors, since these enable to draw a line between what is significant and what should not be interpreted. No new modelling work is done, at least not in the sense of building a model by translating physical processes to equations. The starting point is always an existing model, which is critically evaluated against the available data, keeping in mind their reliability and uncertainty. Only if it turns out that a modification significantly improves the description of reality or reduces the uncertainty associated with the inferences, it is retained.

In this way, this thesis is primarily theoretically-methodologically directed (indeed no measurements were performed), but every step is tightly linked to a real problem encountered “in the field”. Someone may argue that this work is not fully independent since it needed data from other investigators. But this remark ignores the attempt of this study to be generic; including an experimental part would narrow the focus to one particular system’s problem. Furthermore, these are times of data wealth; it is difficult to hold pace with the speed at which new observations fill databases. Observations from expensive and time-consuming measurements are plenty. Nevertheless, in many cases not all information hidden in these data is extracted, which is a waste of money and time. For that reason, scientists should be encouraged to take some time to thoroughly analyse these available data. These analyses could indeed provide useful insights, e.g. helping to improve the design of future expensive and time-consuming experiments.

Structure

This manuscript is an approximately chronological report of the research conducted during my PhD. This structure has the advantage to naturally bisect the whole in two parts based on the application problem considered. In Part A the application of interest is tracer experiments modelled by a particular class of compartmental models. Chapter 3 introduces this topic, while the subsequent chapters of this first part are primarily concerned with the development or application of system identification tools. Similarly, Part B also contains a number of methodological developments but this time in the context of modelling the mixing of water masses in the ocean. The methods presented may sometimes seem completely dedicated to the system under study, but we would like to emphasize that they are more generic and thus are a contribution to the system identification ladder. Hence, all chapters contribute to the same final goal: obtaining reliable, interpretable, quantitative answers to the posed question.

The manuscript is structured as follows:

Chapter 2 gives a general introduction of the basic concepts of the system identification theory. No new information is given here. Instead, the purpose of this chapter is to familiarise the non-expert reader with the concepts that will be used throughout this manuscript.

Part A: Compartmental models for tracer experiments.

It seems appropriate to first reserve a special chapter to sketch the fundamental ideas of tracer studies and introduce the models used (*Chapter 3*). Again no new research is presented here. Alternatively, this chapter is meant to initiate the amateur reader in this special topic that will be the recurring application in this first thesis part. After this introduction, three subsequent chapters discuss three different steps on the system identification ladder: estimation of the model parameters and their uncertainty (*Chapter 4*), identifiability (*Chapter 5*) and model

selection (*Chapter 6*). In every chapter the tracer experiment is the application of interest, but the issues raised are general.

Part B: Ocean mixing model.

The topic of interest in this second part is the empirical modelling of mixing water masses in the ocean. The reference method to do this is called Optimum Multiparameter (OMP) analysis, which is introduced in *Chapter 7*. This chapter ends with a summary of imperfections associated with this approach. The following chapters are an attempt to address these imperfections. In *Chapter 8* a refinement of the OMP analysis is presented, in which the original unknown parameters are parameterised in order to obtain more robust estimates. It appears that the available uncertainty information is insufficient to perform a thorough quality assessment of the modelling results. Therefore, a new method to estimate the noise variances is developed in *Chapter 9*. Finally, in *Chapter 10*, the possibility to independently estimate the source characteristics is investigated, since they are difficult to objectively define a priori. Chapters 8 – 10 present methodological developments in the framework of OMP analysis. To illustrate the consequences on the analysis results one dataset is consistently used as an example. Oceanographical interpretations fall beyond the scope of this study and hence are reduced to a minimum.

The major new contributions developed in this study are summarised in the *General Conclusions (Chapter 11)*.

Finally, as research never ends, a special chapter is dedicated to a discussion of future research paths flowing out of this study (*Chapter 12*).

List of own publications

Peer-reviewed articles in international journals

1. De Ridder, F., R. Pintelon, J. Schoukens, D. Gillikin, L. Andre, W. Baeyens, **A. de Brauwere** and F. Dehairs (2004), Decoding nonlinear growth rates in biogenic environmental archives, *Geochemistry Geophysics Geosystems.*, 5 AR Q12015.
2. **de Brauwere, A.**, F. De Ridder, M. Elskens, J. Schoukens, R. Pintelon and W. Baeyens (2005), Refined parameter and uncertainty estimation when both variables are subject to error. Case study: estimation of Si consumption and regeneration rates in a marine environment, *Journal of Marine Systems*, 55(3-4), 205-221.
3. **de Brauwere, A.**, F. De Ridder, R. Pintelon, M. Elskens, J. Schoukens and W. Baeyens (2005), Model selection through a statistical analysis of the minimum of a weighted least squares cost function, *Chemometrics and Intelligent Laboratory Systems*, 76(2), 163-173.
4. Elskens, M., W. Baeyens, N. Brion, S. De Galan, L. Goeyens and **A. de Brauwere** (2005), Reliability of N flux rates estimated from 15N enrichment and dilution experiments in aquatic systems, *Global Biogeochemical Cycles* 19(GB4028), doi: 10.1029/2004GB0.
5. **de Brauwere A.**, S. H. M. Jacquet, F. De Ridder, F. Dehairs, R. Pintelon, J. Schoukens and W. Baeyens (2007), Water mass distributions in the Southern Ocean derived from a parametric analysis of mixing water masses, *J. Geophys. Res.* 112, C02021, doi:10.1029/2006JC003742.
6. **de Brauwere A.**, R. Pintelon, F. De Ridder, J. Schoukens and W. Baeyens (2007), Estimation of heteroscedastic measurement noise variances, *Chemometrics and Intelligent Laboratory Systems*, 86, 130-138.

7. De Ridder F., **A. de Brauwere**, R. Pintelon, J. Schoukens, F. Dehairs, W. Baeyens and B. H. Wilkinson (in press), Comment on: Paleoclimatic inference from stable isotope profiles of accretionary biogenic hardparts—a quantitative approach to the evaluation of incomplete data, by Wilkinson, B.H., Ivany, L.C., 2002. *Palaeogeogr. Palaeocl. Palaeoecol.* 185, 95–114, *Palaeogeography Palaeoclimatology Palaeoecology*.
8. Elskens M., R. Corvaisier, C. Beucher, N. Savoye, P. Treguer, **A. de Brauwere** and W. Baeyens (in press), Statistical process control in assessing production and dissolution rates of biogenic silica in marine environments, *Marine Chemistry*.
9. Jacquet S. H. M., **A. de Brauwere**, F. De Ridder and F. Dehairs, Investigating dissolved barium behavior in the Southern Ocean using a parametric mixing model, submitted to *Geophysical Research Letters*.
10. De Ridder F., **A. de Brauwere**, R. Pintelon, J. Schoukens and F. Dehairs, Identification of the accretion rate for annually resolved archives, submitted to *Biogeosciences*.
11. Baeyens W., **A. de Brauwere**, N. Brion, M. De Gieter and M. Leermakers, Arsenic speciation in the River Zenne, Belgium, submitted to *Science of the Total Environment*.

Chapter 2:

Basic concepts

“A model is a more or less complex assembly of constituent hypotheses. The purpose of system identification is therefore that of reconciling this set of concepts with the given field data.”

Stigter and Beck [1994].

“System identification is not ‘merely curve fitting’ if that is the end in itself; it is, if anything, curve fitting as a means to an end, where the end is the rigorous, scientific interpretation of field data. (...) without system identification, in its broadest sense, the process of model development and evaluation should not be accorded the label of ‘scientific’.”

Beck [1987].

2.1. Introduction

This chapter introduces the basic concepts of system identification, or statistical modelling, which are used throughout this thesis. The aim is to make the reader familiar with the logical sequence of steps to be taken leading to a quantitative answer to a scientific question. In short, the basic scheme is based on four steps (Figure 2.1):

- (1) collecting data,
- (2) selecting a model or model class to describe these data,
- (3) match the model to the data,
- (4) evaluate the quality of the results.

Note that the sequence given in this chapter may not be complete. Indeed, depending on the specific question, additional or modified steps may be necessary. Most chapters of this thesis consist of a methodological development that can be seen as one of these additional steps completing the system identification ladder.

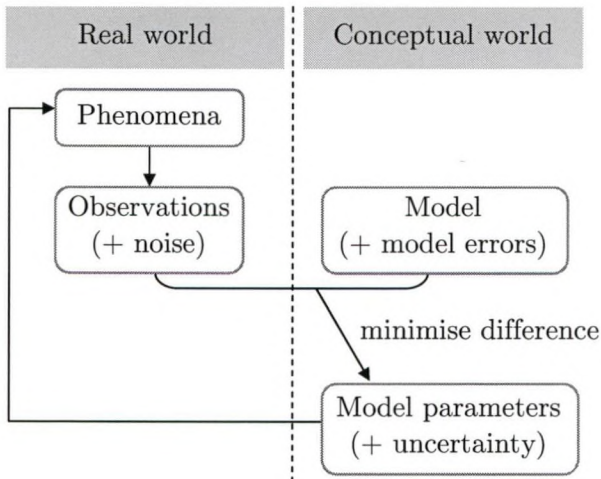


Figure 2.1. Schematic representation of the scientific method or system identification ladder (Adapted from *Soetaert and Herman, 2001*).

2.2. The system

According to *Nihoul* [1975], to model a (marine) system it is first required to define the system without ambiguity, separate it from the outside world and identify the exchanges between the system and the exterior. The system can consist of internal processes as well as interactions with the environment, and as a result some measurable variables fluctuate dependently. This means that if one or several input variables, \mathbf{u}_θ , are varied, one or several output variables, \mathbf{y}_θ , seem to change as a result:

$$\mathbf{y}_\theta = \text{system}(\mathbf{u}_\theta). \quad (2.1)$$

This situation is schematically represented in Figure 2.2a. The question usually is to identify and quantify the processes that cause this relation. Note that it may not always be straightforward to split variables into “input” and “output”. Within the scope of our study, this is not a problem and we will always be concerned with systems in the form of Eq. (2.1).

2.3. The measurements

In order to characterize the system, it is necessary to collect information about it. From the previous paragraph, it follows logically that this is done by measuring the input and output variables. Due to all kinds of fluctuations, the measured values \mathbf{u} and \mathbf{y} are generally not equal to the true values \mathbf{u}_θ and \mathbf{y}_θ (note that a subscript “0” will consistently be used to refer to true values). Consider the following additive noise model,

$$\mathbf{u} = \mathbf{u}_\theta + \mathbf{e}_u \quad \text{and} \quad \mathbf{y} = \mathbf{y}_\theta + \mathbf{e}_y \quad (2.2)$$

where \mathbf{e}_u and \mathbf{e}_y represent the input and output noise vectors.

In many applications, it is explicitly assumed that \mathbf{e}_u is zero or small with respect to \mathbf{e}_y (see *Mandel* [1984] for a meticulous formulation) and as a consequence the input noise contribution can be neglected (Figure

2.2b). In some situations the more general situation is considered: input noise cannot be neglected (Figure 2.2c).

Assumptions 2.1 (noise assumptions)

Throughout this thesis, the noise elements $e_{x,i}$ and $e_{y,i}$ (forming vectors \mathbf{e}_u and \mathbf{e}_y) are assumed to be $\mathcal{N}(0, \sigma_{u,i})$ and $\mathcal{N}(0, \sigma_{y,i})$ distributed random variables ($i = 1, \dots, n_d$), whose values are uncorrelated (pointwise independent).

The variances $\sigma_{u,i}^2$ and $\sigma_{y,i}^2$ are known or priorly estimated, and stored in the vectors σ_u^2 and σ_y^2 .

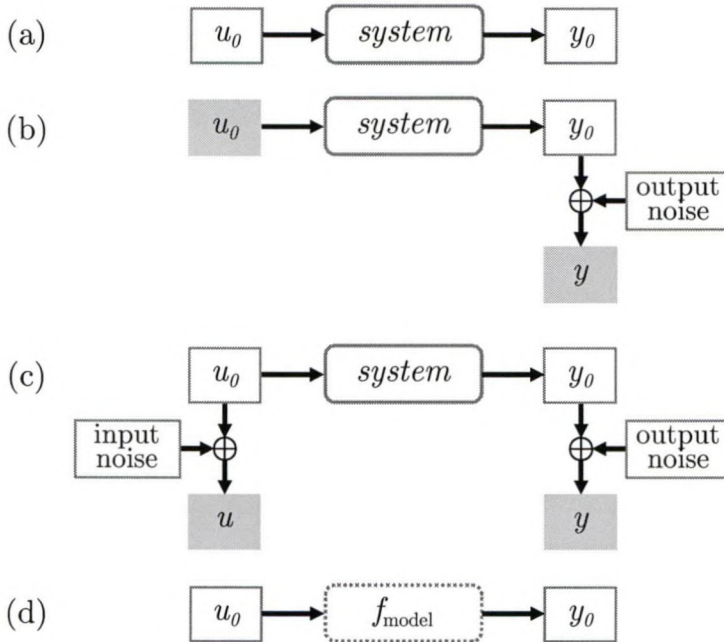


Figure 2.2. Schematic representations of the system, model, measurements and noise concepts. The rounded boxes (\square) stand for transformations, the others are variables. The grey boxes are the measured variables. (a) Illustration of the system concept. (b) The measurements, in the situation of only output noise, or negligible input noise. (c) The measurements, in the situation of non-negligible noise on input and output. (d) Illustration of the “model” concept. The difference with (a) is that the model is a mathematical expression, whereas the system is not (necessarily).

The assumption of normally distributed noise is based on the Central Limit Theorem, which in general terms states that the distribution of an average tends to be Normal as the sample size increases, even when the distribution from which the average is computed is decidedly non-Normal. Therefore, it seems reasonable to assume that most data, which are the result of a sequence of processes, will be normally distributed. However, the normality assumption is in many real life cases only approximate since it would imply that once in while e.g. negative concentrations, population abundances, etc. would be measured, which is obviously not the case. Therefore, some suggest to use the log-normal distribution as a better approximation [e.g. *Limpert et al.*, 2001]. This distribution is skewed and strictly nonnegative, and can be seen as the multiplicative (instead of additive) version of the Central Limit Theorem. Nevertheless, we believe that in most cases the possible error by assuming normally instead of log-normally distributed noise is negligible. Especially as long as close-to-average behaviour is investigated, both distributions are quite similar. Moreover, it must be considered that throughout this thesis secondary data (i.e. not collected and thus controlled by myself) are used. Therefore, the normality assumption seems the best hypothesis.

The assumption of zero mean is an essential feature of noise, because it implies that when enough measurements are considered, the noise terms in (2.2) cancel out, leaving only the true values. The assumption of pointwise independent noise elements stems from the way the measurements are made. Indeed, in the applications that will be considered, every measurement is the result of a separate experiment or analysis.

Remark that the noise terms do not only originate from instrumental errors, but also from all other sources of variability acting on the measurement (e.g. sampling errors) and influencing the interpretation of the measurement in the framework of the model (e.g. environmental variability). An example of the latter occurs when the measurements are made at different scales than those represented in the model (representation error). This aspect of the noise will be important in Part B and will therefore be further discussed then (cf. section 9.5).

2.4. The model

The next step in the identification procedure is to select a mathematical model for the system. Whereas the system is merely a conceptual, qualitative idea, the model is a formal, mathematical description of the system processes under study (Figure 2.2d). Only parametric models will be considered in this thesis. This kind of model relates the true input and output measurements via the mathematical expression f_{model} and the true model parameters θ_0 :

$$\mathbf{y}_0 = f_{\text{model}}(\mathbf{u}_0, \theta_0). \quad (2.3)$$

From expression (2.3) it is clear that the model enables to predict the output measurements, given the input measurements and the model parameters. Or, alternatively, if input and output measurements are known, the model parameter can be computed. It is exactly the quantification of the model parameters that is supposed to answer the original question about the system. Depending on the actual question the model equations can have a different structure, such that the model parameters provide the needed information. Therefore, the choice of the model is a fundamental step, which deserves some elaboration.

Many types of models exist. A fundamental distinction exists between “models which try to give an efficient description of the input-output behavior of a system without relying on hypotheses about how the system works internally, and models which try to give an internal description of the system. Models of the first class are called black box models. They are usually relatively simple, parsimonious and identifiable from the observations available. A modeling approach using this type of model is generally preferred by statisticians and control engineers. (...) Models of the second class are called mechanistic. They are usually more complex and poorly identifiable. A modeling approach relying on this model class is usually preferred by natural scientists. It is characterized by an extensive use of causal hypotheses based on current understanding on how processes work and it is closely linked to a reductionistic world view.” [Brun et al., 2001].

Which kind of model to choose is usually dictated by the question to be answered. Black box models are the preferred vehicles when it comes down to exhaustively analysing and interpreting data. These models are really meant to extract information from the observations at hand under the form of model parameters. Throughout this thesis, the starting point is always the available data and the general question is to extract as much quantitative information from them as possible. Therefore, all models used in this thesis are essentially black box models.

However, my experience is that nowadays in environmental sciences mechanistic (or white box) models are much more popular than black box models. Both in ecology and oceanography, much effort and attention goes to the development of large, physical models, containing the latest hypotheses regarding possibly relevant processes [e.g. *Lancelot et al.*, 2000; *Maltrud et al.*, 1998]. This can probably be appreciated in the context of ever increasing computing power. Of course, these models present a great step in the understanding of internal processes which are not accessible otherwise. Nevertheless, it is rarely admitted that these models essentially are “*formalized archives of hypotheses*” [Beck, 1987] which cannot be validated against observations because they have become so enormously more complex than the available data. As in all models these models contain parameters which must be adjusted to reasonably describe the observations. Contrary to black box models, it is usually impossible to quantify the large number of model parameters from the data at hand. Instead, the model parameters are usually set a priori (e.g. based on literature values) and the model serves to *simulate* reality.

A reconciling view on the black box versus white box discussion is that both can also be combined. For instance, in engineering applications, black box models are often used in a preprocessing step to “clean up” the data by removing most of the non-significant noise. In a subsequent step, physical models may be used to further model the preprocessed data, without any loss of significant information. Alternatively, simple black box models can be used to “postprocess” the the output of large, sophisticated models, enabling a more straightforward understanding and interpretation of this output. The output files produced nowadays

by complex dynamical models are so huge that interpretation of these results becomes increasingly difficult. In this context, simple black box models can be useful to extract the wanted information. Examples of the usefulness of simple box models for interpreting the results of sophisticated three-dimensional hydrodynamical models are given in *Deleersnijder et al.* [1997, 1998].

Once this first choice concerning the purpose of the model is made, still many model structures are possible. A model is said to be *linear-in-the-parameters* (or simply *linear*), if it is a linear function of the model parameters

$$f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta}) = \mathbf{K}(\mathbf{u})\boldsymbol{\theta}. \quad (2.4)$$

For instance a model $f_{\text{model}}(u, \boldsymbol{\theta}) = \theta_1 u + \theta_2 u^2$ is linear (in the parameters), whereas a model $f_{\text{model}}(u, \boldsymbol{\theta}) = (\theta_1 + \theta_2 u)/(\theta_3 + \theta_4 u^2)$ is not.

Even when the model structure is fixed, often some freedom of choice remains with respect to the model complexity or order. For instance, which order of a polynomial model should be used? Once all these choices about the model are made, i.e. f_{model} is fixed, the problem of determining the system processes can be reformulated as quantifying the set of unknown model parameters $\boldsymbol{\theta}$.

2.5. The parameter estimator

After defining the system, carrying out the measurements, and choosing the model, the remaining step is to combine the model and the measurements in order to estimate the value of the model parameters. This is usually done by varying the values of the model parameters until some measure of the distance between the model and the data is minimised. Many “measures of distance” exist and we will call them a *cost function*. The choice of the cost function is important because it will fix all statistical properties of the estimated parameter values. In this section, only a few cost functions will be discussed, all based on the idea of “least squares”, i.e. on minimising the squared difference between

model and measurements. This class of cost functions has the advantage of having well known properties, especially under the noise assumptions made above (cf. Assumptions 2.1).

2.5.1. Least squares estimator

If no noise information is available, often the Least Squares (LS) estimator is used. The LS estimator of $\boldsymbol{\theta}$ is the parameter value, denoted by $\hat{\boldsymbol{\theta}}_{LS}$, that minimizes the LS cost function V_{LS} :

$$\hat{\boldsymbol{\theta}}_{LS} = \arg \min_{\boldsymbol{\theta}} V_{LS} \quad (2.5)$$

with

$$V_{LS} = \sum_{i=1}^{n_d} (y_i - f_{model,i}(\mathbf{u}, \boldsymbol{\theta}))^2 = \sum_{i=1}^{n_d} r_i^2 = \mathbf{r}^T \mathbf{r}. \quad (2.6)$$

n_d is the number of (output) observations, \mathbf{r} stands for the difference, or residual, between the output measurements \mathbf{y} and the model prediction $f_{model}(\mathbf{u}, \boldsymbol{\theta})$. Recalling Eqs. (2.2) and (2.3), it is clear that the residual \mathbf{r} is nothing but noise (\mathbf{e}) if the model prediction is perfect. It is a subtle distinction, but the symbol \mathbf{r} will always be used to denote the difference between output observations and model, whereas \mathbf{e} is only used to explicitly indicate that it concerns only stochastic, random noise. In words, equations (2.5) - (2.6) express that the optimal parameter values correspond to a minimal squared difference between the observations and the model. Obviously, this estimator is only accurate if the model is a close approximation of the observed processes, i.e. if the residuals can be equated with noise.

If the model is linear-in-the-parameters (see Eq. (2.4)), the minimum of V_{LS} can exactly be found by derivation of Eq. (2.6), giving an explicit expression for the associated optimal parameter values:

$$\hat{\boldsymbol{\theta}}_{LS} = (\mathbf{K}^T \mathbf{K})^{-1} \mathbf{K}^T \mathbf{y}. \quad (2.7)^\dagger$$

If the model is nonlinear-in-the-parameters, typically iterative optimisation procedures are necessary to find the parameter values corresponding to a minimal cost function.

A number of interesting properties have been proven for this LS estimator, mostly valid only under the assumption that the model is correct, that input noise is negligible and that the output noise variance is identical for all data [*Pintelon and Schoukens, 2001*].

2.5.2. Weighted Least Squares estimator

If some observations are measured with higher precision than others, it may be useful to emphasize their importance during the estimation of the unknown parameters. This is generally done by attributing to every variable a different weight, inversely proportional with the uncertainty associated with it. In the resulting Weighted Least Squares cost function each term is thus scaled by its associated variance, which should obviously be known in advance

$$V_{WLS} = \sum_{i=1}^{n_d} \frac{(y_i - f_{model,i}(\mathbf{u}, \boldsymbol{\theta}))^2}{\hat{\sigma}_i^2} = \sum_{i=1}^{n_d} \frac{r_i^2}{\hat{\sigma}_i^2} = \mathbf{r}^T \mathbf{W} \mathbf{r}. \quad (2.8)$$

\mathbf{W} represents a diagonal weighting matrix containing the inverses of the variances $\hat{\sigma}_i^2$,

[†] Note that this expression is only valid if the columns of \mathbf{K} are independent because otherwise $\mathbf{K}^T \mathbf{K}$ is not invertible.

$$\mathbf{W} = \begin{pmatrix} \frac{1}{\hat{\sigma}_1^2} & 0 & \cdots & 0 \\ 0 & \frac{1}{\hat{\sigma}_2^2} & \cdots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \frac{1}{\hat{\sigma}_N^2} \end{pmatrix}. \quad (2.9)$$

The fact that \mathbf{W} is diagonal follows from the assumption that the noise elements are uncorrelated (cf. Assumptions 2.1). The exact meaning of the variances $\hat{\sigma}_i^2$ depends on the assumptions made concerning the relative uncertainties of the different observations. Two special cases will be discussed in the next sections (2.5.3 – 2.5.4). Note that if all variances are one, the WLS cost function simplifies to the LS expression. Also if the variances are not one but all equal, the WLS estimator reduces to the LS case, since the cost function is only rescaled by the constant variance and hence it will still reach its minimum for the same $\hat{\boldsymbol{\theta}}$.

In analogy with the LS case, the optimal parameter values are defined by

$$\hat{\boldsymbol{\theta}}_{WLS} = \arg \min_{\boldsymbol{\theta}} V_{WLS}, \quad (2.10)$$

and an explicit expression for the WLS estimator can be derived if the model is linear (and if $\mathbf{K}^T \mathbf{W} \mathbf{K}$ is invertible)

$$\hat{\boldsymbol{\theta}}_{WLS} = (\mathbf{K}^T \mathbf{W} \mathbf{K})^{-1} \mathbf{K}^T \mathbf{W} \mathbf{y}. \quad (2.11)$$

2.5.3. Output Weighted Least Squares

If input noise is negligible and the magnitude of the output noise variances is known, the Output Weighted Least Squares (OWLS) estimator is commonly used [Bard, 1974; Valsami et al., 2000]. This is the WLS estimator with $\hat{\sigma}_i^2 = \hat{\sigma}_{y,i}^2$. So, in the WLS cost function (2.8), the squared error is “weighted” by the corresponding *output* variance. However, by only considering the output variances it is presumed that $\hat{\sigma}_{y,i}^2$ is a good estimate of the variance of the error, i.e. $y_i - f_{\text{model},i}(\mathbf{u}, \boldsymbol{\theta})$, in (2.8) and hence that the input variations are negligible compared with the noise on the output. If this assumption is violated, this can lead to a bias in $\hat{\boldsymbol{\theta}}_{WLS}$ [Powell and Macdonald, 1972]. In the next section, a refinement is suggested that better adapts the WLS estimator for those cases.

2.5.4. Input Output Weighted Least Squares estimator

This section presents a refinement of the OWLS estimator in order to include the input noise effect as well. \mathbf{W} should express the total variance of the error ($\mathbf{y} - f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})$). In the OWLS approach, as only the output is supposed to be noisy, $f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})$ is noise-free and the total variance equals the variance of \mathbf{y} . In that case, \mathbf{W} is indeed composed by the inverse variances of the individual output measurements $\hat{\sigma}_{y,i}^2$.

Now, the aim is to take into account the input variances as well. For that, their effect on the $f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})$ term should be known or estimated. This is achieved by linearisation of the input noise contribution, as was already described by Clutton-Brock [1967] in the framework of likelihood functions. Chandler [1972] has reformulated the problem in terms of Weighted Least Squares. Figure 2.3 schematically represents this adjusted noise conceptualization. In fact, the input noise is transformed into output noise via a linearised approximation, and added to the original output noise, by that way delivering an improved estimation of the total variation. Now the refined weighting matrix is still a diagonal matrix, but composed by inverse variances $\hat{\sigma}_i^2$, which are first order corrections of $\hat{\sigma}_{y,i}^2$:

$$\hat{\sigma}_i^2 = \hat{\sigma}_{y,i}^2 + \sum_{j=1}^{n_d} \left| \frac{\partial y_i}{\partial u_j} \right|^2 \cdot \hat{\sigma}_{u,j}^2 = (\mathbf{W}[i])^{-1}. \quad (2.12)$$

Note that in the special case where input variations appear to be negligible compared to output variations, the refined weighting matrix simplifies to the WLS \mathbf{W} , as the correction term in (2.12) tends to zero. Therefore, (2.12) is generally applicable, and it is not necessary to make a priori checks about whether or not the input noise is negligible. Substituting this expression for \mathbf{W} in the equation of the WLS cost function (2.8), creates a refined cost function which finally takes the input variations into account. The resulting Input Output Weighted Least Squares (IOWLS) cost function and estimator, are easily derived in analogy with the OWLS situation, only adapting the content of the weighting matrix \mathbf{W} .

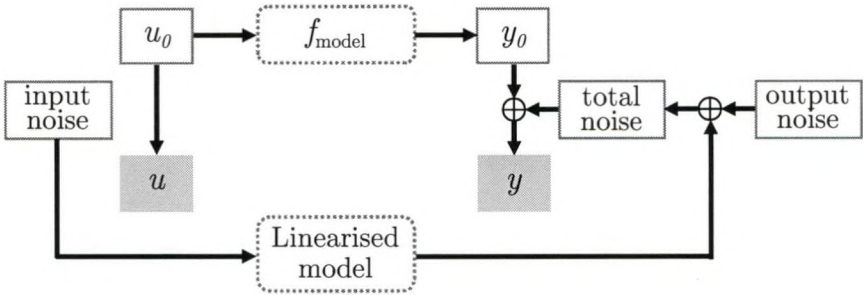


Figure 2.3. Scheme of the noise transformation underlying the refined total noise estimation. It is as if there were no input noise added to the input variables, instead the input noise is transformed into output noise and added to the original output noise. The resulting total noise is treated as if it were merely output noise, but actually it represents an estimation of the total variation. Meaning of boxes is as in Figure 2.2.

Summary

*Assumptions***LS:**

$$\sigma_{u,i}^2 = 0 \text{ and } \sigma_{y,i}^2 = \sigma_y^2 (\forall i).$$

WLS:

$\sigma_{y,i}^2$ are known or priorly estimated ($\forall i$),

and $\sigma_{u,i}^2 = 0 (\forall i)$ **(OWLS)**

or $\sigma_{u,i}^2$ are known or priorly estimated ($\forall i$) **(IOWLS)**

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} V \text{ with}$$

LS:

$$V_{LS} = \sum_{i=1}^{n_d} (y_i - f_{model,i}(\mathbf{u}, \boldsymbol{\theta}))^2 = \sum_{i=1}^{n_d} r_i^2 = \mathbf{r}^T \mathbf{r}.$$

WLS:

$$V_{WLS} = \sum_{i=1}^{n_d} \frac{(y_i - f_{model,i}(\mathbf{u}, \boldsymbol{\theta}))^2}{\hat{\sigma}_i^2} = \sum_{i=1}^{n_d} \frac{r_i^2}{\hat{\sigma}_i^2} = \mathbf{r}^T \mathbf{W} \mathbf{r}.$$

with $\hat{\sigma}_i^2 = \hat{\sigma}_{y,i}^2 = (\mathbf{W}[i])^{-1}$ **(OWLS)**

or $\hat{\sigma}_i^2 = \hat{\sigma}_{y,i}^2 + \sum_{j=1}^{n_d} \left| \frac{\partial y_i}{\partial u_j} \right|^2 \cdot \hat{\sigma}_{u,j}^2 = (\mathbf{W}[i])^{-1}$ **(IOWLS)**

2.6. A note on the residual cost function value

At this point a value can be found for the unknown parameter, i.e. the original question can be answered. However, the system identification ladder does not stop here, since nothing is known about the quality of this value. How much can this estimated value be trusted? Indeed, the observations and the model used to come to this value are generally not perfect, and some quantification of this imperfection could help to evaluate the reliability of the whole procedure. An overall figure of merit is the residual cost function value.

We saw that $\hat{\boldsymbol{\theta}}_{WLS}$ corresponds to that parameter value for which the WLS cost function (2.8) is minimal. Due to measurement noise, even for the optimal parameters there will be a residual difference between the measurements \mathbf{y} and the optimal model values $f_{\text{model}}(\mathbf{u}, \hat{\boldsymbol{\theta}}_{WLS})$. Therefore, the minimal value of the WLS cost function $V_{WLS}(\hat{\boldsymbol{\theta}}_{WLS})$ will not be zero, but equal to a residual value. If the $\hat{\sigma}_i^2$ used in the cost function are a good approximation of the true σ_i^2 , and under the previously stated noise assumptions (cf. Assumptions 2.1), this residual value has the interesting feature of being a sample from a χ^2 distribution with $n_d - n_\theta$ degrees of freedom [Box, 1970; Rod and Hancil, 1980], where n_d stands for the number of measurements and n_θ for the number of free parameters in the model. Formally,

$$V_{WLS}(\hat{\boldsymbol{\theta}}_{WLS}) \sim \chi_{n_d - n_\theta}^2. \quad (2.13)$$

Using this information, it is possible to assess the “probability” of a residual cost function value. If the residual value falls beyond a chosen confidence limit, the results should be rejected for being “unlikely”. For instance, if the residual value is significantly higher than expected, the remaining difference between model and measurements is too high to be explained only by stochastic measurement noise. This significant difference between the expected and observed residual cost function is an indicator of systematic errors. These can be due to

- (i) model errors,

- (ii) artefacts in the measurements,
- (iii) ill estimation of the experimental uncertainties, or
- (iv) numerical issues like a non-global (i.e. local) cost function minimum.

In order to determine exactly which one or combination of these four possibilities is the cause of the unexpected residual cost function, it is necessary to analyze replicate datasets or to use external information.

To summarize, using a WLS cost function has the additional benefit of providing an internal quality control of the modelling results, via interpretation of its residual value. Yet, note that this property is only valid as long as the noise assumptions hold. For more detailed information, the individual residual terms can be inspected, possibly giving more insight in the quality of the observation versus that of the model. Another way to quantify the modelling reliability is by estimating the uncertainties associated with the parameter values found. This will be the subject of next section.

2.7. The parameter uncertainty

In order to interpret the estimated parameter values, it is necessary to assess the uncertainty associated with them. Due to the noise on the data used to estimate the parameters, the parameter values will generally not equal the hypothetical true values. This difference or error resulting from stochastic variations is referred to as *precision*. *Accuracy* describes the closeness of a result to its true or accepted value, i.e. whether systematic errors are present or not [e.g. *Skoog et al*, 2000].

2.7.1. Linearised Covariance Matrix (LCM)

For linear models, straightforward “error propagation” calculations are easily performed since an explicit expression for the parameters as a function of the measurements exists (see Eqs. (2.7) and (2.11)). The parameter covariance matrix (Cov) can then be estimated by

$$Cov(\hat{\boldsymbol{\theta}}) = (\mathbf{K}^T \mathbf{W} \mathbf{K})^{-1}, \quad (2.14)$$

with \mathbf{W} the weighting matrix. Obviously, it is important to use *that* \mathbf{W} that takes into account all significant variances, i.e. both input and output variances if necessary. The diagonal elements of $Cov(\hat{\boldsymbol{\theta}})$ are the parameter variances, the off-diagonal elements present the covariances.

For nonlinear models, such an exact expression does not exist. Alternatively, approximations for Cov are often used, usually providing reasonable parameter uncertainty estimates. We consider the following first order approximation of Cov , assuming that the model can be linearised with respect to the parameters near the optimal values:

$$Cov(\hat{\boldsymbol{\theta}}) = (\mathbf{J}(\hat{\boldsymbol{\theta}})^T \mathbf{W} \mathbf{J}(\hat{\boldsymbol{\theta}}))^{-1}, \quad (2.15)$$

where $\mathbf{J}(\hat{\boldsymbol{\theta}})$ stands for the $(n_d \times n_\theta)$ Jacobian matrix of the model function $f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})$ at $\hat{\boldsymbol{\theta}}$,

$$\mathbf{J}(\hat{\boldsymbol{\theta}}) = \left. \frac{\partial f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = \begin{pmatrix} \frac{\partial f_{\text{model},1}(\mathbf{u}, \boldsymbol{\theta})}{\partial \theta_1} & \dots & \frac{\partial f_{\text{model},1}(\mathbf{u}, \boldsymbol{\theta})}{\partial \theta_{n_\theta}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_{\text{model},n_d}(\mathbf{u}, \boldsymbol{\theta})}{\partial \theta_1} & \dots & \frac{\partial f_{\text{model},n_d}(\mathbf{u}, \boldsymbol{\theta})}{\partial \theta_{n_\theta}} \end{pmatrix}_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} \quad (2.16)$$

If the model is linear-in-the-parameters (cf. Eq. (2.4)), the Jacobian matrix (2.16) equals the linear model matrix \mathbf{K} , and Eq. (2.15) will indeed simplify to Eq. (2.14). Using the Linearised Covariance Matrix (LCM) expression (2.15) to estimate the parameter uncertainty has the benefit of being very easy and fast to compute [Donaldson and Schnabel, 1987], e.g. it can usually be integrated in the parameter optimization algorithm as a final one-step calculation. The main disadvantage is that

in case of nonlinear models, the calculated parameter variances are only approximations. However, expression (2.15) has been proven to be asymptotically exact [e.g. *Pintelon and Schoukens*, 2001], i.e. for increasing number of measurements (2.15) approaches the true covariance matrix, even if the model is nonlinear. If the number of data is small, the accuracy of the approximation varies for each dataset and each model, depending on the “degree of nonlinearity” and the signal-to-noise ratio of the model-measurement combination. In some cases it can become very poor.

Finally, other approximations for Cov exist, e.g. using the Hessian matrix (second derivative) instead of the Jacobian. However, it has been shown that these computationally more expensive and numerically less stable alternative procedures do not improve the uncertainty estimation significantly [*Donaldson and Schnabel*, 1987; *Varah*, 1990].

2.7.2. Monte-Carlo and other resampling methods

“The term ‘Monte Carlo’ was coined by Nicholas Metropolis during the Manhattan Project of World War II, because of the similarity of statistical simulation to games of chance, and the allusion to the capital of Monaco.” [*Minasny and McBratney*, 2002].

We already mentioned that for nonlinear models no explicit expression for the optimal parameters exists and therefore no analytical “error propagation” expression can be derived. Instead, Monte-Carlo, bootstrap or jackknife simulations are often used to estimate the parameter uncertainty [*Meinrath et al.*, 2000 and references therein]. These procedures consist of simulating a large number of datasets, based on the true measurements, but either with a different simulated random noise realization (Monte-Carlo) or with a different subsampling from the original dataset (bootstrap and jackknife). For each such “synthetic dataset” the optimal parameters are estimated. This results in a distribution of best-fit parameter values from which the statistical properties can be examined.

Usually these methods are said to deliver accurate results as long as enough simulations are performed [e.g. *Massart et al.*, 1997]. However, it is often not mentioned that an additional condition for reliable results is the use of an accurate parameter estimator. Indeed, these simulation methods use perturbations around the optimised parameter values. Therefore, it is implicitly assumed that these parameter values are unbiased. For instance, if input noise is negligible, the OWLS estimator (cf. 2.5.2) is unbiased and can be used. However, if input noise is not negligible, this is not true anymore and another, unbiased estimator should be chosen. In short, these simulation methods are useful tools to assess parameter uncertainties, as long as some conditions are met. Their main disadvantages are that they are (i) time-consuming, due to the large number of simulations to be processed, and (ii) more troublesome to incorporate in an automated routine calculation, because e.g. optimisation problems may occur for particular noise realizations or resamplings.

Part A:

Compartmental models

Although this Part A groups a number of quite diverse modelling issues, all are applied to one common application, namely compartmental models for tracer experiments. Therefore, it seems appropriate to first reserve a special chapter to sketch the fundamental ideas of tracer studies and introduce the models used (Chapter 3). This chapter should familiarise the reader with the kind of data and the model class that will be considered in this first thesis part. After this introduction, three subsequent chapters discuss three different steps on the system identification ladder: estimation of the model parameters and their uncertainty (Chapter 4), identifiability (Chapter 5) and model selection (Chapter 6). In every chapter the tracer experiment is the application of interest, but the issues raised here are general. Note that the order of the chapters does not reflect the logical sequence of steps, but are rather an expression of the chronology of the work.

Chapter 3:

Tracer experiments and models

“An isotopic tracer, (also "isotopic marker" or "isotopic label"), is used in chemistry and biochemistry to help understand chemical reactions and interactions. In this technique, one or more of the atoms of the molecule of interest is substituted for an atom of the same chemical element, but of a different (...) isotope. Because the atom has the same number of protons, it will behave in almost exactly the same way chemically as other atoms in the compound, and with few exceptions will not interfere with the reaction under investigation. The difference in the number of neutrons, however, means that it can be detected separately from the other atoms of the same element.

(...)

Isotopic labeling is a technique for tracking the passage of a sample of substance through a system. The substance is 'labeled' by including unusual isotopes in its chemical composition. If these unusual isotopes are later detected in a certain part of the system, they must have come from the labeled substance.”

From Wikipedia, the free encyclopedia.

3.1. Introduction

The purpose of tracer experiments is to follow the path of a specific compound through a natural system in order to retrieve information about internal processes, and more in particular about the rates of these processes. In this chapter an introduction is given both on the experimental setup and on the model class chosen in this thesis to describe this kind of datasets.

3.2. Tracer experiments

3.2.1. Necessity

Tracer experiments originate from the need (or the will) to quantify the different processes transforming and transporting nutrients between various reservoirs. How fast are nutrients taken up by e.g. phytoplankton? And at what rate are they released again to the aquatic environment? There are two main reasons why changes in *total* nutrient concentration do not deliver satisfactory answers, and thus why tracers are needed.

“Most importantly, although changes in the concentration of a particular nutrient with time can provide information on the net balance of production or use of that nutrient, such changes are often the result of multiple metabolic reactions. Tracer procedures allow us to dissect the individual processes contributing to the net flux.” [Glibert and Capone, 1993].

Besides the more specific information tracers provide, isotopic tracers were also introduced because they can be measured with superior analytical precision [Harrison, 1983], enabling the measurement of fluxes in field conditions where rates of exchange are small or in close balance (small net change).

A tracer should satisfy a number of properties to be useful. The necessary properties for an isotopic tracer are summarised below (Assumptions 3.1).

Assumptions 3.1 (*behaviour of isotopic tracer*)

- (i) The tracer undergoes the same transformations as the unlabelled substrate. This implies e.g. that isotopic fractionation is negligible during the exchange processes.
- (ii) There is no exchange of the isotope between the labelled compound and other compounds in the system.
- (iii) The tracer is initially not in isotopic equilibrium with the system under study, otherwise only net fluxes can be detected.
- (iv) The tracer addition does not perturb the state of the system as a whole, i.e. the compartments nor the transformation processes are significantly perturbed.
- (v) The change of the tracer over time is quantifiable.

The validity of these assumptions in real applications is not always straightforward to check. In some cases they have been verified. For silicon, for instance, assumptions (i) and (ii) are not seriously violated [De La Rocha *et al.*, 1997]. For nitrogen, a numerical test has shown that fractionation effects are negligible compared to the elevated level of isotopic tracer present in a tracer experiment (in natural abundances fractionation may be significant) [Elskens *et al.*, 2005]. If incubation times are too long, assumption (iii) may eventually be violated. To satisfy assumption (v), tracer additions $< 10\%$ of the ambient concentration are used, if possible.

3.2.2. Experimental design

Usually the experimental procedure proceeds as follows:

- (i) sampling of water,
- (ii) determination of initial conditions (concentrations),
- (iii) addition of tracer as an isotopically labelled compound,

- (iv) incubation under appropriate conditions,
- (v) filtration of the water to separate particulate fraction (plankton),
- (vi) analysis of the phases of interest.

Note that the composition of the phases are only determined at the beginning and at the end of the experiment. This kind of *endpoint method* (versus a *kinetic* experiment) is used for reasons of experimental feasibility. Through history, the protocol has been adapted mostly concerning when and in which phases the isotopic composition and/or total nutrient concentration should be measured. These decisions depend on the processes assumed (not) to occur during incubation, and hence are tightly related to the model proposed to represent these transformations.

3.2.3. Elements

Especially in the early days of tracer studies, the focus was mainly on nitrogen compounds. This interest originates from the need to better understand the role of the oceans in the sequestration of atmospheric carbon dioxide.

“In order to quantify the carbon dioxide fluxes exchanged at the air-sea interface, there is nowadays a general consensus to integrate physical processes (hydrodynamics), chemical processes (thermodynamic pump) and biological activities (biological pump). The latter are responsible for the photosynthetic carbon fixation (primary production) and carbon dioxide release through respiration and remineralisation of organic matter. On a global scale, the main part of the particulate organic matter (POM) is recycled within the euphotic zone (...). Carbon dioxide returns to the atmosphere within relatively short times scales ranging from days to months.” [Elskens, 1999].

The remaining carbon can escape the photic zone. Quantification of this biologically mediated export of organic matter is crucial because it represents a potential long term sink for atmospheric carbon dioxide. The biological pump, which is principally driven by photosynthesis by

phytoplankton, is usually not limited by the availability of carbon, but by other nutrients. Biological oceanographers have historically tended to think of nitrogen to limit photosynthesis, which is the first reason why to study the nitrogen cycle. Now it is known that other elements can also be limiting, depending on the system under study, e.g. phosphorus, iron, ...

An additional reason to study the nitrogen flow through the ecosystem is that it seemed to hold the key to assess the exportable proportion of primary production. *Dugdale and Goering* [1967] argued that the rate of export of organic nitrogen from the upper ocean is determined by which nitrogen source the phytoplankton community uses (Figure 3.1).

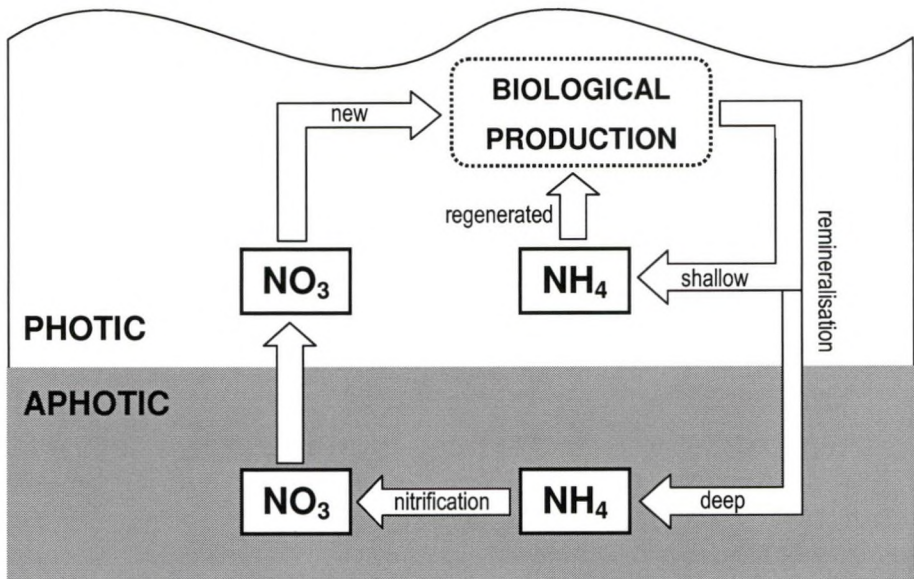


Figure 3.1. This diagram shows the origins of the terms new and regenerated production. To put it simply, new production is that derived from nitrate (NO_3), while regenerated production is derived from ammonium (NH_4). Since nitrate is believed to only be produced at depth (via nitrification), knowing the ratio between new and regenerated production (the so-called f-ratio) allows one to infer the flux of organic nitrogen from the photic zone to the aphotic zone.

Ammonia is the so-called “*regenerated nitrogen*”, because it is formed by recycling within the system. Therefore, if the system is in steady state all primary production associated with ammonia should remain in the system. Only “*new production*” from “*new nitrogen*” like nitrate (and N_2) can be exported to the deep ocean, because the new nitrogen sources are supplied from outside the upper ocean system and can replenish the lost phytoplankton. Although this paradigm is still today the basis of many studies on export production, it should be pointed out that the strict distinction between ammonia uptake representing regenerated production and new production being supplied by nitrate does not always hold.

Finally, nitrogen occurs in a relatively constant ratio to carbon and phosphorus, such that inferences about these latter elements can still be made.

This being said to motivate the study of nitrogen, also silicon tracer experiments will be considered in the chapters of *Part A*. Why study silicon? A special phytoplankton species, called diatoms, need silicic acid for the formation of their frustules (Figure 3.2), so studying silicon dynamics can be used as a clue to better understand the behaviour of diatoms in the oceans. Diatoms are said to make up about a quarter of the world’s plant life’s weight. So it is not surprising that in many regions they are important members of the phytoplankton community and major primary producers [e.g. the Southern Ocean, *Brzezinski et al.*, 2001]. Due to their siliceous skeleton, they are heavier than many other phytoplankton species, and hence are more likely to sink to the deep ocean after dying. This sinking is a key process in the framework of export production or the quantification of the amount of CO_2 that can be sequestered in the deep ocean.



Figure 3.2. Part of a large group of cleaned diatom frustules arranged by the hand of some anonymous Victorian microscopist indulging one of the microscopical pastimes of the day. Most diatoms are much less than half a millimetre in size. This picture illustrates the enormous variety of (beautiful) shapes exhibited by diatoms. Charles Darwin would have said in 1872: “*Few objects are more beautiful than the minute siliceous cases of the diatomaceae*”.

from:

<http://www.micrographia.com/specbiol/alg/diato/diat0100.htm>, last viewed on 6 December 2006.

3.3. Models for tracer experiments

Tracer experiments are generally modelled by a compartmental model scheme. In this scheme every environmental reservoir is symbolised by a compartment, and connecting arrows represent the exchanges between compartments or between compartments and the exterior. The aim of a tracer experiment (and modelling) is to quantify the rates at which these exchanges occur.




Figure 3.3. Schematic representation of an *elemental process*. The box or compartment represents an environmental reservoir, characterised by the concentration of the studied compound C ($[C]$) and its isotopic composition (expressed by the atom% excess æ). The arrows represent the incoming and outgoing fluxes of material. The flowing material has its own isotopic composition (æ_in and æ_out) and is associated with a flux rate (Φ_in and Φ_out).

3.3.1. The elemental process

Every compartmental model, regardless of its complexity, can be decomposed in one or several *elementary processes*, as the one shown in Figure 3.3.

Any compartment is characterised by the concentration and by the isotopic composition for the studied compound. The concentration, usually expressed in μM ($= \mu\text{mol/l}$), is symbolised by rectangular brackets in Figure 3.3, but to avoid unnecessary heavy notation these brackets will usually be dropped. The isotopic composition is expressed in terms of atom% excess α , defined by

$$\alpha = \left(\left(\frac{[\text{special isotope}]}{[\text{sum of all isotopes}]} \right)_{\text{sample}} - \left(\frac{[\text{special isotope}]}{[\text{sum of all isotopes}]} \right)_{\text{natural}} \right) \cdot 100\% , \quad (3.1)$$

where all the isotopes of the element of interest are considered and the “special isotope” is in most cases the heaviest isotope.

The in- and outflowing material can have a different isotopic composition (α_{in} and α_{out}). The flux rates Φ_{in} and Φ_{out} (usually expressed in $\mu\text{M/h}$) are the unknowns to be quantified.

3.3.2. Assumptions and elemental equations

To translate the compartmental scheme to real model equations, a number of assumptions are to be made. The assumptions associated with the elemental process are given below (Assumptions 3.2) and will be the basis for the compartmental models further considered in this study. Note that they are distinct from the hypotheses made regarding the nature of the tracer (Assumptions 3.1). However, they are tightly linked to the experimental design since this delivers the available data to be modelled.

Assumptions 3.2 (*elemental process and equations*)

- (i) At the beginning, only the one spiked pool and the fluxes directly flowing out of this compartment are isotopically enriched.
- (ii) The enrichment of a flow at any time is equal to the abundance of the originating compartment at the moment of leaving this compartment.
- (iii) The flux rates Φ_{in} and Φ_{out} are constant, because with the endpoint experimental design only the average rate over the incubation can be determined. If real kinetic data were available, a mechanism could be postulated and time-varying flux rates could be assumed.

Note that the only exception to assumption (iii) is Chapter 5, where flux rates proportional to the source compartment concentration (first order reaction) will be considered. But for consistency reasons, only the pseudo-zero order rates are further assumed here.

With these assumptions in mind, the elemental process of Figure 3.3 can formally be transformed in mass balance equations, stating the change of concentration and atom% excess over time:

$$\begin{cases} \frac{\partial C}{\partial t} = \Phi_{in} - \Phi_{out} \\ \frac{\partial \alpha_C}{\partial t} = (\alpha_{in}(t) - \alpha_C(t)) \frac{\Phi_{in}}{C(t)} - (\alpha_{out}(t) - \alpha_C(t)) \frac{\Phi_{out}}{C(t)} \end{cases} \quad (3.2)$$

With these equations the basic building block of the compartmental tracer model is fully defined.

3.3.3. Generalised *Elskens et al.* [2002] model

During the last four decades, several models have been proposed in literature, each combining a number of these elemental processes in a different way. It would lead us too far to sum them all up; interesting overviews can be found in *Glibert et al.* [1982], *Harrison* [1983], *Glibert and Capone* [1993] and *Elskens et al.* [2005]. The exact number of compartments and fluxes considered reflect both the experimentally available data and the presumed biological processes active.

All the tracer models (even the small silicon model in Chapter 4) used in this doctoral study are special cases of the generalised model designed by *Elskens et al.* [2002] to model the transfer of ^{15}N into different particulate and dissolved nitrogen pools during a ^{15}N tracer experiment. Many of the models previously proposed in literature can also be interpreted as simplified cases (valid under some additional assumptions) of this generalised model. Therefore, it seems relevant to present this model in some more detail, although it is extensively described and validated in *Elskens et al.* [2002, 2005].

Figure 3.4 shows a schematic outline of the *Elskens et al.* [2002] model. It consists of a number of dissolved nitrogen pools (D_i , with $i = 1, \dots, n$ the nitrogen substrate) and one particulate nitrogen reservoir (P). The associated experimental setup is such that every D_i has been enriched in a separate tracer experiment, such that the influence of each D_i on P has been assessed by measuring the respective resulting isotopic enrichment in P . So for every experiment the following quantities are measured: concentration and isotopic enrichment of the spiked dissolved pool (D_i and α_{D_i}), and concentration and isotopic enrichment of the particulate pool (P and α_P). Note that no subscript i is used for P , because the particulate concentration change is supposed to be the same for all experiments (in practice the average is usually taken).

The compartments are interconnected and related to the exterior (symbolised by the source and the sink) by a number of arrows. The source term represents all processes giving rise to a regeneration of dissolved nitrogen, such as excretion by the plankton communities, bacterial transformation from one nitrogen compound to another (e.g.

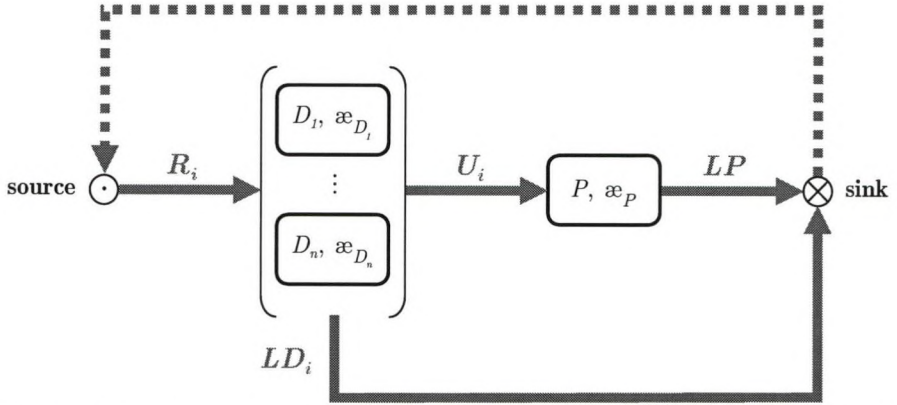


Figure 3.4. Schematic outline of the *Elskens et al.* [2002] model. See text for explanation and meaning of the symbols.

Assumptions 3.3 (*Elskens et al., 2002 model*)

- (i) *Experimental setup:* every dissolved pool is enriched in a separate experiment (n experiments), such that in each experiment only one pool is enriched at a time.
- (ii) At the beginning of each experiment, only the one spiked pool and the fluxes directly flowing out of this compartment are isotopically enriched.
- (iii) The enrichment of a flow at any time is equal to abundance of the originating compartment at the moment of leaving this compartment. Summarised, $a_{U_i} = a_{LD_i} = a_{D_i}(t)$, $a_{R_i} = 0$ and $a_{LP_i} = a_{P_i}(t)$, with i being the enriched pool.
- (iv) The flux rates are constant (see Assumptions 3.2 for motivation).

nitrification), etc. The magnitude of these regeneration flows to each dissolved pool is given by the flux rates R_i . The particulate phase can take up dissolved nutrients from the different D_i -pools. The associated uptake flux rates are symbolised by U_i . The sink term accounts for all processes responsible for missing nitrogen within the time span of the incubation, such as adsorption of nitrogen on container walls, lysis of cells, etc. The material flowing to this sink can either come from the dissolved pools (loss rates LD_i) or the particulate compartment (loss rate LP). The assumptions relevant for this model are almost identical to those for the elemental processes (Assumptions 3.2).

Combining these assumptions, a set of $3n + 1$ differential equations can be attached to the model:

$$\left\{ \begin{array}{l} \frac{\partial D_i}{\partial t} = R_i - LD_i - U_i \\ \frac{\partial \alpha_{D_i}}{\partial t} = (\alpha_{R_i} - \alpha_{D_i}(t)) \frac{R_i}{D_i(t)} \\ \frac{\partial P}{\partial t} = \sum_i U_i - LP \\ \frac{\partial \alpha_{P_i}}{\partial t} = \sum_i (\alpha_{D_i}(t) - \alpha_{P_i}(t)) \frac{U_i}{P(t)} - \alpha_{P_i}(t) \frac{LP}{P(t)} \end{array} \right. , \quad (3.3)$$

three for every tracer experiment (mass balance of D_i , α_{D_i} and α_P), and one for the concentration in the particulate pool, whose change should be the same for any incubation of equal duration, since it is independent of isotopic composition (so usually the average is taken). In its most complete form, there are $3n + 1$ unknown flux rates to be estimated: $R_1, \dots, R_n, U_1, \dots, U_n, LD_1, \dots, LD_n$ and LP . How these can be estimated, as well as their uncertainty is the subject of the next chapter.

In some situations one may have prior knowledge about the system allowing to discard one or more of these flux rates. For instance, one may argue that in some conditions denitrification is unlikely and thus loss from the nitrate pool can be neglected (or set to zero). If no such

prior knowledge is at hand, it would be useful to find a way to objectively determine *from the data* whether such conditions are present, i.e. whether some of the theoretically possible exchanges are not really occurring. This aspect will be discussed in Chapter 6. Note that in all the above statements, we simply take for granted that this set of $3n + 1$ equations contains enough information to estimate all $3n + 1$ unknown flux rates, which is not obvious at first sight. Chapter 5 discusses how this can be verified. In each of these chapters, a different specific model (class) is taken as a case study to illustrate the introduced techniques, mostly chosen based on external reasons, like temporary collaborations.

3.4. Conclusions

In this introductory chapter the basic ideas and assumptions underlying tracer studies were presented. The *Elskens et al.* [2002] model describes how total concentration and isotopic composition change with time, and enables the estimation of the relevant flux rates. The quality of these estimates will obviously depend on the quality of the measurements, but also on the assumptions not being violated. If this happens, a systematic error could be introduced, thus compromising the reliability of the outcome.

To conclude this chapter, it may be interesting to take a closer look at some of these possible limitations of the *Elskens et al.* [2002] model.

1. To construct the model equation α_{R_i} must be known. The assumption of $\alpha_{R_i} = 0$ is made, because in most cases it seems reasonable that it will take some time before enriched incorporated compounds are released back into solution. However, the risk of violating this assumption increases with increasing incubation time, implying that the incubation duration should be minimised. However, the smaller the incubation time, the smaller the changes in concentration and isotopic composition will be, converging to no detectable change at all. The flux rates are based on these changes, so with short incubations, the estimation of the flux rates will also become less reliable. In short, increasing the incubation times

possibly introduces a systematic error, thus reducing the accuracy of the flux rate estimates, while reducing the incubation times degrades their precision. Clearly no perfect solution exists and a compromise has to be struck.

2. A similar situation is encountered for α_{LP} . Here the assumption is made that $\alpha_{LP} = \alpha_P(t)$, which is reasonable if the particulate compartment is homogeneous. However, in reality the particulate phase is defined as that fraction of the sample that is collected on a filter, and hence can be composed by many constituents. For instance, if dead material is present in the particulate phase, it cannot take up any nutrients and therefore cannot be isotopically enriched. To solve this problem, the relative composition of the particulate compartment should be known, along with the isotopic signature of each constituent. This is experimentally impossible.
3. In the model structure (Figure 3.4), no exchanges between the dissolved compartments are assumed. This means that the particulate compartment can only take up enriched compounds from the spiked compartment. However, in some conditions it may be necessary to include the possibility that one dissolved compound is transformed into another dissolved compound before being taken up by the phytoplankton. For instance, in the Scheldt the conditions are such that nitrification (transformation of ammonia into nitrate) is a significant process. This additional flux can be added to the model structure and equations, but to estimate it, an additional measurement has to be included in the experimental design as well. Instead of only measuring the spiked dissolved compound (concentration and α), the connected dissolved compound also has to be measured. In the case of nitrification, this has now been implemented in our department, but this setup will not be further considered in this thesis.

Chapter 4:

Estimation of flux rates and their uncertainty

4.1. Introduction

The goal of this chapter is to estimate the parameters of a simple compartmental model and their uncertainties. In other words, some of the estimation concepts introduced in Chapter 2 are applied to one particular model of the class of compartmental models presented in Chapter 3. For the first time when modelling a tracer experiment, both the input and the output variances are taken into account, and the effect of doing this is assessed. This chapter is a transcription of my first article (but published as second) [*de Brauwere et al.*, 2005b].

4.2. Problem statement

4.2.1. The system

The system under study consists of the consumption and regeneration of silicic acid by phytoplankton (diatoms). When taking up silicic acid, diatoms transform the silicon into biosilica to form their frustules. After dying, the diatom's skeleton can dissolve such that silicic acid is released back into solution. The question is to determine the rate at which these processes occur. A general motivation to study silicon in the oceans, and in particular to better characterise the production and dissolution of biosilica, was given in section 3.2.3.

4.2.2. The measurements

To answer this question, the samples are spiked with ^{30}Si -enriched silicic acid [*Beucher et al.*, 2004a]. Due to the spike, the system is not in isotopic equilibrium anymore. The exchange of silicon is followed by sampling the dissolved pool (silicic acid) and the particulate phase (Si incorporated in phytoplankton under the form of biosilica), just after the spike ($t = 0$) and after a certain incubation period ($t = t_{inc}$). In those samples both total silicon concentration and silicon atom% excess is measured. Recall that atom% excess is a way to express the isotopic composition of the silicon pool. For Si it is defined as

$$\delta = \left(\left(\frac{[^{30}\text{Si}]}{[^{28}\text{Si}] + [^{29}\text{Si}] + [^{30}\text{Si}]} \right)_{\text{sample}} - \left(\frac{[^{30}\text{Si}]}{[^{28}\text{Si}] + [^{29}\text{Si}] + [^{30}\text{Si}]} \right)_{\text{natural}} \right) \cdot 100\% \quad (4.1)$$

4.2.3. The model

In order to extract the unknown uptake and regeneration rates from these measurements a two-compartmental model is used, which is depicted in Figure 4.1.

To translate this conceptual structure into useful model equations, a number of assumptions are necessary, based on the rationale presented in sections 3.3.2 and 3.3.3. The specific assumptions for this simple silicon model are given below.

The rationale behind assumption (iii) is based on the general assumption that only dead diatoms can dissolve, and thus regenerate silicon. Moreover, *Beucher et al.* [2004b] showed that the dissolution (regeneration) rate correlates with the percentage of dead diatoms. It is, therefore, reasonable that a finite period of time is required before any of the tracer will appear in the regenerated material [*Elskens et al.*, in press]. The risk of violating this assumption (and introducing a systematic error) increases with increasing incubation time. However, reducing the incubation time also lowers the observed changes in concentration and isotopic abundance, and hence deteriorates the precision of the final flux estimates, based on these differences. The choice of the incubation time should indeed be made with care, in order to find an acceptable compromise between accuracy and precision [see *Elskens et al.*, 2005].




Figure 4.1. The compartmental model used for estimating the uptake and regeneration rates of silicic acid. Compartments represent the pools of silicic acid (dissolved phase, D) and biosilica (particulate phase, P). The arrows symbolize the rates of silicic acid uptake (U) and regeneration (R).

Assumptions 4.1 (*Silicon model*)

- (i) *Model structure:* all the Si that leaves the dissolved pool appears as particulate biosilica and Si recycling is regarded as a process that transfers Si from the particulate to the dissolved pool.
- (ii) *Experimental setup:* at $t = 0$ only the dissolved pool is enriched with a spike of heavy Si.
- (iii) The regeneration of substrate consists of non-enriched compounds, i.e. the abundance of regenerated species equals the natural abundance ($\alpha_R = 0\%$).
- (iv) The compounds flowing out of the dissolved pool are enriched with the same isotopic abundance as this pool at the moment of leaving the compartment ($\alpha_U(t) = \alpha_D(t)$).
- (v) The flux rates R and U are constant because they represent average values over the incubation experiment duration.

Assuming this, the following differential equations can be written

$$\begin{cases} \frac{\partial D}{\partial t} = R - U \\ \frac{\partial \alpha_D}{\partial t} = -R \cdot \frac{\alpha_D}{D} \\ \frac{\partial P}{\partial t} = U - R \\ \frac{\partial \alpha_P}{\partial t} = R \cdot \frac{\alpha_P}{P} + U \cdot \frac{\alpha_D - \alpha_P}{P} \end{cases} \quad (4.2)$$

When the equations (4.2) are integrated, the actual model equations, i.e. of the form $\mathbf{y} = f_{\text{model}}(\mathbf{u}, \boldsymbol{\theta})$ (cf. Eq. (2.3)), are obtained:

$$\begin{cases} D(t) = D(0) + (R - U) \cdot t \\ \alpha_D(t) = (\alpha_D(0) - \alpha_P(0)) \cdot \left(1 + \frac{R - U}{D(0)} \cdot t\right)^{\frac{R}{U - R}} \\ P(t) = P(0) + (U - R) \cdot t \\ \alpha_P = \frac{(\alpha_D(0) - \alpha_P(0)) \cdot D(0)}{P(0) + (U - R) \cdot t} \cdot \left(1 - \left(1 + \frac{R - U}{D(0)} \cdot t\right)^{\frac{R}{U - R}}\right) \end{cases} \quad (4.3)$$

The unknown parameters are the flux rates R (regeneration rate) and U (uptake rate). From the above equations we can also see that the input variables are the measurements at the beginning of the experiment ($t=0$) or initial conditions: $D(0)$, $\alpha_D(0)$, $P(0)$ and $\alpha_P(0)$. The output variables are measurements of the same quantities but performed after the incubation ($t= t_{\text{inc}}$): $D(t_{\text{inc}})$, $\alpha_D(t_{\text{inc}})$, $P(t_{\text{inc}})$ and $\alpha_P(t_{\text{inc}})$. Due to the second and the fourth equations, this model is nonlinear in the parameters R and U and in the input variables. With four output measurements and two parameters to be estimated, the system has two degrees of freedom.

This model is a special (simplified) form of the *Elskens et al.* [2002] compartmental model, extensively discussed in Chapter 3. It was used only once before this study [*Beucher et al.*, 2004a], but then all analyses were performed assuming negligible input noise.

4.2.4. Noise assumptions

Clearly, input and output measurements are equivalent and thus it is not justified to assume that the input variations are negligible. Also, it should be noted that the input, respectively the output variables are composed of four different quantities, which are each measured in a different experiment. Consequently, the noise variance of one observation is, in general, independent and different from that of another observation (heteroscedastic data). In other words, the weighting matrix is diagonal but each element should have a different value, expressing the combined input-output uncertainty associated with the respective model equation – as defined in Eq. (2.12).

4.2.5. Aims

The first aim of this chapter is to apply the IOWLS formalism to the estimation of the Si flux rates and their uncertainties and to assess the effect of doing this. The Si model is not new (although only used once before), nor is the IOWLS formalism, but it is the first time that both are combined to achieve a more accurate estimation. For the uncertainty estimation, Monte-Carlo simulations can be performed but these can be troublesome due to long calculation times and local minima. Therefore, in this chapter it is suggested to use the Linearised Covariance Matrix (cf. section 2.7.1). This approach has the advantage of being much faster, but as the model is nonlinear, it will only provide approximate estimates of the parameter uncertainties. So, besides assessing the added value of the IOWLS formalism, this chapter is also concerned with evaluating the LCM approach for the parameter uncertainty estimation in this particular case.

The next sections are structured as follows: assessment of the use of an input-output weighting scheme for the parameter estimation (4.3.1), the

interpretation of the residual WLS cost function (4.3.2) and the parameter uncertainty estimation (4.3.3).

4.3. Results

4.3.1. IOWLS parameter estimator consistency

In this section the added value of using the IOWLS estimator versus the more classical WLS estimator is verified. We want to perform this test for one arbitrary silicon tracer experiment, thus estimating the R (regeneration rate) and U (uptake rate) for that experiment. Note that it is not possible to directly assess the accuracy of one estimation, because in real life situations the true values of R and U are unknown. Moreover, the optimised parameter values depend on the specific noise realization of the data. Indeed for replicate experiments different estimated parameter values will be found, while actually the same true processes are occurring and only the noise term added to the data differs.

Instead, simulations can be used to evaluate the *consistency* of the estimation. Simulations are synthetic data created by the user, so they provide the advantage of knowing the true parameter values. Studying the estimation's consistency serves to remove the effect of one noise realisation. This is achieved by taking into account an increasing number of data, such that the influence of noise is progressively cancelled out. Estimation consistency describes whether or not the estimated parameters converge in probability to the true values when the number of data tends to infinity. Formally, an estimator $\hat{\theta}(n_d)$ is consistent if

$$\lim_{n_d \rightarrow \infty} P\left(\left|\hat{\theta}(n_d) - \theta_0\right| > \varepsilon > 0\right) = 0, \quad (4.4)$$

with ε an arbitrary small number.

When both input and output variables are disturbed by noise, the OWLS estimator, $\hat{\theta}_{OWLS}$ as defined in Eqs. (2.8) - (2.10) is inconsistent

[e.g. *Powell and Macdonald, 1972*]. Even with an infinite number of data the exact values will not be retrieved. Using simulations, this inconsistency can be empirically confirmed, and the behaviour of $\hat{\theta}_{IOWLS}$ can be checked. In practice, an increasing number of replicate experiments were simulated (same true R and U values, but different noise realisations on the input and output data), with 12% random noise added to both the input and the output measurements.

Figure 4.2 shows the estimation results using the OWLS and IOWLS cost functions. The following observations can be made:

- (i) The OWLS estimator gives inconsistent results for both parameters R and U : indeed, in Figure 4.2a and 4.2c the mean of estimated parameter values (grey dashed line) does not converge to the true value (solid black line), but to a lower value.
- (ii) When the refined IOWLS estimator is used, this inconsistency is decreased (Figure 4.2b and 4.2d). This illustrates that the proposed refinement represents an improvement of the classical OWLS procedure.
- (iii) The \hat{R}_{IOWLS} values still exhibit a certain bias (Figure 4.2b). This is due to the fact that the proposed refinement takes the input variations into account via a linearization, which is only a first order approximation when the model is nonlinear in the input variables.
- (iv) The observed inconsistency is (for both estimators) larger for parameter R than for U .

Summarising, the simulation test illustrates that the OWLS estimator is inconsistent and that the refinement decreases the inconsistency.

Finally, most will agree that this estimator property is mostly of theoretical importance, since in practice this asymptotical limit is never reached. Nevertheless, we have included this section to show that the input-output refinement also influences the parameter estimation itself.

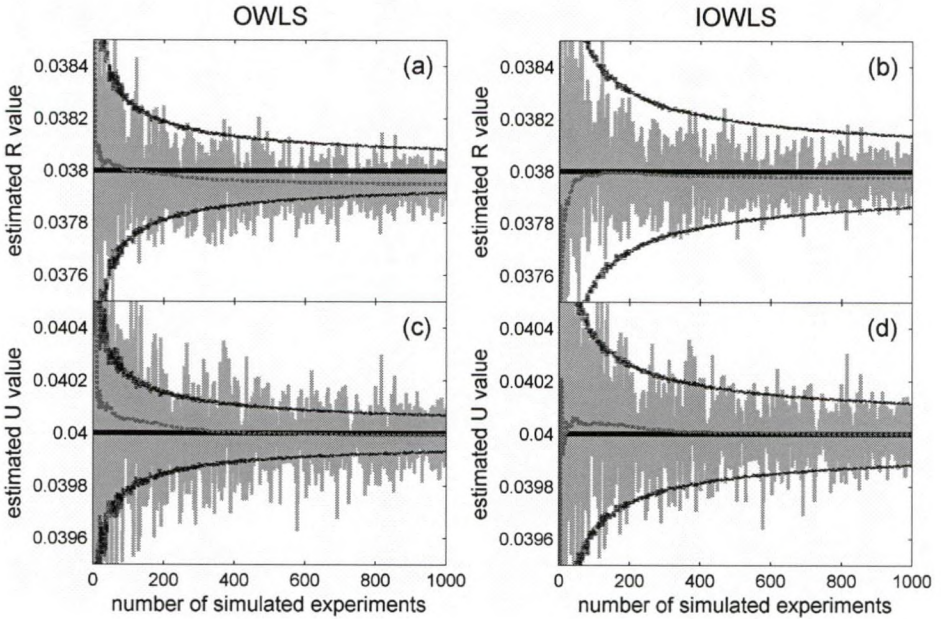


Figure 4.2. Assessment of estimation consistency. The estimated parameter values (in $\mu\text{M}/\text{h}$) are shown as a function of the number of simulated experiments used for the estimation. 12% random noise is added to both the input and the output measurements. Meaning of the symbols: the light grey solid line represents the estimated parameter value, the dashed grey line shows the mean of the estimated parameter values, the bold black line shows the true parameter value and the dotted lines indicate the estimated 95% confidence interval, estimated using the LCM expression. (a) and (b) Consistency of original WLS estimation. The estimation shows a small bias: the estimation converges to a value lower than the true value. (c) and (d) Consistency of refined WLS estimation. The estimation shows a smaller bias.

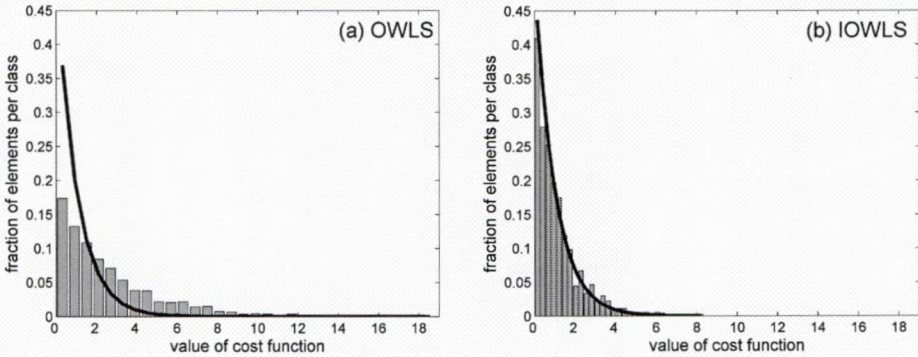
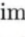



Figure 4.3. Residual cost function values shown in a normalised histogram. The normalized histogram of simulated residual cost function values () is compared to the expected probability density function with two degrees of freedom (). (a) OWLS cost functions: the histogram does not correspond to the expected distribution. (b) IOWLS cost functions: the histogram matches the expected function.

4.3.2. Residual cost function distribution

In the situation where input noise is not negligible, the residual V_{IOWLS} value is expected to be a sample from a χ^2 distribution with $n_d - n_\theta$ degrees of freedom, while the residual V_{OWLS} is expected not to be (section 2.6). These expectations are verified by Monte-Carlo simulations in this section. 1000 different noise realizations were generated for the same dataset. These 1000 artificial datasets were optimized minimizing (i) the OWLS cost function and (ii) the IOWLS cost function. Next, the distributions of the resulting 1000 residual V_{OWLS} and V_{IOWLS} values were compared to the theoretically expected χ^2 distribution. The results are displayed in Figure 4.3. The cost function values are shown in *normalised* histograms, meaning that their values are divided by the number and the width of each histogram bin, to enable a quantitative comparison with the theoretical χ^2 distribution.

Whereas the IOWLS cost function distribution seems to closely resemble the theoretical curve, the OWLS cost function does not. Rather, it

exhibits too few low values and too many high values, which is the expected pattern for an underestimation of the total variance. So, even if, perhaps, the classical OWLS procedure produces quite reliable estimated parameter values (cf. previous section), it is not possible to interpret its residual cost function in order to get information about the “acceptability” of the whole. Moreover, these results confirm that the refined weighting matrix taking into account both input and output variances is an acceptable approximation of the total variation, and clearly better than the output \mathbf{W} .

4.3.3. Parameter uncertainty estimation using the LCM

Two different tests were performed to assess the uncertainty estimation. First, a collection of 42 datasets representing real field experiments were considered, to compare the parameter uncertainty evaluation for a whole range of experimental conditions. Secondly, a simulation test was executed with increasing input noise, to assess the sensitivity of the LCM method to the input linearisation.

Test 1: field data.

The 42 experiments were performed by Charlotte Beucher in 2001-2002 in the Bay of Brest (France), following the tracer experimental setup described above. For full technical details we refer to *Beucher et al. [2004a]*. The experiments were executed over a period of a complete year (approximately one per week), so a natural array of conditions is considered. The dissolved Si concentrations range from 0.85 μM to 10.5 μM , whereas the particulate Si (biosilica) varies between 0.24 μM and 4.0 μM . The experimental relative standard deviations, derived from replicates, range from 2% to 20%. For each dataset, the parameter standard deviations were calculated via

- (i) the LCM using the output \mathbf{W} ;
- (ii) the LCM using the input-output \mathbf{W} ; and
- (iii) 1000 Monte-Carlo simulations, taken as the reference value (associated with an uncertainty of $\pm 4\%$).

The parameter standard deviations estimated with (i) discard the presence of input noise and thus are expected to be lower than those found with (ii). But how different are they, and how different from the reference estimate derived from Monte-Carlo simulations? These questions will be answered in this comparison test.

The results are shown in Figure 4.4. Because the standard deviations are of different magnitude, the results are shown as normalised values or *recoveries*, i.e. the estimated standard deviation is divided by the true one (derived from Monte-Carlo simulations). In other words, the recoveries represent the estimated values, normalised with respect to the accepted true value. Clearly and as expected, the uncertainty estimation using the output weighting matrix systematically underestimates the true parameter variation (the recoveries are systematically below 1). The average recoveries are 65% for parameter R and 62% for parameter U . On the contrary, the refined estimation is mostly a good approximation of the true variation. Sometimes it under- or overestimates the true standard deviation, but no systematic deviation is apparent. This is illustrated by the much better average recoveries: 94% and 92% for R and U respectively. This confirms that the refined uncertainty estimation is, on the whole, a good approximation of the true parameter variation. Nevertheless, it should be kept in mind that the LCM remains a linear approximation. Consequently, the recovery of the estimated standard deviations depends on the experimental conditions, but in an unpredictable way.

Actually, the LCM expression with the input-output \mathbf{W} (2.12) contains two linearisations. The first is with respect to the parameters, which is why it delivers only approximated parameter uncertainties when the model is nonlinear and the number of measurements is small. The other linearization is with respect to the input variables and originates from the input-output \mathbf{W} . If the model is linear in the input variables, this linearization is exact. Otherwise, as is the case with the Si cycling system under study, it is an approximation. The approximation could be improved by taking into account higher order terms. But this would demand more computational time. Another drawback is that higher

order moments should be accurately known, which is not obvious in general.

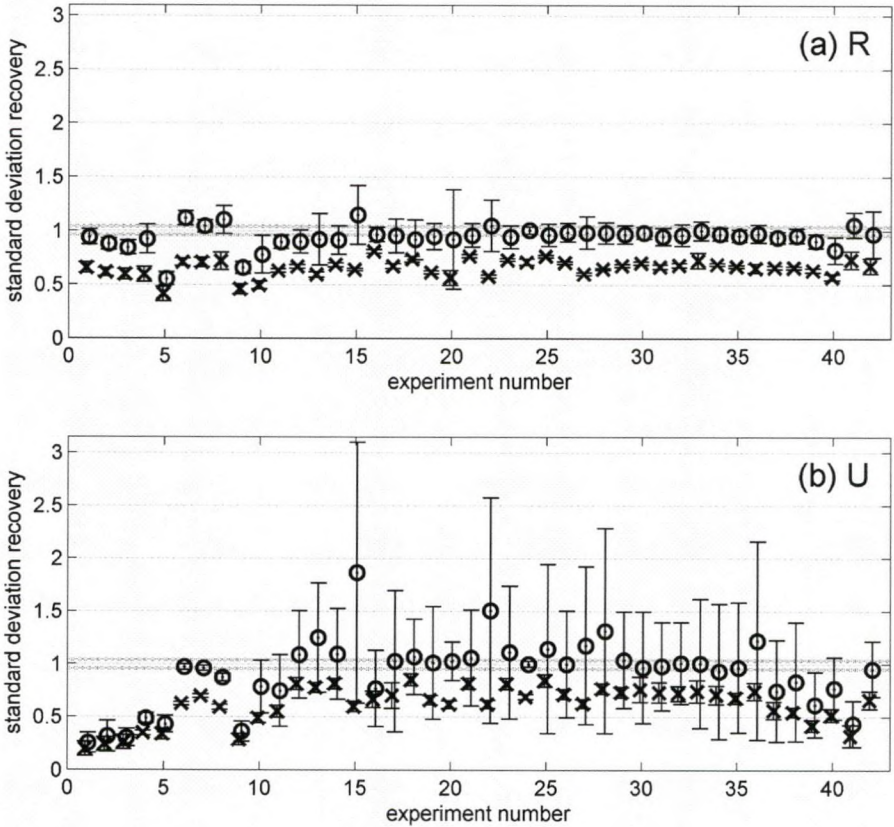


Figure 4.4. Assessment of the uncertainty estimation applied to 42 real data. The recovery achieved by the estimation using the output \mathbf{W} (\times) and using the input-output \mathbf{W} (\circ) are compared to the ideal recovery of one. By recovery is meant: the estimated standard deviation divided by the true standard deviation (as determined by Monte-Carlo simulations). The grey lines represent the approximate 95%

Test 2: simulation with increasing input noise.

In the second test, the performance of the LCM method using both \mathbf{W} s was assessed for one experimental condition (simulated), but with increasing input to output noise ratio. This exercise can help evaluate the sensitivity of the method to the linearization with respect to the input. The results are displayed in Figure 4.5. The output noise amplitude was always 5% of the output, this means that each noise realization was a sample from a $\mathcal{N}(0, 0.05 \cdot \mathbf{y})$ distribution. While the output noise amplitude was fixed, the relative input noise amplitude was varied from 0.05% to 10% of the input.

Figure 4.5 nicely illustrates that for increasing relative input noise

- (i) the true (from Monte-Carlo simulations) parameter standard deviations increase too. This is understandable, because when the measurements used to estimate the parameters are increasingly variable, the parameter estimation itself will also be subject to an enhanced uncertainty.
- (ii) the LCM estimates neglecting input noise remain approximately constant and thus are increasingly inaccurate, as expected.
- (iii) the refined standard deviation estimations do show an increasing trend, which follows relatively well the true standard deviation. Apparently, the refined method produces satisfactory estimations of the parameter standard deviations, even when the input noise level is of the same or higher magnitude as the output noise (i.e. 5 – 10% of the signal). The U uncertainty estimates again appear to be more variable and this time also slightly underestimate the true uncertainty when the input noise level increases. This is possibly a consequence of the linearisations contained in LCM, as discussed above.

Finally, note that when the input noise becomes small with respect to the output noise, the figures confirm that the refined estimation converges to the original estimation. This is expected since the refined weighting matrix converges to the original one as the input noise variance tends to zero.




Figure 4.5. Assessment of the parameter uncertainty estimation for increasing input noise. The estimated standard deviation for U and R estimated using the output \mathbf{W} (\times) and using the input-output \mathbf{W} (\circ) are compared to the standard deviation derived from 1000 Monte-Carlo simulations (—). The error bars represent the approximate 95% confidence intervals of the standard deviation estimates. The relative output noise amplitude is always 5% of the output. So, the relative input noise amplitude varies from 0.05% to 10% of the input.

4.4. Summary

The model studied in this chapter is designed to describe the uptake and regeneration of Si in pelagic waters. When such a nonlinear model is used for describing measurements which are all disturbed by noise, it is not so obvious to accurately estimate the model parameters and, even less obvious, the parameter uncertainties. This problem was tackled using the IOWLS approach (cf. section 2.5.4), combined with the LCM (cf. 2.7.1). However, both contain linearisations which are approximations in the case of the Si model, because this is nonlinear both with respect to the parameters and with respect to the input variables. The aim of this chapter was to assess the improvement achieved by taking into account the input variances thanks to the IOWLS formulation, and at the same time evaluate the error introduced by the linear approximations made.

As a first result, a comparison between the OWLS and IOWLS estimator on a hypothetical data set confirmed that the latter exhibits a smaller bias. Secondly, the analysis of the residual cost function values of both estimators showed that only the IOWLS cost function is nicely χ^2 distributed, corroborating that the input-output weighting matrix is a good reflection of the total variance, and certainly better than the weighting matrix only considering output variances. Using this statistical feature of the IOWLS cost function, it is possible to identify systematic errors in individual measurement-model results. This property is important because it allows an internal quality control of the results and it can be used to select the most appropriate model from a set of models for a given dataset (cf. Chapter 6).

Further, the estimation of the parameter uncertainty using the LCM expression was tested on field data of dissolved and particulate Si in the Bay of Brest (France) from *Beucher et al.* [2004a]. The estimated standard deviations were compared to values derived from Monte-Carlo simulations taken to be the “true” standard deviations. Overall, the recovery was good for the input-output method (average recoveries of 94% and 92% for R and U respectively), whereas with the output \mathbf{W} the

estimations systematically and significantly underestimated the parameter variations (average recoveries of 65% and 62%).

It can be concluded that the use of an IOWLS cost function improves the consistency of the estimation and enables the statistical interpretation of the residual cost function. The first property may not be of much practical importance, the second certainly is. In addition, parameter uncertainty estimates derived from the LCM combined with the input-output weighting matrix appear to be reasonably reliable. Although these estimates are probably still less accurate than those derived from extensive Monte-Carlo simulations, the LCM estimation is attractively fast and easy to automate. Therefore, the IOWLS – LCM combination seems to be a satisfactory approach for the estimation of parameters and their uncertainties in the case of this biogeochemical model.

Chapter 5:

Identifiability of two-compartments models

“It is a capital mistake to theorize before one has data.”

Sir Arthur Conan Doyle, in: ‘The adventures of Sherlock Holmes’.

5.1. Introduction

Before executing any experiment or parameter estimation, it should always be checked whether the model parameters are fully identifiable with the intended measurements. This issue is called *a priori identifiability* and “deals with the problem of determining whether a planned input-output identification experiment is able to supply the desired information about the unknown parameters of the (...) model” [Cobelli et al., 1979]. For a more rigorous definition we refer to Cobelli et al. [1979]. This is valuable knowledge because, if known in advance, it avoids that any worthless experiments are performed. A priori identifiability depends only on the assumed model structure and the chosen experimental setup. Therefore, it is a *structural* property and it is often referred to as *structural identifiability*. In the logical steps of identification (cf. Chapter 2), the analysis of a priori identifiability comes after (i) the necessary measurements are decided *but not made* and (ii) the model structure is chosen. In other words, it is set in the context of ideal perfection by considering noiseless observations and an error-free model. If the test concludes that the intended combination is unidentifiable, either the model structure or the experimental design has to be modified (or both). Alternatively, if the result of the test confirms a priori identifiability, the experiment can be executed and the model parameters can be estimated.

Besides a priori identifiability, one can also define an *a posteriori* identifiability, to be determined *after* the data are collected and the model parameters are estimated. Due to particular data values and/or measurement noise, some parameters may in practice be associated with such a large uncertainty that they can be said to be *a posteriori* unidentifiable – although they were checked to be *a priori* identifiable. The following example from bivariate regression illustrates both concepts: fitting a line through two distinct points is perfectly *a priori* identifiable, but will be poorly *a posteriori* identifiable if the points are associated with high uncertainty and lie close to each other.

In this chapter, we are interested in the a priori identifiability of a class of compartmental models, relevant in the framework of modelling tracer

experiments. For the ease of formulation, we will from now on simply speak of identifiability, implying a priori identifiability. For simple model structures its analysis may be trivial, but in the case of nonlinear model structures it is not always straightforward to determine identifiability simply from looking at the model equations. The family of compartmental models is a special case of nonlinear models where identifiability is a recognised issue. Especially in the domains of pharmacology and ecology a number of studies on this topic have been published [e.g. *Cobelli and Caumo*, 1998; *Brun et al.*, 2001]. However, it appeared difficult to directly use the results from these studies for the environmental tracer models, due to a number of reasons.

- (i) The experimental design in pharmacological and ecological applications is usually different from the one used for environmental tracer studies. For instance, often real kinetic experiments are performed, i.e. the pools are sampled at many time instances between the start and the end of the experiment (versus only two measurements in time in the environmental tracer studies). Also, in pharmacological applications rarely more than one tracer is considered (versus two natural isotopes in environmental applications). In ecological studies usually no tracer is added at the beginning of the experiment, so actually the experiment is just a monitoring of the ecological system starting at an arbitrary time.
- (ii) Only a few specific (widely used) model structures are analysed in the pharmacological identifiability literature.
- (iii) The major problem in ecological studies is that the models are overparameterised compared to the number of observations available. No elaborate analysis is necessary to see that this situation is highly unidentifiable. However, these models are often intended for prediction, requiring an extrapolation outside the observed domain. Poorly identifiable parameters may be relevant to processes that become important in the extrapolation domain but are of less significance in the observed domain. "*Discarding such processes so that only models with identifiable parameters are considered can lead to a significant underestimation of the prediction uncertainty*" [*Reichert and Omlin*, 1997]. Therefore, the identifiability issue in ecological modelling is rather to optimally

select subsets of identifiable parameters, while fixing the other parameters at a priori values [Reichert and Omlin, 1997; Brun et al., 2001; Andersson et al., 2006].

Consequently, it seemed useful to perform our “own” general identifiability study. We decided to focus on the class of two-compartment models, because it is still manageable to consider all (63) of them. We combined these 63 models with 12 different experimental set-ups, containing the ones used in pharmacology as well as in environmental studies. The result is a “catalogue” with identifiability information for all models consisting of two compartments. In the future, it can now easily be checked whether the intended experiment-model combination delivers an identifiable entity. Below is explained how this identifiability study was achieved and the resulting “catalogue” is shown.

5.2. Methods

5.2.1. Testing identifiability

“A necessary and sufficient condition for structural identifiability is that the number of unknown parameters (n_θ) equals the number of independent equations (...) among them” [Cobelli et al., 1980].

The number of independent equations can be quantified as the rank of the Jacobian matrix (cf. Eq. 2.16). Therefore, this identifiability study will be based on a comparison of the Jacobian rank with the number of unknown parameters. The minimalism of this procedure is its main advantage. Indeed, it is easy to automate such that the large number of tests necessary for our catalogue can be executed (almost) automatically. However, a number of additional remarks are useful.

- (i) For complicated models, this matrix may be cumbersome to evaluate. Therefore, some necessary conditions for structural identifiability have been proposed in compartmental analysis which are based on the topology of the compartmental diagram [Cobelli et al., 1980; Saccomani et al., 1994]. If the necessary conditions are

not satisfied, the investigator knows that the system is not identifiable and is spared from the more elaborate matrix evaluation. Other methods proposed in literature to determine identifiability of compartmental models include the analysis of the Laplace transforms of the observed variables in terms of the unknowns of the model [*Cobelli and DiStefano*, 1980] and investigation of the Jacobian matrix of the Markov parameter matrix [*Bellman and Åström*, 1970].

- (ii) Even if all model parameters are identifiable, note that their uniqueness is not assured. Indeed, only the *local* identifiability is assessed by the Jacobian rank. This means that in the “neighbourhood” of the parameter values used to evaluate the Jacobian matrix (cf. Eq. 2.16) there is no other set of parameter values that describes the output observations equally well. Hence, it is still possible that a finite number of equally good solutions exists in parameter space, for instance as a result of internal symmetry in the model. This “uniqueness” of the solution, or the *global* identifiability of the system is important to assess, especially when the parameters are analogues of physical attributes and the model is meant to quantify and eventually interpret them. For black box models, whose role is rather to optimally describe observations, this issue may be of less importance.
- (iii) If the rank of the Jacobian matrix is lower than the number of unknown parameters, some structural relation exists between the parameters and they cannot all be uniquely identified, not even locally. The disadvantage of the Jacobian rank approach is that this is the only information directly available. To characterise the existing parameter interrelations, further analysis is needed. When analysing a particular system, this is certainly important, but not within the scope of this study. Instead, the aim is to provide a first identifiability diagnosis for a class of model-experiment setups, returning only the number of solutions: infinite (unidentifiable), one, two, etc.

5.2.2. Investigated models and experimental setups

The aim is to investigate all two-compartments model structures, only differing in the number and position of the connecting fluxes (see Figure 5.1 and Tables 5.3 and 5.4).

In addition, twelve different experimental setups are considered (Table 5.1), comprising the situations often encountered in pharmacological studies, as well as in

“our” environmental tracer study. In all studies it is assumed that at $t = 0$ compartment 1 (Figure 5.1) is spiked with a tracer, which can either be a radioactive or stable isotopic analogue of the species of interest (e.g. glucose or ammonium). In pharmacological applications, often only the labelled species is measured (i.e. not the natural, unlabeled species), either in one or both of the compartments. At the same time, it is customary to perform real kinetic experiments, i.e. sampling the tracer at several time instances. In contrast, in the environmental tracer studies, both the labelled and the unlabelled species are measured, but only once, at the end of the incubation experiment (cf. Chapter 3). As a reflection of this diversity of existing experimental schemes, different combination of sampling frequencies, sampled compartments and measured species are considered (cf. Table 5.1). A total incubation time of 24 (hours) is chosen, which is realistic in environmental tracer studies. As a result, the endpoint (environmental) design only considers one sampling after this time period, while in the kinetic (pharmacological) setup, samples are supposed to be taken once every hour – explaining why the a priori strange choice of 24 sampling times was used (cf. Table 5.1). An additional sampling frequency is considered of two samplings during the experiment (resp. after 12 and 24 (hours)). This setup was added, because for the more complex models performing only one final measurement simply does not provide enough observations to identify all rate constants. This result of unidentifiability due to fewer observations than unknowns is rather trivial; indeed it is more interesting to investigate the effects of model structure on the system identifiability.




Figure 5.1. Diagram of the full two-compartment model. The two compartments (circles) are not interchangeable, because only compartment 1 is spiked at $t = 0$ (represented by the open arrow). The six grey arrows symbolize all possible exchanges between the compartments or between a compartment and the exterior. The exchanges are assumed to be first order processes (characterised by the rate constants k), except for those coming from the exterior which are constant (characterised by the flux rates Φ).

Table 5.1. Overview of the twelve experimental setups considered. Q stands for the labelled species concentration and q for the unlabeled species concentration, with the subscript indicating which compartment was measured.

# species measured	Sampled compartments	Sampling times (excl. initial spike)	Measurements
1	Comp. 1	1	$Q_1(t_1)$
		2	$Q_1(t_1), Q_1(t_2)$
		24	$Q_1(t_1), \dots, Q_1(t_{24})$
	Comp. 2	1	$Q_2(t_1)$
		2	$Q_2(t_1), Q_2(t_2)$
		24	$Q_2(t_1), \dots, Q_2(t_{24})$
	Comp. 1 & 2	1	$Q_1(t_1), Q_2(t_1)$
		2	$Q_1(t_1), Q_2(t_1), Q_1(t_2), Q_2(t_2)$
		24	$Q_1(t_1), Q_2(t_1), \dots, Q_1(t_{24}), Q_2(t_{24})$
2	Comp. 1 & 2	1	$Q_1(t_1), q_1(t_1), Q_2(t_1), q_2(t_1)$
		2	$Q_1(t_1), q_1(t_1), Q_2(t_1), q_2(t_1), Q_1(t_2), q_1(t_2), Q_2(t_2), q_2(t_2)$
		24	$Q_1(t_1), q_1(t_1), Q_2(t_1), q_2(t_1), \dots, Q_1(t_{24}), q_1(t_{24}), Q_2(t_{24}), q_2(t_{24})$

Combining the model structure (Figure 5.1) and the most complete experimental setup (last row of Table 5.1), one can formalise the problem in the following differential equations

$$\begin{cases} \frac{\partial Q_1}{\partial t} = -(k_{01} + k_{21})Q_1 + k_{12}Q_2 + \alpha_{nat}\Phi_{10} \\ \frac{\partial q_1}{\partial t} = -(k_{01} + k_{21})q_1 + k_{12}q_2 + (1 - \alpha_{nat})\Phi_{10} \\ \frac{\partial Q_2}{\partial t} = -(k_{12} + k_{02})Q_1 + k_{21}Q_2 + \alpha_{nat}\Phi_{20} \\ \frac{\partial q_2}{\partial t} = -(k_{12} + k_{02})q_1 + k_{21}q_2 + (1 - \alpha_{nat})\Phi_{20} \end{cases} \quad (5.1)$$

These equations express how the concentrations of two species (Q and q) change with time in both compartments, after compartment 1 has been spiked with a certain amount of labelled species (Q) at $t = 0$. The incoming (constant) flux rates Φ_{10} and Φ_{20} have the isotopic composition of the natural species. Therefore, they are multiplied by α_{nat} , which is the natural isotopic abundance. Taking the example of environmental nitrogen studies, where ^{15}N analogues of the nutrients are used as tracer, α_{nat} is 0.365%, meaning that in nature 0.365% of the nitrogen atoms are ^{15}N rather than ^{14}N .

The differential equations (5.1) can be summarised in matrices

$$\dot{\mathbf{y}} = \mathbf{A}\mathbf{y} + \mathbf{b}, \quad (5.2)$$

with

$$\mathbf{y} = \begin{pmatrix} Q_1 \\ q_1 \\ Q_2 \\ q_2 \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} -k_{01} - k_{21} & 0 & k_{12} & 0 \\ 0 & -k_{01} - k_{21} & 0 & k_{12} \\ k_{12} & 0 & -k_{12} - k_{02} & 0 \\ 0 & k_{12} & 0 & -k_{12} - k_{02} \end{pmatrix},$$

$$\mathbf{b} = \begin{pmatrix} \alpha_{nat} \Phi_{10} \\ (1 - \alpha_{nat}) \Phi_{10} \\ \alpha_{nat} \Phi_{20} \\ (1 - \alpha_{nat}) \Phi_{10} \end{pmatrix} \quad (5.3)$$

and $\dot{\mathbf{y}}$ symbolises the time-derivative of \mathbf{y} .

To estimate the unknown parameters (Φ_{10} , k_{01} , k_{21} , k_{12} , k_{02} , Φ_{20}), these equations must be integrated to the time instances at which the species are measured, delivering the model values for the measured concentrations. For the most complete experimental setup this means that model values for $\mathbf{y}_t = [Q_1(t_1) \ q_1(t_1) \ Q_2(t_1) \ q_2(t_1) \ \dots \ Q_1(t_{24}) \ q_1(t_{24}) \ Q_2(t_{24}) \ q_2(t_{24})]^T$ are computed. The integrated model is not linear in the parameters, i.e. it cannot be written as Eq. (2.4). Therefore, a nonlinear optimisation procedure is needed to find the optimal parameter values. This requires the construction of the Jacobian matrix (cf. Eq. (2.16)), whose rank also gives information about identifiability.

5.2.3. Procedure

Briefly, these were the different steps of the procedure:

1. Choose arbitrary initial conditions for \mathbf{y} , i.e. $\mathbf{y}_0 = [Q_1(t=0) \ q_1(t=0) \ Q_2(t=0) \ q_2(t=0)]$. Two initial situations were considered:
 - 1a. The modelled species is not naturally present in the system, or at least not present in compartment 2 at the beginning of the experiment: $Q_2(t=0) = q_2(t=0) = 0$.
 - 1b. $Q_2(t=0)$ and $q_2(t=0) \neq 0$.
2. Define all models in terms of parameters, and assign arbitrary test values to each model parameter set.
3. Repeat the following for each of the 63 model structures.
4. Repeat the following for each of the 12 experimental setups.

5. Compute the Jacobian matrix \mathbf{J} . This is done by numerically integrating equations (5.1), which are first analytically differentiated with respect to the parameters:

$$\mathbf{J} = \frac{\partial \mathbf{y}}{\partial \boldsymbol{\theta}} = \int_{t=0}^{t=t_{inc}} \frac{d}{dt} \left(\frac{\partial \mathbf{y}}{\partial \boldsymbol{\theta}} \right) dt = \int_{t=0}^{t=t_{inc}} \frac{\partial}{\partial \boldsymbol{\theta}} \left(\frac{d\mathbf{y}}{dt} \right) dt = \int_{t=0}^{t=t_{inc}} \frac{\partial \dot{\mathbf{y}}}{\partial \boldsymbol{\theta}} dt. \quad (5.4)$$

The numerical integration is performed for the chosen initial conditions and parameter values, and simultaneously with the integration of the model equations (5.1).

6. Identifiability test: if the $rank(\mathbf{J}) < \text{number of model parameters}$ or if the $cond(\mathbf{J}) > 10^{9(**)}$, then the model-experiment is classified as unidentifiable. Otherwise it is identifiable, but maybe not uniquely.
7. For every identifiable design, a crude test was performed to evaluate for the presence of multiple identifiable parameter solutions. 100 simulations were performed in which the parameter values were really optimised, in each simulation starting from a different (random) initial guess for the parameters. If in each of the 100 simulations, the final optimised parameter values approximate the true values, it is concluded that the model-experiment combination is uniquely identifiable. Otherwise, there are multiple parameter sets that satisfy the data equally well. From the 100 simulations, a first estimate can be made concerning the number of optimal parameter sets, but clearly the test is too crude to guarantee these results. For all the results it should be noted that they remain only local, because the initial parameter values were generated as perturbations around the true values (and not from $-\infty$ to ∞).

(**) The reference value to determine a well-conditioned Jacobian is chosen here in relation with the precision of the numerical integration of the equations. A relative precision (option "RelTol" for the ode45 MATLAB integration function) of 10-9 is used, so the elements of the Jacobian matrix will be computed within this precision. If the matrix is badly conditioned this will also be seen within this tolerance.

5.3. Results

The results of this identifiability study of two-compartments models are summarised in Tables 5.3 and 5.4, one table for each initial condition considered (cf. 1a and 1b of the procedure in the previous section). The simulation conditions are summarised in Table 5.2. The results in Tables 5.3 and 5.4 were validated by cross-checking with the results available in the literature (only 7 designs).

The first large-scale observation is that no pattern is apparent. Secondly, even when neglecting the model-experiment combinations which are unidentifiable simply due to fewer observations than unknowns, still a considerable number of designs appears to be unidentifiable. This is real structural unidentifiability, and may not always be obvious at first glance. It is also worth noting that the identifiability results are not identical for the two initial conditions, which was not really expected at first. Obviously, this is why the two initial conditions are considered and their results are presented separately. Considerably more structures are identifiable if initially compartment 2 is not empty.

Tables 5.3 and 5.4 form the “catalogue” mentioned in the introduction. Now it is possible to check any intended model-experimental setup before executing any expensive action. Unfortunately, there hasn’t been an opportunity to take this last step to the real application. The main reason is that it was not relevant for the compartmental models used in the framework of environmental tracer studies in our research group ANCH. First of all because at the time of this work the models assumed “zero order” processes (cf. section 3.3.2), instead of first order exchanges assumed here (cf. section 5.5.2). Moreover, thanks to experimental advances the models are progressively enlarged to contain more than two compartments. Only a few models still contain as little as two compartments, and the existing two-compartments models (e.g. the Si model in Chapter 4) were already tested individually.

Table 5.2. Overview of numerical values used. Initial conditions represent concentrations, so they could be expressed in μM . The parameters flux rates and rate constants, so their dimension could be respectively $\mu\text{M}\cdot\text{h}^{-1}$ and h^{-1} , if the incubation times are expressed in hours (see section 5.2.2). Results for situation (1a) are shown in Table 5.3; those for situation (1b) are given in Table 5.4.

	<i>Initial conditions</i>				<i>Parameter values (if flux is</i>					
	$Q_1(0)$	$q_1(0)$	$Q_2(0)$	$q_2(0)$	Φ_{10}	k_{01}	k_{21}	k_{12}	k_{02}	Φ_{20}
(1a)	19	37	0	0	0.5	0.01	0.02	0.03	0.04	0.5
(1b)	43	11	37	7						

Table 5.3. Overview of identifiability results for all 63 two-compartments models and 12 experimental setups, with initial conditions 1a (compartment 2 is empty at the beginning). On the second row, the measured species are given; on the third row the sampling time instances: [24] = only once after 24h, [12 24] = two samplings, resp. after 12h and after 24h, [1:24] = 24 samplings, after 1h, ..., 24h. Time units are arbitrary since only simulations are considered. The symbols in the table represent the number of parameter solutions for the model-experiment combination of interest: ∞ means infinite number of possible parameter solutions, i.e. unidentifiable design; 1 means one parameter solution, i.e. a (locally) uniquely identifiable design; 2 or more counts the number of parameter solutions. The unidentifiable combinations are shaded: light grey background if the unidentifiability is due to a too low number of observations (lower than the number of unknown parameters), dark grey if the unidentifiability is caused by structural features.

No.	diagram	identifiability											
		Q_1			Q_2			Q_1			Q_2		
		[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]
6.01		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.01		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.02		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.03		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.04		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.05		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.06		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
4.01		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
4.02		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.03		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	2
4.04		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
4.05		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.06		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.07		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1

Table 5.3 continued.

No.	diagram	identifiability											
		Q_1 q_1 Q_2 q_2			Q_1 Q_2			Q_1			Q_2		
		[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]
4.08		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
4.09		1	1	1	∞	1	1	∞	∞	1	∞	∞	3
4.10		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.11		∞	1	1	∞	1	1	∞	∞	1	∞	∞	∞
4.12		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
4.13		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.14		1	1	1	∞	1	1	∞	∞	∞	∞	∞	2
4.15		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
3.01		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.02		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
3.03		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
3.04		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.05		1	1	1	∞	1	1	∞	∞	2	∞	∞	1
3.06		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.07		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.08		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
3.09		∞	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.10		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.11		∞	1	1	∞	1	1	∞	∞	1	∞	∞	2
3.12		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	2
3.13		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.14		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
3.15		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.16		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞

Table 5.3 continued.

No.	diagram	identifiability											
		Q_1 q_1 Q_2 q_2			Q_1 Q_2			Q_1			Q_2		
		[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]
3.17		∞	1	1	∞	1	1	∞	∞	1	∞	∞	1
3.18		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
3.19		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.20		∞	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
2.01		1	1	1	∞	1	1	∞	1	1	∞	∞	∞
2.02		1	1	1	1	1	1	∞	1	1	∞	1	1
2.03		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
2.04		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
2.05		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.06		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.07		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
2.08		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
2.09		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.10		∞	1	1	∞	1	1	∞	1	1	∞	1	1
2.11		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.12		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.13		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
2.14		1	1	1	1	1	1	∞	1	1	∞	1	1
2.15		∞	1	1	∞	1	1	∞	∞	∞	∞	1	1
1.01		1	1	1	1	1	1	1	1	1	∞	∞	∞
1.02		1	1	1	1	1	1	1	1	1	∞	∞	∞
1.03		1	1	1	1	1	1	1	1	1	1	1	1
1.04		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
1.05		∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
1.06		1	1	1	1	1	1	∞	∞	∞	1	1	1

Table 5.4. Same as Table 5.3, but with initial conditions 1b (compartment 2 is not empty at $t = 0$). Identifiability results are slightly different than in Table 5.3.

No.	diagram	identifiability											
		Q_1 q_1 Q_2 q_2			Q_1 Q_2			Q_1			Q_2		
		[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]
6.01		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.01		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.02		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.03		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.04		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.05		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
5.06		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
4.01		1	1	1	∞	1	1	∞	∞	1	∞	∞	1
4.02		1	1	1	∞	1	1	∞	∞	2	∞	∞	∞
4.03		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	2
4.04		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
4.05		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.06		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.07		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
4.08		∞	1	1	∞	∞	1	∞	∞	∞	∞	∞	∞
4.09		2	1	1	∞	1	1	∞	∞	1	∞	∞	1
4.10		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.11		1	1	1	∞	1	1	∞	∞	2	∞	∞	∞
4.12		∞	1	1	∞	∞	1	∞	∞	1	∞	∞	∞
4.13		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
4.14		1	1	1	∞	1	1	∞	∞	∞	∞	∞	2
4.15		2	1	1	∞	1	1	∞	∞	1	∞	∞	1

Table 5.4 continued.

No.	diagram	identifiability											
		Q_1 q_1 Q_2 q_2			Q_1 Q_2			Q_1			Q_2		
		[24]	[12:24]	[1:24]	[24]	[12:24]	[1:24]	[24]	[12:24]	[1:24]	[24]	[12:24]	[1:24]
3.01		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.02		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.03		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.04		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.05		2	1	1	∞	1	1	∞	∞	1	∞	∞	1
3.06		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.07		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.08		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.09		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.10		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.11		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
3.12		1	1	1	∞	1	1	∞	∞	∞	∞	∞	2
3.13		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.14		1	1	1	∞	1	1	∞	∞	2	∞	∞	∞
3.15		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
3.16		1	1	1	∞	1	1	∞	∞	∞	∞	∞	∞
3.17		1	1	1	∞	1	1	∞	∞	2	∞	∞	1
3.18		1	1	1	∞	1	1	∞	∞	1	∞	∞	2
3.19		1	1	1	∞	1	1	∞	∞	∞	∞	∞	1
3.20		1	1	1	∞	1	1	∞	∞	1	∞	∞	∞
2.01		1	1	1	∞	1	1	∞	1	1	∞	∞	∞
2.02		1	1	1	1	1	1	∞	1	1	∞	1	1
2.03		1	1	1	1	1	1	∞	1	1	∞	∞	∞

Table 5.4 continued.

No.	diagram	identifiability											
		Q_1 q_1 Q_2 q_2			Q_1 Q_2			Q_1			Q_2		
		[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]	[24]	[12 24]	[1:24]
2.04		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.05		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.06		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.07		1	1	1	1	1	1	∞	1	1	∞	∞	∞
2.08		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.09		1	1	1	1	1	1	∞	∞	∞	∞	∞	∞
2.10		1	1	1	∞	1	1	∞	1	1	∞	1	1
2.11		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.12		1	1	1	1	1	1	∞	∞	∞	∞	1	1
2.13		1	1	1	1	1	1	∞	1	1	∞	∞	∞
2.14		1	1	1	1	1	1	∞	1	1	∞	1	1
2.15		1	1	1	∞	1	1	∞	∞	∞	∞	1	1
1.01		1	1	1	1	1	1	1	1	1	∞	∞	∞
1.02		1	1	1	1	1	1	1	1	1	∞	∞	∞
1.03		1	1	1	1	1	1	1	1	1	1	1	1
1.04		1	1	1	1	1	1	1	1	1	1	1	1
1.05		1	1	1	1	1	1	∞	∞	∞	1	1	1
1.06		1	1	1	1	1	1	∞	∞	∞	1	1	1

5.4. Conclusions

Although it hasn't been possible to take this chapter to the final step of application to a real world problem, I absolutely wanted to include it. Identifiability is an important and interesting issue, not yet recognised in all communities. It is an essential step in the identification process, especially considering that it can avoid useless but expensive data gathering. Furthermore, in the second part of this thesis, the issue of identifiability will recur (Chapter 10), showing its generic importance.

So far for the relevance of investigation identifiability. Now some final remarks can be made regarding the method used. First, it could be argued that the Jacobian rank is an “ugly”, numerical method, in contrast to more appealing theoretical or a priori methods only based on the model structure. This kind of methods exists for some types of model structure but their drawback is exactly that they are only practical for that particular model class. The Jacobian rank approach, on the contrary, is very general and therefore easy to automate. For instance, it will be used again in Chapter 10 in the context of the very different water mass mixing model. Also, the Jacobian matrix can be used to rank the importance of the model parameters measured as their impact on the model output (i.e. sensitivity). Some formulas have been suggested to enable such rankings in the framework of large, overparameterised models [Brun *et al.*, 2001], but this kind of parameter ordering can be interesting in any work on experimental design.

Finally, the Jacobian matrix delivers information only about the *local* identifiability of *individual* parameters. For global analyses and to make statements about the joint influence of groups of parameters a more sophisticated analysis is needed [Turányi, 1990].

Chapter 6:

Model selection

“All models are wrong but some are useful.”

Chatfield [1995].

6.1. Introduction

The ultimate question in applied modelling is to determine which is the true mechanism that generated a given dataset. Unfortunately, answering this question for real life applications is illusional. Instead, the question is often reformulated as choosing that model from a set of available models that, according to the available data, seems most suitable to describe the processes. This is the central question to be answered in this chapter and it will be referred to as the issue of *model selection*.

The difficulty of model selection is that no straightforward and universal criterion exists that defines which model is most suitable. Goodness-of-fit of a model is one aspect that should be considered, but it has to be balanced with the model parameters' variability. Indeed, if attention is only concentrated on the models' ability to fit the data, one will tend to select too complex (i.e. over-parameterised) models. These models have the attractive feature to fit the given observations very well. But the price paid is a high sensitivity to small changes in the data, because they misinterpreted at least a part of the coincidental noise as relevant underlying mechanism. In the limit, the model would simply store the data values, and be nothing more than a data replicator [*Pitt et al.*, 2003]. Consequently, overparameterised models will fail to account for new data. In this view, the problem of model selection can be reformulated as determining what part of the measurements is due to stochastic variations, which should be ignored, in order to capture only those variations caused by the processes of interest. The model that disentangles signal from noise most effectively is then, based on the information at hand, closest to the true model and should be selected as the best model. As a consequence, even for a given process, the model selected as the best one depends on the available data and the knowledge about their uncertainty.

If the quality and quantity of the data was sufficient, the selected model will be able to accurately predict future data. Model selection is indeed

tightly related to prediction^{**}. For instance, M. R. Forster states that “*the goal of model selection is to maximize the predictive accuracy of (...) rival models*” [Forster, 2001]. However, model selection only has a predictive power as far as the processes under study are stationary, and the selected model will thus accurately predict new data as far as they are either replicates or interpolations.

Model selection is performed intrinsically in every parameter estimation problem, since unavoidably a model has been chosen to represent the data. Yet, this is often done in a subjective way: based on experience or subjective judgement. Clearly, there is a need for a well-defined, objective model selection strategy. Several such schemes are described in literature, depending on the available noise information [Schoukens *et al.*, 2002], sample size [De Ridder *et al.*, 2005] and prior knowledge about the true underlying data generating mechanism [de Luna and Skouras, 2003].

In this chapter, we consider the situation where the noise variances associated with the data are known or estimated in a pre-processing step. Furthermore, we are especially interested in the performance of model selection strategies in the case of small sample sizes. First, a review of existing methods is given (section 6.2), followed by the presentation of a new method (section 6.3). This method will be tested on simulations, and compared to some of the existing methods in section 6.4. Finally, the new method is applied to a set of field data revealing some interesting patterns (section 6.5).

The contents of this chapter, in particular the method description and the results, are based on my second publication, *de Brawwere et al.* [2005a].

^{**} Note that here we use the term “prediction” in the common and somewhat vague meaning of “saying something about new data”. Nothing is implied about the method of deriving these new data.

6.2. Existing methods

6.2.1. AIC and MDL

Intuitively, a good model fits the data well. Hence, it will be associated with low residuals or, expressed in one number, with a low residual cost function (sum of squared residuals)

$$V_{LS}(\hat{\boldsymbol{\theta}}) = \mathbf{e}(\hat{\boldsymbol{\theta}})^T \mathbf{e}(\hat{\boldsymbol{\theta}}). \quad (6.1)$$

$\hat{\boldsymbol{\theta}}$ represents the estimated model parameters, found by minimizing the cost function (6.1)

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} V_{LS}(\boldsymbol{\theta}), \quad (6.2)$$

and $\mathbf{e}(\hat{\boldsymbol{\theta}})$ is the vector of associated minimal residuals. One model selection strategy considers the minimization of a criterion of the general form:

$$V(\hat{\boldsymbol{\theta}}) \cdot p(n_d, n_\theta), \quad (6.3)$$

with $V(\hat{\boldsymbol{\theta}})$ the minimal value of the cost function as defined in Eq. (6.1)-(6.2), n_d is the number of observations and n_θ the number of free model parameters. The penalty factor $p(n_d, n_\theta)$ corrects $V(\hat{\boldsymbol{\theta}})$ for overfitting, therefore a more complicated model will only be accepted if its improvement in goodness-of-fit is large enough, i.e. balances the penalty factor. To put it simply, expression (6.3) articulates in a formal way that the best model should be the best compromise between goodness-of-fit ($V(\hat{\boldsymbol{\theta}})$, decreasing with n_θ) and model variability ($p(n_d, n_\theta)$, increasing with n_θ). Or, recalling the previous chapter, the penalty factor ensures that the selected model is a posteriori identifiable.

The criteria of the form of (6.3) considered in this chapter are listed in Table 6.1. AIC is usually referred to as being the abbreviation of “Akaike information criterion”, although its original author Hirotugu Akaike explains it differently: “*IC stands for information criterion and*

A is added so that similar statistics, BIC, DIC etc., may follow [Akaike, 1974].

MDL is based on the minimum description length principle, from which it received its name. Grünwald [2000] explains the fundamental idea behind this principle: “any regularity in the data can be used to compress the data, i.e. to describe it using less symbols than the number of symbols needed to describe the data literally.” The MDL criterion was derived by Jorma Rissanen in 1978 [Rissanen, 1978] in the framework of coding theory. In the same year, Gideon Schwarz published an (asymptotically) equivalent model selection criterion [Schwarz, 1978]. This criterion is often referred to as Bayesian information criterion, BIC, because it follows from a Bayesian approach, selecting the model associated with the highest a posteriori probability. Historically, MDL is more commonly used in applied and computer sciences, whereas statistical circles usually refer to BIC.

Initially, the AIC and MDL rules were presented for cases where no noise information is available [Akaike, 1974; Rissanen, 1978; Schwarz, 1978]. Although the actual aims of the methods are different [Forster, 2001; Kieseppä, 2003], they were all constructed with the goal to perform well in the limit of large samples ($n_d \rightarrow \infty$). Later, these criteria were adapted for situations where noise variances are known [Schoukens et al., 2002], which are the AIC and MDL criteria used here (Table 6.1). In this case, $V_{LS}(\hat{\theta})$ is replaced by $V_{WLS}(\hat{\theta})$, the minimal value of the Weighted Least Squares (WLS) cost function, defined by Eq. (2.8).

De Ridder et al. [2005] have further modified these model selection criteria in order to improve their performance for short data records ($n_d \rightarrow n_\theta$). Since this is also our concern, we will consider these criteria as well, and will compare their performance to that of the cost function approach presented in section 6.3 (see AIC_S and MDL_S in Table 6.1).

Table 6.1. AIC and MDL criteria considered in this chapter. All are designed for the situation that the noise variances are known or priorly estimated.

Name	Expression	Applicability	Reference
	$V_{WLS}(\hat{\boldsymbol{\theta}}) \cdot p(n_d, n_\theta)$		
AIC	$\frac{V_{WLS}(\hat{\boldsymbol{\theta}})}{n_d} \cdot \left(1 + \frac{2n_\theta}{n_d}\right)$	$n_d \rightarrow \infty$	<i>Schoukens et al.</i> [2002]
MDL	$\frac{V_{WLS}(\hat{\boldsymbol{\theta}})}{n_d} \cdot \left(1 + \frac{n_\theta \cdot \ln(n_d)}{n_d}\right)$	$n_d \rightarrow \infty$	<i>Schoukens et al.</i> [2002]
AIC _S	$\frac{V_{WLS}(\hat{\boldsymbol{\theta}})}{n_d - n_\theta} \cdot \left(1 + \frac{2n_\theta}{n_d - n_\theta}\right)$	$n_d \rightarrow n_\theta$	<i>De Ridder et al.</i> [2005]
MDL _S	$\frac{V_{WLS}(\hat{\boldsymbol{\theta}})}{n_d - n_\theta} \cdot \left(1 + \frac{n_\theta \cdot \ln(n_d)}{n_d - n_\theta}\right)$	$n_d \rightarrow n_\theta$	<i>De Ridder et al.</i> [2005]

Note that AIC and MDL serve to *rank* models, such that the “best” model simply is the one with the lowest AIC or MDL value. However, this ranking merely states something about the relative quality of a number of models. Nothing is implied regarding the absolute suitability of any model, nor is a given difference in AIC or MDL value directly related to a difference in probability of being the best model. To address this last aspect, *Wagenmakers and Farrell* [2004] demonstrated how AIC values can be transformed in so-called Akaike weights which can be directly interpreted as probabilities for each model.

6.2.2. Other model selection criteria

In the statistical literature, a number of alternative criteria exist which can be used in the same way as AIC and MDL, i.e. as a ranking criterion for alternative models. Without going into detail, it may be of interest to briefly present some of them.

A criterion that was originally proposed to select which variables should be included in a linear regression is called C_p [Mallows, 1973]. The best variable subset should minimise

$$C_p = \frac{1}{\hat{\sigma}^2} V_{LS}(\hat{\boldsymbol{\theta}}) - n_d + 2n_\theta. \quad (6.4)$$

In this expression $\hat{\sigma}^2$ is an estimate of the experimental error σ^2 , assumed constant, such that the C_p 's expected value is n_θ .

The root mean squared deviation (RMSD) is yet another corrected cost function criterion

$$\text{RMSD} = \sqrt{\frac{V(\hat{\boldsymbol{\theta}})}{n_d - n_\theta}}. \quad (6.5)$$

Although the RMSD has been repeatedly used a heuristic criterion to select among mathematical models [e.g. Friedman *et al.*, 1995], “no statistical justification exists as to what the criterion attempts to estimate” [Myung, 2000].

Hamparsum Bozdogan recently developed a new measure of complexity, called ICOMP (Information COMPLexity) [e.g. Bozdogan and Haughton, 1998]. It is given by

$$\text{ICOMP} = -2\log L(\hat{\boldsymbol{\theta}}) + 2C(\text{Cov}(\hat{\boldsymbol{\theta}})), \quad (6.6)$$

with $L(\hat{\boldsymbol{\theta}})$ the likelihood function associated with the parameters $\hat{\boldsymbol{\theta}}$, C is a measure for complexity. As for AIC and MDL, ICOMP includes a second term which penalises for the model complexity. However, instead of penalizing the number of free parameters directly, ICOMP estimates the model complexity from the covariance of the parameter vector. This setup appears to be especially suited for Principal Components Regression (PCR) models, because the covariance matrix of the regression vector is well defined for these models [e.g. Capron *et al.*, 2005].

6.2.3. Cross-validation

A different strategy for selecting the most appropriate model more directly assesses the prediction capacity of the models. These methods are based on the idea of cross-validation [Stone, 1974], which consists of calibrating the models on a part of the data and next evaluating their (predictive) performance on the remaining part of the observations. The primary advantage of cross-validation is that it does not involve distributional assumptions about the data [Browne, 2000]. Different resampling strategies and performance criteria have been discussed in literature [Shao, 1993; Browne, 2000; Li *et al.*, 2002]. Robust alternatives have been proposed, considering that in the unaccounted presence of outliers overly complex models are selected [e.g. Wisnowski *et al.*, 2003]. It has been shown that cross-validation is under some specific assumptions asymptotically equivalent to AIC [Stone, 1977]. One major disadvantage of these methods is that they need a considerable amount of data. Therefore, we will not further consider this kind of model selection methods.

6.2.4. Note on the choice of the model set

The relevance of the model selection decision depends on the representativeness of the considered model set. Ideally, all possible models should be tested, their complexity only being constrained by the available number of data. Obviously, this rapidly becomes an unfeasible task and other procedures must be used. Usually, first a model *class* is chosen, on more or less subjective basis. Little has been written about how to objectively or systematically choose the models to consider [Chatfield, 1995].

Once a model class has been defined (typically forming nested models), all models of this class can be tested, but even this restricted model set can become too extensive. In this case, forward or backward procedures can be useful. These are stepwise procedures where in every step it is investigated which model parameter can best be added/removed [e.g. Massart *et al.*, 1997, p. 280; Kadane and Lazar, 2004] without starting to over/undermodel. The forward and backward procedures do not necessarily lead to the same optimal model, nor to same model that

would have been selected if all models had been considered. In particular when model parameters are correlated, an optimal model may be overlooked using these stepwise procedures.

6.3. Model selection based on the statistical interpretation of the minimal WLS cost function

6.3.1. No model errors

The proposed model selection strategy is based on the property of the minimum of the WLS cost function (2.8) to be a sample from a χ^2 distribution with $n_d - n_\theta$ degrees of freedom [Box, 1970; Rod and Hancil, 1980]:

$$V_{WLS}(\hat{\boldsymbol{\theta}}_{WLS}) \sim \chi_{n_d - n_\theta}^2. \quad (2.13)$$

The assumptions necessary for this property to be valid are listed below.

Assumptions 6.1 (*statistical interpretation of minimal WLS cost function*)

A. The measurements are disturbed by random zero mean normally distributed noise with known (or priorly estimated) variance, and are not subject to systematic errors. The noise information is accurately taken into account by the weighting matrix in the expression of the WLS cost function.

B. No model errors are present.

If assumption A is valid and assumption B is not, this will appear in the value of the residual cost function $V(\hat{\boldsymbol{\theta}})$, which will be significantly too large compared with the expected value of the appropriate $\chi_{n_d - n_\theta}^2$ distribution. Summarised, the analysis of the residual WLS cost function enables the detection of model errors.

In the framework of selecting the best model out of a given set, this assessment can be repeated for each of the models, eliminating those models that give too high a residual cost function. After that, a set of models (possibly only one or even none) remain, which are all validated not to “undermodel” the data (within some confidence level). From those models, the simplest one should be preferred.

The latter rule is nothing but the application of the “principle of parsimony” to model selection. If a simpler model fits the observations, why should parameters be added? Isaac Newton already stated as First Rule of Reasoning in *Principia*: “*We are to admit no more causes of natural things than such as are both true and sufficient to explain their appearances.*” And since the model selection procedure explicitly tests for undermodelling, we could say this quote from an unknown author describes the method quite well: “*Everything should be made as simple as possible but not simpler.*”

The cost function (CF) procedure shows some similarities with AIC and MDL strategies:

- (i) it can be run fully automatically, without interference from the user: it lets the data talk.
- (ii) It is supposed that the *global* minimum of the cost function has been reached.
- (iii) The same assumptions about the noise are made.
- (iv) Obviously, in order to deliver satisfactory results, at least one model in the model set should be a good approximation of the true mechanisms. (If this is not the case, the CF procedure will reject all models, and this aspect is therefore explicitly included in the method.)

The main “theoretical” differences are:

- (i) the CF method needs the subjective choice of the confidence limit, i.e. the $\chi^2_{n_d - n_\theta}$ percentile, to which the minimal cost function $V(\hat{\theta})$ has to be compared.

- (ii) To use the CF strategy, it is necessary to know the noise variances. Instead, expressions exist for AIC and MDL when this noise information is not available.
- (iii) In contrast with the classical AIC and MDL rules, the CF approach is not based on the assumption that the sample size is large.
- (iv) The CF procedure delivers not only the best model from the tested model set, it also provides the guarantee that the selected model is in itself appropriate to describe the data (within the desired level of confidence). Criteria like AIC and MDL only make statements about the models in relative terms (is a model better than another model?), and do not assess their absolute quality.

The CF method as just described will be tested and compared to the criteria listed in Table 6.1, by applying them to three different simulations (section 6.4) and to field data (section 6.5).

The performance of this CF model selection approach obviously depends on the validity of the two major assumptions (**A** and **B**, cf. Assumptions 6.1). Some remarks can be made regarding possible violations of these assumptions.

Violation of assumption A. The noise assumptions can be violated in several ways. If the noise distribution is known but not normal, the WLS cost function will not be $\chi_{n_d - n_\theta}^2$ distributed anymore. But in this case, the actual cost function distribution can still be determined empirically, i.e. by simulations. For the model selection application, not all properties of the distribution must be known, it is enough to identify the main percentiles. Another possible violation of the noise assumptions is that the variances are not well known. A strategy to estimate them in a preprocessing step is presented in Chapter 9, such that the presented CF method is still usable.

Violation of assumption B. If model errors are present, i.e. if none of the models tested can describe the data adequately, no model will be selected as “best model”. In that case, either the model set has to be extended as to include at least one “good model”, or another model

selection strategy has to be used. For instance methods using ranking criteria (cf. 6.2.1 – 6.2.2) can be used since they rank models according to their relative suitability and do not require that a *good* model is comprised in the tested set. However, the performance of e.g. AIC also depends on the assumption that the *true* model belongs to the study set. Another possible strategy, still based on the residual cost function value, is outlined below.

6.3.2. Model errors present

Although one always strives to achieve the perfect model, the presence of model errors cannot always be avoided. This results in residuals that contain a systematic component besides the stochastic noise. Hence, the residual cost function does not satisfy the assumptions anymore to enable the above statistical interpretation. However, this does not mean that any model selection is impossible. In real applications, it is often observed that when increasing the model complexity (in terms of number of model parameters), the cost function value declines slowly until it suddenly decreases considerably (see Figure 6.1). The gradual decrease symbolises how the model fit increases simply by better modelling the noise in the data. By contrast, the abrupt jump often indicates an optimal model complexity, since a small increase in complexity allows to describe the data considerably better. Sometimes, multiple of these jumps occur, signifying distinct structures in the data are subsequently captured by the model.

In practice, this intuitive approach remains rather subjective or at least qualitative when it comes down to determine which cost function decrease is really significant. It has been suggested to accept a cost function change as significant if it is larger than a certain percentage of the simplest model's cost function [Wisnowski *et al.*, 2003]. If the stochastic noise component is well known, the variance associated with the WLS cost function can be estimated [Pintelon *et al.*, 1997]. Comparing cost function changes with this variance allows to identify which decrease in cost function is significant for a given increase in model complexity.

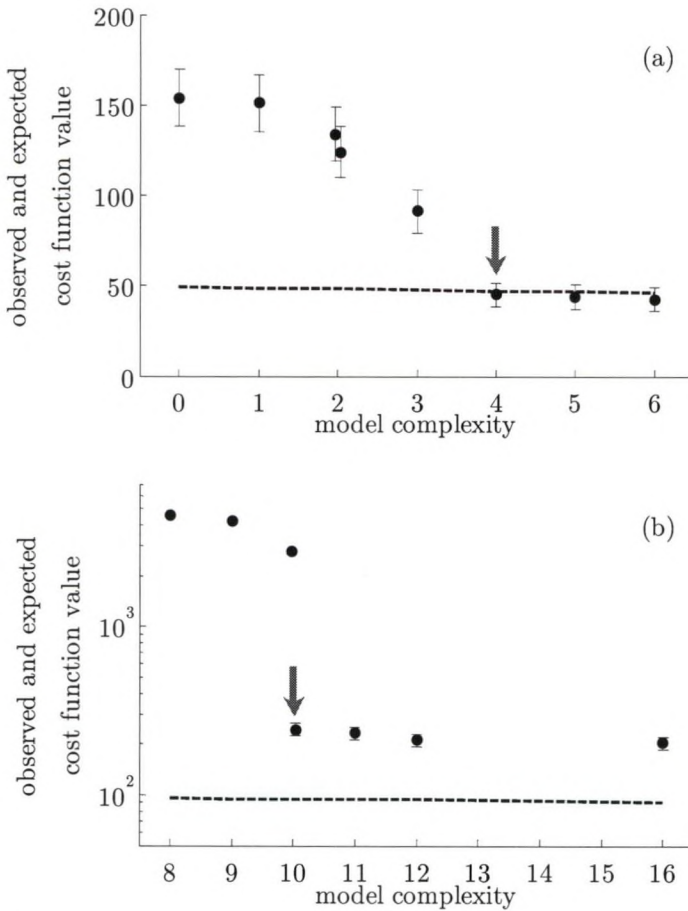


Figure 6.1. Two plots showing the evolution of the observed cost function values (black dots) with errorbars representing $\pm 1\sigma$, compared to the expected values (dotted line). The model to be selected is indicated by the arrow. (a) Situation without model errors, i.e. the true model with complexity 4 belongs to model set. The true model is selected either by comparison with the expected value or by identifying the significant jump. (b) Situation with model errors, i.e. the observed cost function never reach the expected value. In this case a comparison with the expected value is of little use, but the sudden jump indicates the optimal model. [after Pintelon et al., 1997]

6.4. Simulation tests

In order to test the performance of the proposed cost function (CF) approach and compare it with the other model selection strategies, we have applied them to three distinct simulation examples. Using simulations it is possible to check the results of the outcome (i.e. the selected model) since the “true” model is known, which is never the case for real measurements. On the other hand, a simulation can often be manipulated in order to display (or not) the weakness of a given method. To avoid this criticism, two of the simulations are performed on models previously used in literature about model selection: the FIR model taken from *De Ridder et al.* [2004] in 6.4.1 and the polynomial model from *de Luna and Skouras* [2003] in 6.4.3. In the third simulation test (6.4.2) a biogeochemical compartmental model (cf. Chapter 3) is considered, since this is our special interest, and it is used to model the field data in 6.5.. In all simulations, no model errors are present, i.e. the true model belongs to the tested model set.

6.4.1. FIR model

As a first case study, an electrical engineering example from *De Ridder et al.* [2005] is repeated, in order to have an objective comparison between the CF strategy and their AIC_s and MDL_s criteria. It concerns the identification of a discrete-time, finite impulse response (FIR) system (using their notation in the whole section)

$$G(z^{-1}, \boldsymbol{\theta}) = \sum_{r=0}^R g_r z^r \text{ with } \boldsymbol{\theta} = [g_0, g_1, \dots, g_R]^T \quad (6.7)$$

from noisy measurements of the input/output discrete Fourier transform spectra $U(k)$ and $Y(k)$. For a given order R , the model parameters are estimated by minimizing the maximum likelihood cost function

$$V(\boldsymbol{\theta}, Z) = \sum_{n=0}^{N-1} \frac{|Y(n) - G(z_n^{-1}, \boldsymbol{\theta})U(n)|^2}{\sigma_Y^2(n) + \sigma_U^2(n)|G(z_n^{-1}, \boldsymbol{\theta})|^2}, \quad (6.8)$$

with $\sigma_Y^2(n)$ and $\sigma_U^2(n)$ the known noise variances (the covariances were set to zero in the simulation). Expression (6.8) represents a WLS cost function as in expression (2.9), where the weighting matrix \mathbf{W} takes into account noise on both the input and output data. The true FIR model $G_\theta(z^{-1})$ used in the simulations is of order 19 ($R = 19$ in (6.7)); i.e. the true number of parameters is 20. The values of the parameters in θ are $[1.2101\text{e-}004, 1.2603\text{e-}003, 6.3357\text{e-}003, 2.0665\text{e-}002, 4.9573\text{e-}002, 9.3808\text{e-}002, 1.4610\text{e-}001, 1.9189\text{e-}001, 2.1454\text{e-}001, 2.0249\text{e-}001, 1.5490\text{e-}001, 8.2985\text{e-}002, 6.3557\text{e-}003, -5.4093\text{e-}002, -8.3676\text{e-}002, -7.8804\text{e-}002, -4.7431\text{e-}002, -5.2004\text{e-}003, 3.0909\text{e-}002, 4.8779\text{e-}002]^T$.

One thousand runs of a Monte-Carlo simulation were performed with $N = 80$. For each independent run, circular complex normally distributed errors were added to the true input and output spectra $U_\theta(k)$ and $Y_\theta(k)$, and the model parameters θ are estimated for model orders $R = 14, 15, \dots, 24$.

Table 6.2 shows how many times a model of order R has been selected by the different model selection methods. Columns one to four are similar to the results reported in Table 6.2 of *De Ridder et al.* [2005]; the slight differences are due to different random noise realizations in the two simulation tests. The classical AIC and MDL rules perform (much) worse than the modified AIC_s and MDL_s criteria. The first significantly overestimate the model order, the latter do so to a much lesser extent. Further, *De Ridder et al.* [2005] noted that it is expected that AIC selects higher orders than MDL, which is indeed observed here, both for the classical and the modified criteria. Their conclusion on this example is that MDL_s outperforms the other criteria.

On the other hand, the CF approach was tested for different significance levels (75%, 90%, 95% and 99%). It appears that each significance level picks out the true model approximately with a probability as expected (74.8%, 89.2%, 94.1% and 97.2, respectively). For that reason and to enable a comparison with the information criteria, we suggest using the 99th percentiles as confidence limits. Doing so, the CF approach performs even better than MDL_s (see Table 6.2):

- (i) the true model order was selected more often (972 times versus 911 times by MDL_s),
- (ii) the mean value of selected orders is a little closer to the true value, and
- (iii) the dispersion of selected model orders is smaller and more symmetric around the true order.

Table 6.2. Number of times a FIR model (6.7) of order n has been selected over 1000 simulations. The true model order = 19.

Model order	Model selection strategy							
	AIC	MDL	AIC _s	MDL _s	CF			
					75%	90%	95%	99%
14	0	0	0	0	0	0	0	0
15	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	1
18	0	0	0	0	0	1	2	17
<u>19</u>	<u>351</u>	<u>606</u>	<u>830</u>	<u>911</u>	<u>748</u>	<u>892</u>	<u>941</u>	<u>972</u>
20	104	123	86	59	14	12	7	3
21	103	74	43	21	7	9	3	1
22	99	67	24	5	9	3	7	0
23	138	60	10	2	8	7	5	0
24	205	70	7	2	2	3	1	0
Count	1000	1000	1000	1000	788	927	966	994
Mean value	21.1	20.1	19.3	19.1	19.1	19.1	19.1	19.0
Standard deviation	2.0	1.6	0.84	0.50	0.61	0.53	0.44	0.17

Finally, remark that 6 times no model order was selected at all by the CF approach, because none of the proposed ones provided an acceptable fit (at 99% significance level). For these noise realizations, the other criteria did choose a best model. They will always choose a model from a model set, even if all models are inappropriate, because they do not assess the absolute quality of the selected model (is it a good model?), only the relative superiority compared to the other models (is it better than the other models?).

The fact that sometimes none of the tested models is selected by the CF method, should not be seen as a disadvantage. In practice, this is rather an indication to the user that “something is wrong”, either due to

- (i) a spurious or systematic error in the measurements,
- (ii) an error in the inclusion of the measurement noise (i.e. assumption \mathbf{A} in section 6.2.2 is violated), or
- (iii) the model set does not contain the true model (or a good approximation).

By closer examination of the experiment-model design, it should generally be possible to determine which of these possibilities is the cause of the problem, and next to remediate the identified deficiencies.

From these results, we conclude that in this example the CF approach outperforms the other tested model selection criteria.

6.4.2. Compartmental model

The second example concerns a biogeochemical compartmental model used to estimate the exchange rates between different dissolved nitrogen pools, phytoplankton and the environment [Elskens *et al.*, 2002]. The schematic representation of the complete model is given in Figure 6.2. It consists of a special case of the general model presented in Chapter 3 with two dissolved nitrogen pools and seven potentially physically significant exchanges. They are estimated using measurements from time-consuming and expensive isotopic enrichment experiments, which deliver (only) seven observations. This means $n_d = n_{\theta, \max} = 7$, hence this




Figure 6.2. Schematic representation of a biogeochemical compartmental model used to describe the exchange between dissolved nitrogen nutrients (ammonium NH_4^+ and nitrate NO_3^-) and phytoplankton (P) in aquatic environments. The compartments are specified by two measurable quantities: total concentration and isotopic composition (atom% excess æ). The arrows represent the exchange rates, which are the parameters to be estimated. R 's stand for regeneration rates, U 's for uptake rates and L 's for loss rates. For more information, we refer to *Elskens et al.* [2002].

is a situation really on the edge of statistical significance and methods designed to have good asymptotic properties as the classical AIC and MDL rules, are doomed to fail. Since all measurements of such an experiment are subject to errors, the flux rates are estimated using a WLS cost function, adapted to take the input and output noise into account (similar to (6.7), cf. Chapter 4). Although the input noise is included via linearisation, the residual WLS cost function still is well approximated by a $\chi_{n_d - n_\theta}^2$ distributed variable (cf. section 4.3.2 and *de Brauwere et al.* [2005b]).

The question to be answered here is whether or not all the processes are necessary to account for the observations. Can any of the parameters be cancelled without losing the essence hidden in the measurements?

To answer this, we used a backward procedure of model selection. Optimising all possible parameter combinations (127) and then selecting the best model (as was done in the FIR model example) involves too long calculation times. Instead, an iterative *backward elimination procedure* is followed: starting from the complete model ($n_\theta = n_{\theta, \max}$), the complexity is decreased step by step until no simpler model is acceptable anymore:

1. $n_\theta = n_{\theta, \max}$ Optimise θ for the most complex (here even saturated) model. The system has no degrees of freedom, hence the value of $V(\hat{\theta})$ carries no statistically interpretable information (it is approximately zero regardless of true processes and data), but $\hat{\theta}$ can be useful as starting value in the next optimisations.
2. $n_\theta := n_\theta - 1$. Optimise all models of complexity n_θ . Keep only that model with the lowest $V(\hat{\theta})$. This is the best model of complexity n_θ .
3. Compare this $V(\hat{\theta})$ to the chosen $\chi^2_{n_d - n_\theta}$ percentile.
 - 3a. If the model fit is acceptable, go to 2.
 - 3b. If not, select the best model of complexity $n_\theta + 1$.

Using this procedure, the saturated model is selected if none of the models with one degree of freedom is acceptable. Although the risk exists to overlook the optimal model by this kind of partial procedure, the backward procedure has been reported to be more reliable than the forward alternative [Mantel, 1970].

For the AIC and MDL criteria, the procedure differs only in that the calculated AIC or MDL values are compared to those from the best model from the previous loop (lowest AIC or MDL value). The AIC_s and MDL_s criteria cannot assess the saturated model (zero degrees of freedom), for this would require a division by zero (see Table 6.1). Therefore, for them the backward procedure is adapted to start with the 6 parameter model with the smallest residual cost function. Finally, note that performing a cross-validation strategy (cf. 6.2.3) to select the best model is quite impossible in this example. There are too few data (7!) and they are not interchangeable, because they represent measurements

of different properties (concentrations and atom% excesses of different compounds, cf. Table 6.3).

A series of true models were simulated, with $n_\theta = 2, \dots, 7$. The simulation conditions are summarised in Table 6.3. For each model, 100 Monte-Carlo simulations were performed. The number of times each method selected the right model is shown in Table 6.4.

Table 6.3. Summary of the simulation conditions.

		<i>Simulated model (number of nonzero parameters)</i>					
		7	6	5	4	3	2
<i>Initial conditions</i> ($t=0$)	[NH ₄]				0.55 μ M		
	α_{NH_4}				8.9668 %		
	[NO ₃]				0.55 μ M		
	α_{NO_3}				8.9668 %		
	[P]				1 μ M		
	$\alpha_{\text{P,NH}_4 \text{ exp}}$				0		
	$\alpha_{\text{P,NO}_3 \text{ exp}}$				0		
	t_{inc}				24 h		
<i>Parameter values</i> ($\mu\text{M/h}$)	R_{NH_4}	0.01	0.01	0.01	0.01	0.01	-
	R_{NO_3}	0.008	0.008	0.008	0.008	-	-
	U_{NH_4}	0.02	0.02	0.02	0.02	0.02	0.02
	U_{NO_3}	0.01	0.01	0.01	0.01	0.01	0.01
	L_P	0.02	-	0.02	-	-	-
	L_{NH_4}	0.008	0.008	-	-	-	-
	L_{NO_3}	0.007	0.007	-	-	-	-

Table 6.4. Number of times the true compartmental model (i.e. only containing the parameters also present in the true model) has been selected over 100 simulations. Between brackets, the parameters that are equal to zero are given for each true model.

<i>True model</i> (number of nonzero parameters)	<i>Model selection strategy</i>							
	AIC	MDL	AIC _s	MDL _s	CF			
					75%	90%	95%	99%
7	100	100	imposs.	imposs.	89	83	70	49
6	0	0	77	79	65	80	70	55
5	0	0	31	30	67	86	94	99
4	0	0	12	11	64	81	93	95
3	0	0	5	14	73	90	94	96
2	0	0	4	4	77	91	95	99

The following observations and remarks can be made:

- (i) The classical AIC and MDL rules are absolutely unreliable in this example. This is not specified in the table, but they always select the saturated model with no degrees of freedom, obviously including the case where this was indeed the true model. Actually, this is a trivial result since a model with $n_\theta = n_d$ gives $V(\hat{\theta}) = 0$, and both AIC and MDL will also be zero (see Table 6.1). Hence, AIC and MDL are always minimized by the saturated model, independently of the true model.
- (ii) AIC_s and MDL_s also tend to select too complex models. In fact, for each of the true models simulated, AIC_s and MDL_s selected a 6 parameter model most often (detailed results not shown). So, even if they are designed to perform well in case of small sample sizes, in this border case they do not exhibit satisfactory results.

- (iii) The performance of CF has again been tested for different levels of confidence: 75%, 90%, 95% and 99%. When the saturated model is the true one, the true model is selected most often using the smaller significance levels. This can be explained remembering that the saturated model is only accepted if all simpler models are rejected. Consequently, using larger limits of acceptance (as for the 99% confidence level), it is more probable that a simpler model is accepted and hence the saturated (true) model is selected at less occasions. For models with more than one degree of freedom, the trend is opposite: with the 99% significance level, the true model is selected most often, just as for the FIR model (see 3.1.1). Moreover, it seems that as the degrees of freedom of the system increase, the different levels of significance tend to select the true model with their expected probabilities. Overall, using the 95th percentiles provides the highest probability of selecting the true model. This information will be useful when the CF model selection is applied to field data, as we will do in section 6.5.
- (iv) Due to fact that not all models are compared, the backward model selection procedure potentially selects a suboptimal model. Indeed, a more optimal parameter combination can remain undetected depending on the order in which parameters are removed. To address this caveat a combined forward-backward procedure could be used. Starting from the simplest model, at each step the cost function values for all parameters not yet in the model are checked and the parameter giving the best model fit is entered. After each step, the cost functions if any one parameter is removed from the model are tested. If a parameter is detected that does no longer significantly contribute to explain the data (i.e. the associated cost function is acceptable), it is rejected [after *Massart et al.*, 1997, p. 280]. The procedure is continued until no new parameter can be added without being removed afterwards. However, this more complicated scheme was not applied or tested because the backward procedure seems to provide satisfactory results.

This simulation example shows that even when the degrees of freedom approach zero, or even become zero, the CF approach can offer a

reliable way to select a model (provided that the level of confidence is chosen with care). This is in contrast with the AIC and MDL criteria, including the ones specifically adapted for small sample sizes.

6.5. Results on field data

The above simulation tests provided some confidence in the proposed CF model selection method. Now the final step to the real data is taken. Model selection is performed for a set of field experiments fitted to the compartmental model as described in section 6.4.2. The field data consist of 15 sets of nitrogen measurements made from May 1993 to April 1994 in the surface waters at several stations along the continental margin in the North-East Atlantic Ocean (*Elskens et al.* [1997] and personal communication).

Unlike in the simulation example, the true values of the parameters are unknown. However, the selection results can be interpreted to a certain extent using prior knowledge about the system and by intercomparing the results for the different experiments. Based on the simulation results of section 6.4.2 the model selection was performed using the CF approach with a 95% significance level. The optimised parameter values obtained before and after model selection (abbreviated as MS) are shown in Table 6.5. Obviously, the model selection step was not redundant since for all datasets three to five parameters (out of seven) were cancelled.

Several further remarks can be made:

- (i) Note that after model selection the standard deviation on the remaining parameters is mostly lower than before model selection. This is one of the rewards using model selection: by reducing the number of parameters, the uncertainty on the remaining parameters decreases.
- (ii) In section 6.4.2, we suggested to use a significance level of 95%, since on the whole it gave the highest probability to select the true model. In other words, what we could call type I errors (rejecting

the true model) are minimized. In order to check on type II errors (accepting a wrong model), we have compared the models selected for the field data (i.e. with two, three or four parameters) with the detailed results of simulations with a 95% significance level. This means we have verified how large the chance is (based on simulations) that the model complexities selected here were selected while actually the true model had another complexity. This probability appears to be small: at most 6%.

- (iii) The following trend can be noted: if before applying model selection, a given parameter value was not higher than its standard deviation, that parameter would be cancelled after model selection. This observed trend provides some confidence in the model selection results. Indeed, if by optimising the complete model, some parameters appear not to be significantly different from zero, we could expect that simply omitting them would equally well explain the data.
- (iv) The exceptions to this trend in Table 6.5 are underlined. Apparently, both kinds of error occur: (a) expected, but not cancelled by model selection (4 times, see underlined parameter values in Table 6.5) and (b) not expected, but cancelled by model selection (5 times, see x in Table 6.5).^(††) This means that simply replacing the model selection procedure by cancelling the parameters that appear to be approximately zero after total optimisation is not allowed. One reason for this is that cancelling all parameters that are not significantly different from zero at once, ignores “synergetic” effects between the different parameters: e.g. cancelling one may affect another parameter making it significant. *Elskens et al.* [2005] already made the more general observation that a parameter not being significantly different from zero does

^(††) Here, we have chosen to compare the parameter values with (one time) the estimated standard deviation. Comparing them to two or three times the standard deviation would be equally defensible. For the sake of curiosity, we have checked the results when the parameter values are compared to those limits. More differences compared to model selection results are noticed (resp. 10 and 11 exceptions, versus 9 when one time the standard deviation was used), yet they are all of the type “expected, but not cancelled by model selection”.

not necessarily imply that this parameter is not significant for modelling the data.

- (v) Another consistent trend through Table 6.5 is that the loss rates L_{NH_4} and L_P are systematically deleted after model selection. This also gives faith in the results, because in fact the loss rates do not have a univocal physical meaning but are rather fitting parameters representing all possible loss processes, used to constrain the model. If they are cancelled by model selection, it means that only considering the physically (biologically) meaningful parameters (regeneration and uptake rates) is sufficient, which is at least encouraging.
- (vi) It is striking that the only times that loss rate L_{NO_3} is not cancelled (cases are marked with an asterisk*), their optimised values are negative. Clearly, negative rates are impossible and these values point out that for these datasets something is wrong. The explanation was found with the experimenters. During the experimental work it was assumed that there was no nitrate regeneration (i.e. $R_{NO_3} = 0$), which enabled them to omit some measurements ($\varpi_{NO_3}(t)$ was set = $\varpi_{NO_3}(0)$). Indeed, R_{NO_3} is consistently left out by model selection. But the negative values for L_{NO_3} indicate that the assumption of no nitrate regeneration was not always valid, especially for the September measurements. In these cases, the only way to fit the measurements is to let parameter L_{NO_3} compensate by becoming negative.

These observations linked to knowledge about the system and experiment give some faith in the results, even though it is impossible to verify whether or how well the selected models correspond to the true mechanisms that generated the data.

Table 6.5. Results before and after applying the CF model selection (MS) to 15 field datasets, modelled by the compartmental model depicted in Figure 6.2. The values represent the estimated parameter values ± 1 standard deviation. The underlined results after model selection are exceptions to the “rule” that a parameter whose value is smaller than its standard deviation can be neglected. Those parameters that remain negative after model selection are marked with an asterisk*.

Exp. No	Sampling date		Estimated rates (parameters) in nM/h						
			R_{NH_4}	R_{NO_3}	U_{NH_4}	U_{NO_3}	L_{NH_4}	L_{NO_3}	L_P
1	03/05/93	Before MS:	8.0 \pm 6	0 \pm 10	10.3 \pm 3	23.0 \pm 4	-0.6 \pm 6	-6.4 \pm 10	8.3 \pm 10
		After MS:	7.1 \pm 3	×	8.9 \pm 3	20.4 \pm 3	×	×	×
2	24/09/93	Before MS:	3.8 \pm 4	0 \pm 6	5.7 \pm 2	2.2 \pm 0.3	1.9 \pm 3	-1.4 \pm 6	7.9 \pm 5
		After MS:	<u>3.8 \pm 2</u>	×	6.2 \pm 1	2.2 \pm 0.5	×	×	<u>×</u>
3	25/09/93	Before MS:	8.3 \pm 5	0 \pm 5	9.6 \pm 2	7.3 \pm 1	3.2 \pm 4	-4.4 \pm 5	4.5 \pm 6
		After MS:	7.9 \pm 2	×	11.4 \pm 3	7.7 \pm 1	×	-4.4 \pm 5*	×
4	26/09/93	Before MS:	9.8 \pm 4	0 \pm 5	11.0 \pm 2	7.2 \pm 1	2.2 \pm 4	-5.1 \pm 4	5.7 \pm 6
		After MS:	9.1 \pm 2	×	11.8 \pm 3	7.4 \pm 1	×	-5.4 \pm 4*	×
5	27/09/93	Before MS:	3.8 \pm 2	0 \pm 7	6.9 \pm 1	10.3 \pm 2	-1.4 \pm 2	-6.9 \pm 6	4.6 \pm 6
		After MS:	3.5 \pm 1	×	6.4 \pm 1	10.2 \pm 2	×	-6.9 \pm 6*	×
6	27/09/93	Before MS:	3.1 \pm 2	0 \pm 6	4.7 \pm 1	9.9 \pm 1	0.5 \pm 3	-5.8 \pm 6	2.1 \pm 8
		After MS:	3.2 \pm 2	×	4.9 \pm 1	10.0 \pm 2	×	-5.9 \pm 6*	×
7	28/09/93	Before MS:	2.9 \pm 3	0 \pm 8	3.8 \pm 1	5.6 \pm 1	1.6 \pm 3	1.0 \pm 8	5.3 \pm 6
		After MS:	<u>3.1 \pm 2</u>	×	4.2 \pm 1	6.0 \pm 1	×	×	×
8	21/04/94	Before MS:	52.2 \pm 16	0.2 \pm 160	8.6 \pm 1	10.9 \pm 2	7.2 \pm 17	0.3 \pm 164	37.7 \pm 53
		After MS:	46.6 \pm 15	×	8.7 \pm 1	11.3 \pm 3	×	×	×
9	22/04/94	Before MS:	42.0 \pm 12	0 \pm 120	13.4 \pm 2	14.2 \pm 2	1.5 \pm 13	0.1 \pm 122	-4.1 \pm 51
		After MS:	40.8 \pm 10	×	13.3 \pm 2	14.2 \pm 2	×	×	×
10	22/04/94	Before MS:	4.1 \pm 4	0.9 \pm 235	4.9 \pm 1	6.5 \pm 1	2.5 \pm 8	4.4 \pm 242	28.0 \pm 22
		After MS:	<u>×</u>	×	4.8 \pm 1	6.9 \pm 1	×	×	×
11	22/04/94	Before MS:	5.0 \pm 4	0.4 \pm 180	6.0 \pm 1	9.7 \pm 1	5.7 \pm 8	0.7 \pm 184	15.7 \pm 23
		After MS:	<u>×</u>	×	4.2 \pm 1	6.0 \pm 1	×	×	×
12	24/04/94	Before MS:	22.0 \pm 7	0.7 \pm 145	9.4 \pm 1	9.7 \pm 1	6.9 \pm 8	0.2 \pm 150	7.6 \pm 14
		After MS:	17.9 \pm 6	×	9.3 \pm 2	9.9 \pm 2	×	×	×
13	24/04/94	Before MS:	46.1 \pm 14	0.6 \pm 168	11.8 \pm 2	11.1 \pm 2	3.3 \pm 15	1.6 \pm 172	5.6 \pm 32
		After MS:	43.6 \pm 13	×	11.7 \pm 2	11.1 \pm 2	×	×	×
14	27/04/94	Before MS:	28.0 \pm 8	0 \pm 57	22.1 \pm 3	62.3 \pm 9	9.2 \pm 8	-7.9 \pm 55	6.6 \pm 56
		After MS:	25.6 \pm 5	×	24.4 \pm 4	61.5 \pm 9	<u>×</u>	×	×
15	27/04/94	Before MS:	15.9 \pm 5	0 \pm 63	11.5 \pm 2	52.0 \pm 7	6.6 \pm 6	-2.0 \pm 62	19.1 \pm 53
		After MS:	13.6 \pm 4	×	11.7 \pm 2	52.5 \pm 8	<u>×</u>	×	×

6.6. Conclusions

Combining a statistical interpretation of the minimum of a Weighted Least Squares cost function and the principle of parsimony, a method to select the “simplest good” model was proposed. Different simulation tests suggest that the CF method outperforms the AIC and MDL-like criteria, especially when the number of data is very close to the number of parameters to be estimated. We hope that these results may give some confidence in the usefulness of this model selection method.

We would like to emphasize that the choice of the significance level remains a delicate step, as for all hypothesis tests. From the simulation results, it becomes clear that it is not straightforward to propose one absolute best confidence level. Instead, we propose to interpret the results for different confidence limits in a more qualitative way. For instance, the classification as proposed by the Analytical Methods Committee [1992] could be used:

- if $V(\hat{\theta}) \leq \chi_{95\%}^2$, then the model is “satisfactory”,
- if $\chi_{95\%}^2 < V(\hat{\theta}) < \chi_{99\%}^2$, then the model is “questionable”,
- if $V(\hat{\theta}) \geq \chi_{99\%}^2$, then the model is “unsatisfactory”

(where $\chi_{95\%}^2$ and $\chi_{99\%}^2$ represent respectively the 95th and 99th percentiles of the χ^2 distribution with the appropriate ($n_d - n_\theta$) degrees of freedom).

This would allow an interpretation and thus a model selection with more nuance, since the distinction between acceptance and rejection of a model is less discrete.

Making the link to the environmental applications, it seems appropriate to note how relevant but unrecognised this issue of model selection is today, considering that in most studies of environmental modelling the trend is to make ever larger models.

“For if we know of something of potential relevance, and computational power is increasing, why should that something be left out? Those who use the results of such models are probably reassured by this imprematur, of having supposedly based their decisions on the best available scientific

evidence. Our models, and certainly those we would label 'state-of-the-art', seem destined always to get larger.” [Beck, 1999].

This effort is completely opposed to the whole rationale presented in this chapter. Admittedly, the purpose of these large mechanistic models is not the same as the black or grey box models considered here (cf. remark in section 2.4). Nevertheless, performance criteria such as identifiability and robustness of the model output are equally important for both applications, although they may have to be tackled in different ways. In a sense, the difference between the two approaches lies in the value attributed to the scientific experience or other subjective and qualitative prior knowledge. In the argumentation of some large models this criterion may be abused, but maybe it is too much ignored in black box modelling.

As a final remark, a question of more philosophical nature may be raised, putting the whole model selection subject in a broader perspective. How realistic is the assumption that one true model exists for the observed, especially natural, phenomena? Most will agree that even the best model is not reflecting all true processes, but is rather a simplification giving an adequate approximation of the data at hand. Indeed, we already argued that a model should primarily be useful in describing the current (and future) data. But in that case, there may be more than one model that is “useful”, i.e. a sufficiently close approximation of the observations for the purpose of the study [Chatfield, 1995]. In other words, some uncertainty is associated with the selection of a best model. This uncertainty is independent of the selection procedure (the model may even be chosen based on prior knowledge), it is merely the consequence of the lack of one true model. This uncertainty is rarely taken into account in the further identification steps. The one selected model is taken as a priori known and its parameters (and their uncertainty) are estimated ignoring a possible uncertainty regarding the model structure. Several authors argue that this model uncertainty is likely to be one of the more relevant sources of model prediction uncertainty [Chatfield, 1995; Brun *et al.*, 2001]. Obviously, this uncertainty is difficult to quantify and it would require other assumptions to find estimates. For instance, a

Bayesian model averaging approach has been suggested, combining a number of plausible models such that it is not necessary anymore to select a single best model [Chatfield, 1995; Draper, 1995; Reichert and Omlin, 1997]. But this approach needs the specification of prior probabilities for the various models, which is a tricky business. Furthermore, for many applications it is still important to have *one* model (e.g. if the physical interpretation of the parameters is of interest), and not a, possibly very large, set of models. Therefore, we will continue to use the classical scheme involving only one model, keeping in mind that no solution is perfect.

Part B:

Ocean mixing model

In Part A, the issues of parameter and uncertainty estimation, identifiability and model selection were introduced and solutions were suggested in the framework of tracer experiments modelled by compartmental models. In Part B, some of these issues will return, new ones will be launched, but all in the context of a different application. The common application in all following chapters is the modelling of water mass mixing in the ocean.

Oceanic circulation and water mass mixing processes can be studied via complex physical models [e.g. *Maltrud et al.*, 1998]. Such approaches have the advantage to be highly interpretable since every aspect can be traced back to fundamental physical laws. However, due to their dimension validation of these models using extensive datasets is not straightforward. On the other hand, a conceptually more simple, but completely data-based approach is often used: Optimum Multiparameter (OMP) analysis [*Tomczak*, 1981; *Thompson and Edwards*, 1981; *Mackas et al.*, 1987; *Tomczak and Large*, 1989]. This type of model reconstructing distributions of mixing water masses will be the recurring application of Part B. The main aspects of OMP analysis will be introduced first (Chapter 7), followed by an extensive description, argumentation, validation and application of an improvement we propose (Chapter 8). This improvement cannot solve all problems, so in the next chapters two additional developments are investigated. In Chapter 9 a refined method to estimate the noise variances of the measurements is derived, in an attempt to enable a better quality check of the estimation results. Finally, identifiability will be further investigated in Chapter 10.

It is probably useful to allocate an additional paragraph to the motivation of this work. Indeed, modelling the ocean dynamics does not all fall within the traditional focus of the ANCH (Analytical and Environmental Chemistry) department and the reconstruction of the mixing fields wasn't really the result of interest. The actual incentive of the work presented in this part was to enable the detection of non-mixing processes involving dissolved barium. This question is relevant for Stéphanie Jacquet, whose PhD was concerned with a better understanding of barium cycling in the Southern Ocean. By accurately reconstructing mixing fields, we hoped that those regions in the ocean could be detected where barium concentrations are not behaving as expected if they were only influenced by mixing processes. The barium question will not further be raised in this thesis part, but knowing the reason why we wanted to go through this trouble may put some things into perspective. For the interested reader, a section in the chapter on future perspectives is entirely devoted to a more detailed discussion of this barium issue.

Chapter 7:

Optimum MultiParameter analysis

“Some notes on the history of water mass analysis. (...) The period before the First World War was dominated by the collection of primary data. Between the wars it was followed by a period of systematic study of ocean basins characterized by ocean-wide expeditions (...). The study of water masses formed the major tool for the discovery and description of core aspects of ocean dynamics during this period. (...) The development of the autonomous recording current meter and of techniques for long term deployment of oceanographic moorings shifted the focus of research toward dynamic phenomena (...). The discovery of strong time-variable ocean currents, of meso-scale eddies, fronts and internal waves dominated the first decades after the Second World War. (...) Classical water mass analysis was ill equipped to contribute to this phase of oceanography. (...) The last two decades have witnessed the reappearance of water mass analysis as a respected branch of oceanographic research (...). This new phase of oceanography is characterized by our growing awareness of the ocean’s role in climate variability and climate change. (...) [T]he study of water masses has become an important component of international programs such as WOCE [World Ocean Circulation Experiment] and CLIVAR [Climate Variability] which have produced and will continue to produce a large amount of high quality hydrographic data (...). The range of methods useful for water mass analysis has also grown notably over the last few decades. Numerical modeling (...). Inverse methods (...). These examples show that water mass analysis is on the threshold of a new era.”

Tomczak [1999]

7.1. Introduction

“Water mass visualization will of course always rely on temperature-salinity and other property-property diagrams. Such graphical methods are invaluable as orientation tools, but their use in quantitative analysis is very limited. (...) The products most useful for climate analysis will eventually be maps of water mass content (...), water mass age and water mass formation rate” [Tomczak, 1999].

Optimum MultiParameter (OMP) analysis is a data-based tool offering insights into especially the first of these aspects. If time can be introduced in the analysis, also the two other questions are reachable [Karstensen and Tomczak, 1998], but this will not be considered in this thesis.

In this chapter an introduction to OMP analysis is given, emphasizing its valuable aspects, but also pointing out some of its imperfections, some of which will be addressed in the next chapters.

7.2. Method

Classical OMP analysis usually considers a water parcel sampled on a grid (e.g. latitude versus depth) with a total of N grid points. At every sampling position of the grid, n_v physical and/or chemical properties are measured, such as temperature, salinity, dissolved oxygen concentration, etc. These observed fields are considered to be the result of mixing of a number (n_s) of source water masses, whose characteristics are assumed to be known. Following the definition given in Tomczak [1999], a set of properties corresponding to such source water mass is called a *source water type* (SWT). Defining the SWTs is a delicate job and the final outcome of the analysis highly depends on the choices made here. Once the SWTs are characterized, each measured property can be written as a linear combination of the SWTs' values for that hydrographic property. For example, potential temperature measured at location k can be modelled by

$$\Theta_{obs}(k) = \sum_{i=1}^{n_S} x_i(k) \Theta_i + e_{\Theta}(k), \quad (7.1)$$

where x_i is the i^{th} SWT contribution or fraction and Θ_i its potential temperature, and $e_{\Theta}(k)$ the disturbing noise. Only conservative hydrographic properties are considered in the analysis and it is assumed that only linear mixing occurs such that for each property a similar equation can be written. If enough properties are measured, the unknown fractions x_i at position k (the same in each equation) can be estimated by simultaneously optimising all mixing equations of form (7.1). In practice this is done taking into account two constraints. First, the mass balance equation

$$\sum_{i=1}^{n_S} x_i(k) = 1 \quad (7.2)$$

should be satisfied at any position k , and therefore this equation is optimised simultaneously with the mixing equations. In addition, usually a nonnegativity constraint is imposed on the mixing fractions, for obvious reasons of physical relevance. The procedure can be repeated for all grid points, resulting in optimal values for the mixing fractions at each sampling location. These are the basic ideas of the OMP approach.

Assumptions 7.1 (*OMP analysis*)

- (i) Only linear mixing processes occur, i.e. all water mass properties undergo the same mixing processes.
- (ii) Only conservative water mass properties are considered, i.e. their values are not influenced by biogeochemical processes.
- (iii) All relevant source water masses are included in the analysis, such that at any position $\sum_{i=1}^{n_S} x_i = 1$ and $x_i \geq 0$ ($\forall i$).
- (iv) The source water types are accurately known (but can be associated with a standard deviation).

This method has first been presented more than two decades ago by Tomczak [1981]. Several improvements have been proposed since then. The most objective way to assign a weight to each variable in order to reflect differences in reliability of each variable has been discussed [Thompson and Edwards, 1981; Mackas et al., 1987; Tomczak and Large, 1989]. Later, multivariate methods, like cluster analysis and principal components analysis, have been applied in an attempt to more objectively determine the source characteristics (i.e. the Θ_i in Eq. (7.1)) [You and Tomczak, 1993; de Boer and van Aken, 1995; You, 1997; Fogelqvist et al., 2003]. Although these methods are efficient in reconstructing multivariate datasets, they lack the physical interpretability that OMP analysis has. In particular, the “SWTs” are mere mathematical constructs representing the directions in multivariate space with maximal variation, but their characteristics generally do not correspond to real source water masses. Therefore, these methods are not further considered here. The OMP methodology has been widely used in oceanographic literature to investigate source water contributions in various regions and systems (e.g. You, 1997; Castro et al., 1998; Coatanoan et al., 1999; Perez et al., 2001; You et al., 2003; Lo Monaco et al., 2005).

This is probably due to its main advantage: its simplicity. It is indeed a very intuitive approach, not requiring many assumptions about the physics of the investigated system. The outcome of the computations (i.e. the fraction fields) is directly visualisable and interpretable. This does not mean that the analysis input should not be chosen carefully, but rather that the quality of the results can often be checked even visually. Yet, this OMP method exhibits a number of imperfections, especially from a statistical point of view. These are further discussed in the next section.

7.3. Imperfections

A. No position dependence. Since every grid point is optimised individually and independently, the position dependent variation of the source contributions over the grid is not taken into account. In other

words, it is not considered that if a source water contribution is high in one point it will probably be quite high in a neighbouring point too. A consequence of the current working scheme is that the mixing fractions for a given point only depend on the data associated with this point, and hence are highly influenced by errors in that point.

B. Few or no residual degrees of freedom. To estimate all unknown fractions x_1, \dots, x_{n_s} from the set of equations (mixing equations of form (7.1) and the mass balance equation (7.2)), the number of unknown fractions, n_s , may not exceed the number of measured variables (n_v) + 1. Otherwise the problem is degenerated, which means that there is an infinite number of solutions for x_1, \dots, x_{n_s} which all fit the data equally well. In the extreme situation where $n_s = n_v + 1$, there is exactly one mathematically correct solution for the unknown fractions, but it is of little statistical significance because there are no residual degrees of freedom. This situation can be compared to a linear regression (two unknown parameters) through two data points (two equations): a straight line can perfectly be fitted through two points, but the noise contained in the data is entirely modelled as significant information. Consequently, the fitted line will certainly not equal the *true* line through the *true* data values. To summarise, to be statistically relevant the OMP method should be performed in overdetermination, i.e. with $n_s \leq n_v$, and preferably even $n_s \ll n_v$. However, due to the reality that only a small number of conservative properties are usually available, it is common practice to apply OMP analysis to the extreme situation with $n_s = n_v + 1$. The main consequence is that the resulting mixing fraction values are highly influenced by stochastic errors in the data and thus associated with a higher uncertainty.

C. Mass balance is not strictly imposed and inconsistent weighting. In the OMP procedure mass balance (Eq. (7.2)) is treated as if it were a mixing equation (cf. Eq. (7.1)). Conceptually this is not entirely correct since the mixing equations are actually modelling a measurement, whereas the mass balance equation only states a physical constraint that should (exactly) be obeyed. As a consequence of this handling, the fractions optimised by OMP analysis only approximately obey mass balance. An additional difficulty with this approach arises

when weights have to be chosen for each equation. Indeed, before optimising the mixing equations, each equation is multiplied by a weight, usually expressing the reliability of the measurement described in the equation. This is a healthy thing to do since it is known that some variables are measured with much higher precision than others. The values of the weights highly influence the final fractions so they should be chosen with care. However, as mentioned just before, it is not possible to attribute such a weighting factor to the mass balance equation, because there is no measurement involved in this equation. In practice, *“the largest of the weights (...) is allocated to the mass conservation equation”* [You and Tomczak, 1993]. This dilemma was already noted by Tomczak and Large [1989]: *“More thought has to be given to the most suitable weights. How important in this context is conservation of mass?”*.

D. Subjectivity of source water type definition. Maybe the most important assumption made in the whole analysis is the choice of the SWTs. A few attempts have been made to make the SWT definition more objective, in particular by using multivariate methods such as Cluster Analysis and Principal Components Analysis [You and Tomczak, 1993; de Boer and van Aken, 1995; Fogelqvist et al., 2003], but these methods seem unable to provide useful definitions in general. Consequently, the selection and definition of SWTs is typically said to be based on data inspection, combined with literature values and prior knowledge about the local oceanographical features and processes, in other words the SWTs are defined in a subjective way.

7.4. Concluding remarks

In this chapter the main aspects of classical OMP analysis were discussed. The emphasis was on the mathematical side of the story, culminating with a listing of technical imperfections. The awareness of these imperfections was the driving force behind the developments in the remaining chapters.

However, this may be a suitable moment to look further than the maths and stats.

“The knowledge of exact mixing fractions of water masses in the ocean is necessary for various applications, especially in the analysis of transient tracer fields or biogeochemical cycling. The distribution of tracers is controlled by a combination of transport processes associated with the oceanic circulation and mixing and by reactive processes associated with the major biogeochemical cycles (...). To evaluate the distribution of nutrients and tracers in the sea one has to resolve effects of mixing and of biogeochemical cycling. OMP analysis is a tool to analyse the water mass mixture in a water sample by calculating the contributions from the original water masses (so called source water masses) to the sample.” [Karstensen and Tomczak, 1999].

Indeed, except for potential temperature and salinity, most available hydrographic parameters are not strictly conservative, but possibly involved in biogeochemical processes. Although these processes impede the straightforward application of OMP analysis, these are exactly the processes of interest in many studies, eventually even this one (see section 12.5). If it is possible to extract the contribution of mixing to the observed profiles of tracers in the ocean, knowledge of the importance of biogeochemistry comes within reach.

The problem of including possibly nonconservative variables, such as oxygen or nutrient concentrations, in an OMP analysis has been tackled in several ways. One way, is to convert the nutrients into “preformed” nutrients, like NO and PO, before the analysis. These are corrected nutrient concentrations, considering that nutrients and oxygen are utilised and formed following a predefined Redfield ratio: $\text{NO} = r_{\text{O/N}}\text{NO}_3 + \text{O}_2$ and $\text{PO} = r_{\text{O/P}}\text{PO}_4 + \text{O}_2$. Another strategy, also based on the Redfield ratios, is to add a “biogeochemistry term” in the mixing equations, introducing an additional unknown ΔP to be estimated:

$$\left\{ \begin{array}{l}
 \Theta_{obs} = \sum_{i=1}^{n_S} x_i \Theta_i + e_{\Theta} \\
 S_{obs} = \sum_{i=1}^{n_S} x_i S_i + e_S \\
 \vdots \\
 O_{2,obs} = \sum_{i=1}^{n_S} x_i O_{2,i} - r_{O/P} \Delta P + e_{O_2} \\
 PO_{4,obs} = \sum_{i=1}^{n_S} x_i PO_{4,i} + \Delta P + e_{PO_4} \\
 NO_{3,obs} = \sum_{i=1}^{n_S} x_i NO_{3,i} + r_{N/P} \Delta P + e_{NO_3} \\
 1 = \sum_{i=1}^{n_S} x_i + e_{mass}
 \end{array} \right. \quad (7.3)$$

This procedure is called “*extended OMP analysis*” and allows to simultaneously quantify the changes due to biogeochemical processes. However, a limitation of both strategies is that one must be sure that the changes occur according to fixed Redfield ratios and one must know the values of these ratios. Even for well-known nutrients like nitrate and phosphate, this is still a topic of active debate. As a result, this procedure can certainly not be used for other tracers whose biogeochemical behaviour is not well understood. The incentive of this second thesis part was to assess the importance of biogeochemical processes in controlling the observed fields of dissolved barium. How we used OMP analysis to address this question is discussed in section 12.5. But the first step seemed to tackle the imperfections listed in the previous section, in order to have more reliable results. The technical developments needed for this goal appeared to be important enough by themselves to form the chapters of this second thesis part.

Chapter 8:

Parametric Optimum MultiParameter analysis

8.1. Introduction

In the previous chapter, the main ideas, advantages and imperfections of the classical OMP analysis have been summarised. It is the aim of this chapter to present a modified OMP analysis which addresses some of the imperfections while keeping the advantages. The method we propose is still based on the same mixing equations, but includes a parameterisation of the mixing fractions as a function of position allowing a more accurate tuning of the model complexity. First, the method is presented in detail, including some technicalities and a schematic algorithm (section 8.2), followed by some simulation tests intended to validate it (section 8.3). In section 8.4 a oceanographical dataset from the Southern Ocean is analysed using the new method, demonstrating its performance on real data, but also indicating the need for better estimates of the measurement uncertainties. This chapter is largely based on *de Brauwere et al.* [2007a]

8.2. Method

8.2.1. Parameterisation to increase precision

If n_v hydrographic properties are considered in OMP analysis, a mixing equation like Eq. (7.1) can be written for each variable at position k . This set of mixing equations for position k can formally be combined in one matrix equation

$$\mathbf{y}_k = \mathbf{S} \cdot \mathbf{x}_k + \mathbf{e}_k, \quad (8.1)$$

with \mathbf{y}_k the $(n_v \times 1)$ measurement vector at position k , \mathbf{S} is the $(n_v \times n_s)$ SWT definition matrix, which is independent of position, \mathbf{x}_k is the $(n_s \times 1)$ fractions vector for position k and \mathbf{e}_k stands for the error vector at position k . If this is repeated for all (N) positions of the grid where the hydrographic properties are measured, one can gather all equations in a similar matrix equation

$$\mathbf{Y} = \mathbf{S} \cdot \mathbf{X} + \mathbf{E}, \quad (8.2)$$

where \mathbf{Y} is the ($n_t \times N$) measurement matrix, \mathbf{S} is still the SWT definition matrix, \mathbf{X} is the ($n_s \times N$) fractions matrix and \mathbf{E} stands for the noise matrix. At this point \mathbf{X} contains the unknown fractions to be estimated and there is still one set of fractions for every grid point as mentioned in imperfection A. To make the model more robust, the discrete fraction field of each SWT (i.e. each row of \mathbf{X}) will from now on be described by a continuous function. We consider the situation where the samples lay on a two-dimensional (e.g. latitude-depth) grid, hence two-dimensional functions are used. So, instead of having values for the fractions of a given SWT on a discrete grid of points, we now consider a two-dimensional fraction *function*, which is function of two position coordinates (e.g. latitude l and depth z). We propose to parameterise the fractions of SWT i by a number (n_B) of two-dimensional spline basis functions B [Dierckx, 1995]. Formally, this means that the mixing contribution of SWT i , x_i , at position (l, z) is modelled by a linear combination of n_B basis functions evaluated in (l, z)

$$x_i(l, z) = \sum_{b=1}^{n_B} c_{i,b} B_b(l, z). \quad (8.3)$$

In practice, the two-dimensional functions B are formed by combinations of two one-dimensional B-spline functions [Vanlanduit, 2001]

$$B_b(l, z) = N_{i,k_l}(l) M_{j,k_z}(z), \text{ with } i \in \{1, \dots, n_l\} \text{ and } j \in \{1, \dots, n_z\}. \quad (8.4)$$

$N_{i,k_l}(l)$ is the i^{th} (out of n_l) B-spline, of order k_l , in the latitude direction evaluated in l . Similarly, $M_{j,k_z}(z)$ is the j^{th} (out of n_z) B-spline, of order k_z , in the depth direction evaluated in z . These B-splines are defined by knots (fixing their number) and by an order. $n_l n_z$ is the total number of basis functions used to parameterise the mixing fractions, i.e. $n_l n_z = n_B$ (see Eq. (8.3)). Between two consecutive knots the B-splines $N_{i,k_l}(l)$ and $M_{j,k_z}(z)$ are polynomials of degree k_l and k_z respectively. We use the basic configuration where the knots defining the transition between B-splines are ordered equidistantly over the study area, but in principle their position can be optimised to suit specific features of the

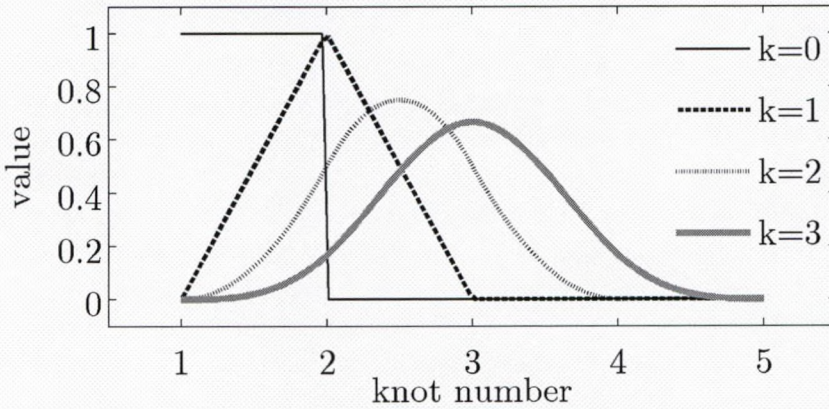


Figure 8.1. Shapes of B-splines of varying order k , illustrating that for one B-spline $k + 2$ knots are necessary.

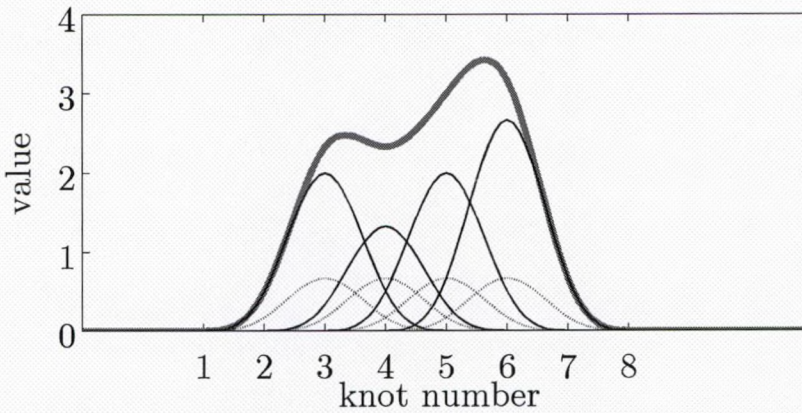


Figure 8.2. Construction of a spline as a weighted sum of B-splines, as a one-dimensional analogue of the procedure proposed here (Eq. (8.3)). Dashed lines represent B-splines of order 3 on an arbitrary knots grid; thin black lines are the individual B-splines multiplied with a given factor (resp. 3, 2, 3 and 4), forming the bold grey line when summed.

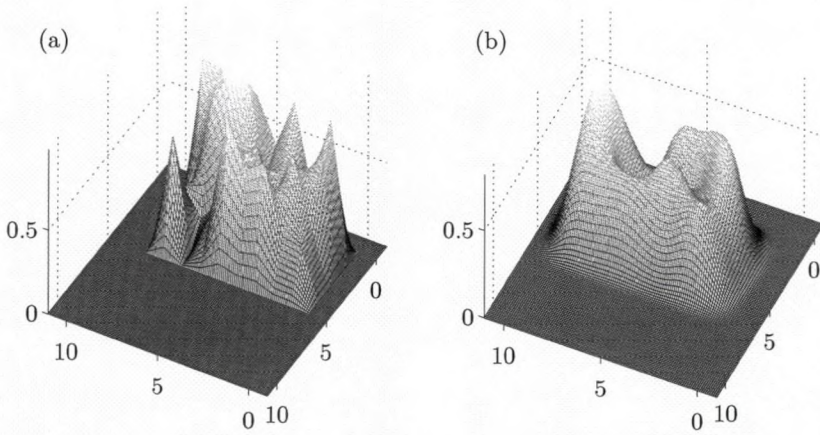


Figure 8.3. Examples of two-dimensional splines, formed by multiplying and summing B-splines of orders $k_l = k_z = 1$ (a) and 3 (b). Surfaces shown in (a) and (b) are arbitrary and independent.

investigated system. Between two consecutive knots the B-splines $N_{i,k_l}(l)$ and $M_{j,k_z}(z)$ are polynomials of degree k_l and k_z respectively.

More details about the choice of appropriate orders (k_l and k_z) and numbers of the B-splines (n_l and n_z) are given in the next section.

Merged in matrices, Eq. (8.3) becomes

$$\mathbf{X} = \mathbf{C} \cdot \mathbf{B}^T. \quad (8.5)$$

Combining this with Eq. (8.2), we find the central equation in this chapter:

$$\mathbf{Y} = \mathbf{S} \cdot \mathbf{C} \cdot \mathbf{B}^T + \mathbf{E}, \quad (8.6)$$

with

$$\mathbf{Y} = \begin{pmatrix} y_1(l_1, z_1) & \cdots & y_1(l_N, z_N) \\ \vdots & \ddots & \vdots \\ y_{n_y}(l_1, z_1) & \cdots & y_{n_y}(l_N, z_N) \end{pmatrix} \quad (8.7)$$

the $(n_y \times N)$ matrix containing the measurements of all variables (y_1 may be potential temperature, y_2 salinity, etc.) at all sampling positions;

$$\mathbf{S} = \begin{pmatrix} y_{1,swt\ 1} & \cdots & y_{1,swt\ n_S} \\ \vdots & \ddots & \vdots \\ y_{n_y,swt\ 1} & \cdots & y_{n_y,swt\ n_S} \end{pmatrix} \quad (8.8)$$

the $(n_y \times n_S)$ SWT matrix holding the characteristics of a source in every column; \mathbf{C} represents the $(n_S \times n_B)$ spline coefficients matrix

$$\mathbf{C} = \begin{pmatrix} c_{1,swt\ 1} & \cdots & c_{n_B,swt\ 1} \\ \vdots & \ddots & \vdots \\ c_{1,swt\ n_S} & \cdots & c_{n_B,swt\ n_S} \end{pmatrix} \quad (8.9)$$

and

$$\mathbf{B} = \begin{pmatrix} B_1(l_1, z_1) & \cdots & B_{n_B}(l_1, z_1) \\ \vdots & \ddots & \vdots \\ B_1(l_N, z_N) & \cdots & B_{n_B}(l_N, z_N) \end{pmatrix} = \begin{pmatrix} N_1(l_1)M_1(z_1) & \cdots & N_{n_l}(l_1)M_{n_z}(z_1) \\ \vdots & \ddots & \vdots \\ N_1(l_N)M_1(z_N) & \cdots & N_{n_l}(l_N)M_{n_z}(z_N) \end{pmatrix} \quad (8.10)$$

is the $(N \times n_B)$ basis function matrix (the same for every source). Now, the unknown parameters to be estimated are the spline coefficients contained in \mathbf{C} . More details about how to optimize these coefficients is given in sections 8.2.3 – 8.2.5.

The approach essentially consists of a parameterization of the original OMP equations. Therefore we propose to call it **Parametric OMP analysis** or **POMP**. A number of remarks can be made.

- (i) By parameterising the source fractions, they are made explicitly function of position (cf. Eq. (8.3)), thus addressing imperfection A. In other words, truly a continuous fraction field is estimated, and not a discrete fraction grid. However, it is always possible to get the values of the mixing fractions on the sampling grid, using Eq. (8.3). Note also that to visualize the discrete results of classical OMP analysis in a contour plot, as is usually done, it is anyway necessary to fit a function through the fraction values in order to interpolate and form the contour lines. POMP performs this parameterization a step earlier.
- (ii) Generally the number of basis functions n_B will be much smaller than the number of sampling positions N . Hence the number of unknowns (spline coefficients \mathbf{C}) to be estimated, namely $n_B n_S$, is much smaller than the original number of fractions that had to be estimated, which was $N n_S$. This implies there will be a certain “smoothing”, comparable to what happens in classical linear regression, which reduces the influence of individual noise contributions and makes the resulting mixing fraction estimates more robust. On the other hand, too much smoothing should be avoided, otherwise risking to discard significant features present in the data. Making this trade-off is an inclusive part of the POMP methodology and it will be further discussed in section 8.2.2.
- (iii) POMP possibly creates an opening to objectively optimise the SWT characteristics \mathbf{S} . As a consequence of the reduction in unknowns to be estimated (cf. (ii)), enough residual degrees of freedom are left over to extract additional information from the data, possibly about the elements of matrix \mathbf{S} . This topic will be the subject of Chapter 10, so it will not be further discussed here. Let us only mention that recently some work on the characterization of SWTs has been done in the framework of OMP analysis applied to time series [*Henry-Edwards and Tomczak, 2006*], so this issue is certainly of current interest.
- (iv) The main assumption underlying the POMP method is that the mixing contributions systematically vary with position. This implicitly constrains the possible values of the mixing fractions and therefore reduces their sensitivity to individual measurement errors.

Pointwise variation is interpreted as insignificant random effect which should be removed. So, clearly POMP analysis is not suitable to study small-scale (grid scale) phenomena. Exactly how much “detail” will be provided by the POMP analysis depends on the particular function used to describe the mixing patterns. This issue is addressed in the next section.

8.2.2. Model selection to ensure accuracy

The quality of the POMP analysis outcome obviously depends on the choices made concerning the functions parameterising the mixing fractions. Spline functions were chosen to describe the fraction fields, because they are very flexible in shape and stable to compute [*Dierckx*, 1995]. But an essential choice to be made, regardless of the type of functions used, concerns the complexity of the functions. Indeed, if a too simple function (or model) is chosen, some significant features present in the data will not be captured, so the results will be biased although they are robust. Conversely, if a too complex model is used (like OMP analysis in most occasions), not only the significant underlying processes are modelled, but a part of the stochastic noise as well. As a consequence, such a model will closely fit the data but is highly variable due to the noise’s influence. Summarizing, a model should have an “intermediate” complexity, such that it optimally balances accuracy and precision.

In the case of spline functions, this boils down to carefully selecting the number and order of the B-splines. This is comparable to choosing the best order in polynomial regression, or the number of variables to consider in multivariate regression. All these models are purely empirical (i.e. they do not have a physical interpretation) and the most suitable model complexity only depends on the data they describe. The complexity of the spline functions is expressed by the numbers (n_1 and n_z) and orders (k_1 and k_z) of the underlying B-splines can be chosen by performing the POMP analysis for all relevant combinations (n_1, n_z, k_1, k_z) and inspecting the results according to statistical rules designed to select an optimal model complexity (see Chapter 6). The only restriction is that $n_B = n_1 n_z \leq N$. The extreme situation when $n_B = N$ represents the

special case that the POMP model reduces to the original OMP solution, where the mixing fractions are estimated for each sampling position separately. Since it is not always obvious whether all assumptions underlying the statistical tests (e.g. normality of noise distributions) are satisfied with this kind of field data, we suggest to apply a number of these tests (see Chapter 6) and to use them together as a guide, in combination with common sense.

The goal of the parameterization is to make the resulting mixing fractions more robust, by only cutting off that part of the variability which is due to noise. Depending on the sampling density and the noise level, this “filtering” will also cut away a certain portion of some systematic, but small-scale, processes occurring (e.g. eddies). If these processes are of particular interest, they can be further investigated in the residual plots. In this perspective, the aim of a POMP analysis (combined with a statistically based selection of the splines complexity) is to only include those features in the mixing model which are most likely to be significant, while all the remaining information will be in the residuals. With classical OMP, almost everything is packed in the model.

The complexity of the splines functions, or the degree of smoothing, could also be interpreted in terms of ocean dynamics. If some processes are associated with specific scales, they could possibly be isolated using a POMP-like technique. The remaining processes could then be visualized by subtracting these results from those obtained with classical OMP analysis. This remark was first suggested by Matthias Tomczak during my visit to the Flinders University in summer 2006. It is included here as a hint for future applications, in this chapter and the remainder of this thesis the selection of the splines complexity is based on the statistical criteria.

To briefly summarize the POMP developments presented so far, parameterization of the unknown mixing fractions, combined with an effective choice of the splines complexity improves imperfections A and B mentioned in the introduction. How to handle remark C will be discussed in the next sections.

8.2.3. Weighting scheme

The final values for the mixing fractions can be greatly influenced by weighting the equations before optimisation. For instance, temperature and salinity measurements are usually associated with lower relative uncertainty than the other variables. This knowledge should be incorporated in the method such that these measurements are treated with greater importance. To achieve this every equation is multiplied with a weight symbolizing the reliability of the variable in question. Since this is not quantitatively possible for the mass conservation equation (uncertainty is zero?), we suggest to treat this equation as a constraint rather than as an equation to be optimised. This aspect is discussed in more detail in the next section.

In analogy with the approach followed in Chapters 4 and 6 we choose to use a weighting scheme compatible with the statistical Weighted Least Squares (WLS) framework [Pintelon and Schoukens, 2001]. This implies that the weights are the inverses of the uncertainty (in terms of standard deviation) associated with each equation. In other words, optimal values for the unknown spline coefficients \mathbf{C} (and hence the unknown mixing fractions) can be found by minimising the squared residuals multiplied by a weighting factor (w) which is the inverse of the standard deviation (std), i.e.

$$\begin{aligned} \hat{\mathbf{C}} &= \arg \min_{\mathbf{C}} \sum_{i=1}^N \sum_{j=1}^{n_y} \left(w_{ij} \left(\mathbf{S}_j \mathbf{C} \mathbf{B}_i^T - \mathbf{Y}_{ij} \right) \right)^2 \\ &= \arg \min_{\mathbf{C}} \sum_{i=1}^N \sum_{j=1}^{n_y} \left(\frac{\mathbf{S}_j \mathbf{C} \mathbf{B}_i^T - \mathbf{Y}_{ij}}{std_{ij}} \right)^2 . \end{aligned} \quad (8.11)$$

Here \mathbf{S}_j stands for the j^{th} row of \mathbf{S} , \mathbf{B}_i is the i^{th} row of \mathbf{B} and w_{ij} represents the weighting for the i^{th} sampling point for the j^{th} hydrographic property. In Eq. (8.11) it can also be seen that the weighting simultaneously serves to normalise the variables, so that they become dimensionless.

At this point we propose to consider only the measurements as source of uncertainty. This means that actually the mixing fractions (spline coefficients) are optimised in OWLS sense (cf. section 2.5). We suggest to assume the input, i.e. the source water types, to have negligible uncertainty, because of their conceptual nature. Since they are not really measured, but postulated, it is difficult to associate them with an uncertainty in terms of stochastic variance, as necessary in the WLS framework. An additional difficulty to include the input uncertainties, e.g. as proposed in Chapter 4, is that the source water types are not independent from sample to sample. This implies that covariances should be included as well, which are even more impossible to quantify. To put it as a paradox: in the name of objectivity, we propose to neglect the input uncertainty in the weighting, although it is undoubtedly significant.

In the most general case, the weighting factor varies with the property considered (e.g. a different weight for temperature than for salinity) and the location (e.g. a different weight for measurements made at the surface than for the deep observations). For simplicity and conform to common OMP practice, in this chapter one weight per variable will be attributed (no location dependency).

In practice, even when ignoring the SWT uncertainty, the difficulty remains to quantify the uncertainties to be used as weights. An accurate estimate of the standard deviation associated with each hydrographic property should be known. Ideally, this uncertainty is assessed based on the standard deviation of a number of replicates. However, this information is not always available because it is expensive and time-consuming to repeat experiments. In section 8.4.2 more details are given about the strategy used in this chapter to retrieve the uncertainty information. However, the results obtained using these weights will reveal the need for a better estimation of the measurement uncertainties, also allowing spatial dependency. In Chapter 9 a new strategy to estimate these uncertainties is presented.

8.2.4. Constrained optimisation scheme

If the mass balance equation (7.2) is not included in the mixing model, this has the advantage that the weighting can be performed following the same logic for all equations (cf. imperfection C in section 7.3). On the other hand, the optimised fractions *must* satisfy mass balance, otherwise they are not realistic. Both requirements can be assured simultaneously when the mass balance equation is imposed as an equality constraint during the optimisation. Generally in OMP analysis, an optimisation procedure is used that constrains the fractions to be positive (e.g. nonnegative least-squares algorithm “lsqnonneg” in MATLAB [OMP2, 2005]). The mass balance imposes an additional constraint that can be inserted in the optimisation algorithm (e.g. using the “lsqlin” MATLAB function). More details about the algorithm are given in section 8.2.5. The main consequences of this approach are:

- (i) Mass balance will be exactly (i.e. within computation precision) satisfied. In other words, only the mixing equations will have residuals. This is a healthy situation, because the mass balance equation does not represent the modelling of a measurement and is therefore not affected by measurement errors or environmental variability.
- (ii) The only equations really optimised are the mixing equations, so a consistent weighting scheme can be applied to all equations (see section 8.2.3).

Apart from these positive changes, the method and results are hardly influenced by treating mass conservation this way. For instance, the condition for mathematical solution (identifiability) that $n_s \leq n_v + 1$ still holds. Also, the resulting fractions will not differ much from those found with the original procedure, especially when the weight attributed to the mass balance equation was very large, as is the custom procedure. This confirms the observation of *Tomczak and Large* [1989]: “*Trial runs with (arbitrary chosen) larger weights for mass conservation show that the OMP analysis converges rapidly to nearly perfect mass conservation, with only minor adjustments in the mixing ratios*”. To summarize, imposing the mass conservation as a strict constraint does not influence

the results much. It is rather a technical issue allowing to attribute weights to all equations in a more consistent way.

8.2.5. POMP algorithm

Combining all suggestions made above to improve the classical OMP analysis, the POMP algorithm was designed. In this section a schematic overview of this algorithm is presented. All functions were written in MATLAB [MATLAB, 2005]. The goal is to estimate the spline coefficients using Eq. (8.6). Once these are known, the source contributions at each sampling point can be calculated with Eq. (8.3).

1. Input

- data matrix \mathbf{Y} ;
- source definition matrix \mathbf{S} ;
- weights for all variables;
- spline parameters (order, number).

2. Calculations

- Construct spline matrix \mathbf{B} .
- Eq. (8.6) must be vectorised, because the optimisation functions are written for this kind of input. Using the *vec* operator, all columns of a matrix are stacked underneath each other, giving one long column vector (see A.4 – A.6 in Appendix). For Eq. (8.6)

$$\mathbf{Y} = \mathbf{S} \cdot \mathbf{C} \cdot \mathbf{B}^T \Rightarrow \text{vec}(\mathbf{Y}) = \mathbf{B} \otimes \mathbf{S} \cdot \text{vec}(\mathbf{C}), \quad (8.12)$$

where $\text{vec}(\mathbf{Y})$ and $\text{vec}(\mathbf{C})$ stand for the vectorised matrices. \otimes is the Kronecker product operator which multiplies each element of the first matrix with the second matrix [Brewer, 1978], as defined in A.5 (Appendix).

- The equality constraints must be specified by matrices \mathbf{A}_{eq} and \mathbf{b}_{eq} defined as

$$\mathbf{A}_{eq} \text{vec}(\mathbf{C}) = \mathbf{b}_{eq} \quad (8.13)$$

For the mixing fractions to sum up to one in each point^{††}:

$$\mathbf{A}_{eq} = \mathbf{B} \otimes \mathbf{1}(1, n_s) \text{ and } \mathbf{b}_{eq} = \mathbf{1}(N, 1), \quad (8.14)$$

with $\mathbf{1}(\dots)$ a matrix of ones of the specified dimensions.

- The inequality constraints must be specified by matrix \mathbf{A}_{ineq} and vector \mathbf{b}_{ineq} defined as

$$\mathbf{A}_{ineq} \text{vec}(\mathbf{C}) \leq \mathbf{b}_{ineq}. \quad (8.15)$$

For all mixing fractions to be nonnegative:

$$\mathbf{A}_{ineq} = -\mathbf{B} \otimes \mathbf{I}_{n_s} \text{ and } \mathbf{b}_{ineq} = \mathbf{0}(n_s N, 1), \quad (8.16)$$

with $\mathbf{0}(\dots)$ a matrix of zeros of the specified dimensions.

- Optimal values for $\text{vec}(\mathbf{C})$ can now be found by solving the linear system Eq. (8.12), taking into account the mass balance and nonnegativity constraints. For this optimisation, we used the MATLAB function “lsqlin” [MATLAB, 2005] for constrained linear least squares optimisation.
- This function does not have the possibility to input a weighting matrix. Hence, the weighting has to be done before starting the optimisation: each element in $\text{vec}(\mathbf{Y})$ and in $\mathbf{B} \otimes \mathbf{S}$ has to be divided by the according standard deviation (see expression (8.11)).
- An optional input argument for this optimisation function is a starting point or initial guess value for the parameters to be

^{††} Note that the equality constraints are only imposed in the sampled points. However, from our tests it appears that if the constraints are satisfied in these points, they will be satisfied in any other point too. This is probably a consequence of the special property of splines to sum up to one in any given point in the interval where the splines are defined [Dierckx, 1995, p. 11].

estimated. We used the basic starting point for \mathbf{C} that is equivalent to all mixing fractions having the same values (all sources contribute equally over the domain).

3. Output

- optimal values for the spline coefficients $\text{vec}(\mathbf{C})$;
- optimal values for the source contributions at the sampling positions can be found by applying Eq. (8.3) or (8.5);
- residuals can be computed.

4. Spline complexity selection

To select the optimal spline complexity, the above calculations should be repeated for all relevant number of splines. For each fixed number of splines, the analysis should also be repeated for a number of orders (we tested orders 1 to 3). The best order for a fixed number of splines is that one giving the smallest residuals (residual sum of squares is a good measure), since varying the order does not change the number of model parameters or unknowns. Now the optimal number of splines still has to be selected, which we suggest to do based on the criteria developed and presented in Chapter 6.

8.3. Simulation tests

8.3.1. Synthetic dataset

A synthetic dataset was constructed based on the CLIVAR-SR3 data which will be analysed in section 8.4, such that it resembles a real situation. To obtain the synthetic dataset, the real data are slightly modified as to be perfectly modelled by 5 sources. The mixing fractions can exactly be parameterised by 10 splines in latitude direction and 4 in depth direction (this complexity will from now on be abbreviated as (10, 4)), and strictly satisfy both the equality and inequality constraints.

8.3.2. Error free situation

The trivial test of the POMP method is to apply it to the noiseless synthetic data and using the exact simulation model (the five true sources and a (10, 4) splines complexity). In that “perfect” case, the model should describe the data within computation precision. This is verified, and indeed, the residual cost function is smaller than 10^{-20} and the reconstructed mixing fractions (almost) perfectly match the simulated mixing fields. In short, if no errors are added, the POMP method works perfectly, which is the least we can expect. How does it behave when either systematic errors or (stochastic) noise is added?

8.3.3. Noiseless situation with wrong splines

In a second series of tests, still no noise is added to the data, but the mixing fractions are modelled by splines with a wrong complexity. The situations of both overmodelling (splines complexity of (14, 6) instead of (10, 4)) and undermodelling (complexity of (8, 2) instead of (10, 4)) are investigated.

Both overmodelling and undermodelling deliver residual cost function values much higher than the computation precision. In the case of undermodelling this is expected since not all features in the mixing fractions can be modelled by this simpler model and thus significant residuals will remain. In the case of overmodelling, this result may be more surprising, but it can be explained by the nature of the splines. Because the more complex splines are defined by different knots (automatically generated to be equidistant over the grid), they will not be able to capture the specific patterns of the simulated mixing fractions, although these were produced by “simpler” splines. This illustrates that model complexity is not always simply equated with the number of parameters, but that the structure of the model can also be a significant factor. In this case, the position of the knots also play an important role in the fitting power of the model. However, this is an extreme situation because splines were simulated. In the long term it is expected that by adding parameters the overall trend of the cost function will be decreasing.

A final important observation from these simulation tests concerns the shape of the residuals. For both the overmodelling and the undermodelling case, the residual histograms are relatively symmetric, which would hide the fact that model errors are present in this analysis. However, when inspecting the spatial distribution of the residuals (i.e. the residuals plotted as a function of position), systematic patterns appear which do point towards a non-optimal spline complexity (or knots position). Concluding from these observations, it may be misleading to judge the residuals whiteness only from histograms.

8.3.4. Noisy data and no systematic errors

In real applications, the measurements will always be disturbed by noise. So, it is of interest to investigate the effect of adding a stochastic component to the synthetic data on the performance of the POMP method. To ensure that the observed effects are only due to the added noise, no other errors are added, i.e. the true model (sources and splines) is used in the analysis.

It is meaningless to interpret the results of one such simulation, because it consists of only one noise realisation which does not infer anything about another realisation. The best strategy when investigating the behaviour of a stochastic component is to repeat the experiment a large number of times, each time with a different noise realisation, and to interpret the mean results. Accordingly, 1000 repeated simulations were performed, adding random, normally distributed noise to the synthetic data. The cost function was weighted correctly, i.e. with the inverses of the noise variances, such that it is expected that the average value of the residual cost function approximates the residual degrees of freedom. However, the average value of the residual WLS cost function (965) is higher than the number of degrees of freedom (892). Although the difference is not large, it is significant because it concerns an average value over 1000 realisations (the standard deviation associated with the average residual WLS cost function value = 1.3).

It appears that the presence of active inequality constraints (mixing fraction must be nonnegative) is the reason of this discrepancy. Due to

the added noise, the optimal solution, associated with the minimal cost function, may include negative mixing fractions. However, the inequality constraints in the algorithm make this solution unacceptable, and the final solution will be associated with a higher residual cost function. If we repeat the 1000 simulations, but omitting the inequality constraints, we do find the expected average value for the residual cost function: 893 ($df = 892$).

8.3.5. Conclusions

Although in the error free situation the POMP method works perfectly, these simulation tests indicate that elevated residual cost function values may be expected in real applications. First of all, because the spline parameterisation will never be perfect (real mixing distributions are probably not exactly spline-shaped). And even if the true model is a spline, it has been shown that (except for the true spline complexity) a model error is present and increasing the residual cost function. Secondly, even if no model errors are present, some inequality constraints are likely to become active due to the noise on the measurements. This will exclude the solution associated with the truly minimal cost function and thus force the solution to higher cost function values. Finally, a third reason why the residual cost function will be higher than expected has already been mentioned in section 8.2.3. Indeed, a OWLS cost function is used, implying that in real applications the uncertainty associated with the source water types is ignored. As a result the weights will underestimate the total uncertainty and the weighted cost function will be inflated compared to what is expected. This knowledge will be useful when analysing the real dataset in the next section.

8.4. Results on field data

In the previous sections of this chapter a refinement is proposed for the classical OMP analysis. After having tested it on simulations, this adapted procedure is now applied to a real dataset from the Southern Ocean.

8.4.1. Dataset and hydrological setting

The dataset of interest was collected during the CLIVAR-SR3 cruise (October – December 2001) with the R/V Aurora Australis. Figure 8.4 shows the study area.

Samples were taken at a number of depths at stations along the WOCE SR3 line ($\sim 142^\circ\text{E}$) in the Southern Ocean between Tasmania and Antarctica. The dataset represents the situation of late winter - early spring 2001. We focus on the Subantarctic region (northern part of the transect), roughly between 44°S and 53°S , and on the upper 1000 m of the water column.

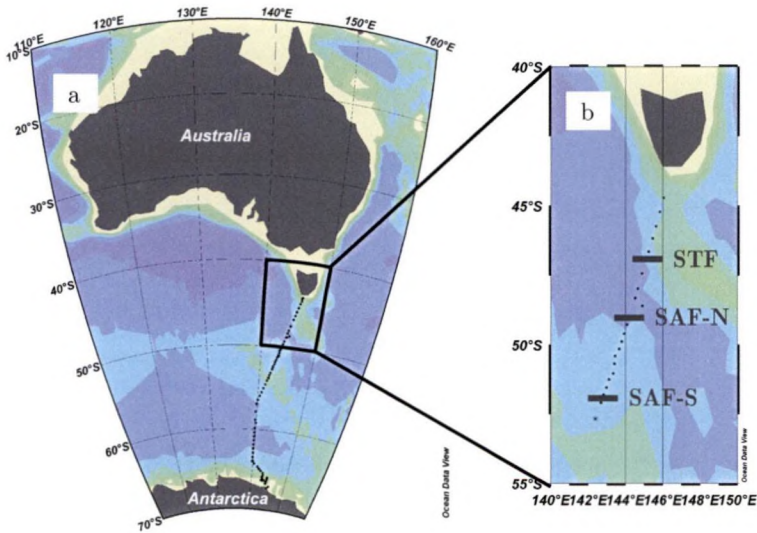


Figure 8.4. Maps of the study area with the positions of the stations from the CLIVAR-SR3 cruise: (a) the whole cruise, (b) only those stations considered in this study. The specific front positions are shown in grey along the cruise track and the meaning of the abbreviations is: STF = Subtropical Front, SAF-N and SAF-S = northern and southern branches of the Subantarctic Front.

At this point it may be of interest to discuss the specific hydrological setting of the area in some more detail, because it could facilitate the interpretation of the reconstructed mixing fields in the next section. The study area is characterised by a number of specific fronts and water masses, mostly defined based on potential temperature and salinity values and gradients. From north to south along the WOCE SR3 line one can find: the Subtropical Zone (STZ), the Subtropical Front (STF), the Subantarctic Zone (SAZ), the Subantarctic Front (SAF, northern and southern branch) and the Polar Front Zone (PFZ). The fronts along the SR3 section are observed in approximately the same location year to year [Rintoul *et al.*, 1997] but seasonal and interannual variabilities are apparent. The STF is generally reported to be located around 45°S [Rintoul and Bullister, 1999; Sokolov and Rintoul, 2002], with specific locations of 47°S for the CLIVAR cruise in spring 2001 [Jacquet *et al.*, 2004]. The SAF is the major front in this section. Sokolov and Rintoul [2002] consistently find two branches of the SAF along the SR3 line (SAF-N and SAF-S), at mean latitudes of 50.5°S and 52°S respectively. During spring 2001, the northern branch appeared at 49°S [Jacquet *et al.*, 2004]. In Figure 8.4 the specific front positions during the CLIVAR-SR3 cruise are indicated.

These displacements of front locations are related to changes in water mass formation. Seasonal and interannual changes of water mass properties have been identified in Aoki *et al.* [2005a], Chaigneau and Morrow [2002] and Rintoul *et al.* [1997]. The major water masses in the study area are: the Subtropical Surface Water (STSW), extending north of the STF, the Subantarctic Mode Water (SAMW), located between the STF and the SAF, the Antarctic Intermediate Water (AAIW), subducting between the PFZ and the STF, the Circumpolar Surface Water (CSW) in the PFZ and the Upper Circumpolar Deep Water (UCDW), located deeper in the PFZ.

8.4.2. Input into POMP

The conservative variables. For the mixing analysis, four conservative variables were used: potential temperature (Θ), salinity, PO and NO. PO and NO are tracers based on phosphate and nitrate, but corrected for respiration, according to Broecker [1974]: $PO = 170PO_4 + O_2$ and $NO = 9NO_3 + O_2$. Redfield ratios are similar to those used in other Southern Ocean studies [You, 2002; You et al., 2003; Lo Monaco et al., 2005].

The source water types relevant for the spring 2001 data are listed in Table 8.1 and shown in a Θ -S diagram in Figure 8.5. They are defined based on the dataset itself, usually on the boundaries of the study area, and ensuring that the data are well encircled by the SWTs. This way, the SWTs represent local water masses, which is justified since the study area is small and we are concerned with local processes. The selected SWTs are in good agreement with values reported in literature, listed in Table 8.2.

The weights. The weight attributed to each variable is given in Table 8.1. PO and NO weights are based on analytical precision information obtained from Mark Rosenberg [Rosenberg et al., unpublished report]. The values for temperature and salinity were taken larger than the analytical precision, because preliminary tests showed that just taking into account analytical precision greatly underestimates the total variability associated with these data. To assess these overall variabilities of temperature and salinity, raw CTD data (downcast and upcast) were used, but from a different cruise. For several depth horizons the standard deviation for temperature and salinity measurements within a 2 meter water layer was computed and the average of these numbers is given in Table 8.1. This approach implies that the water within a 2 meter layer is homogeneous, or at least reflects the same (local) water mass. It is also assumed that the variability is similar throughout the water column (because we take one averaged standard deviation), but this only confirms the earlier decision to take one weight per variable and not to take any location dependency into account (see section 8.2.3). In this case, the standard deviation of

the water characteristics within a 2 m layer is a better measure of the total variability than the analytical precision. *Chaigneau and Morrow [2002]* use a similar procedure but with latitudinal neighbouring data to estimate sea surface temperature and salinity precisions (giving resp. $\sim 0.01^\circ\text{C}$ and ~ 0.02).

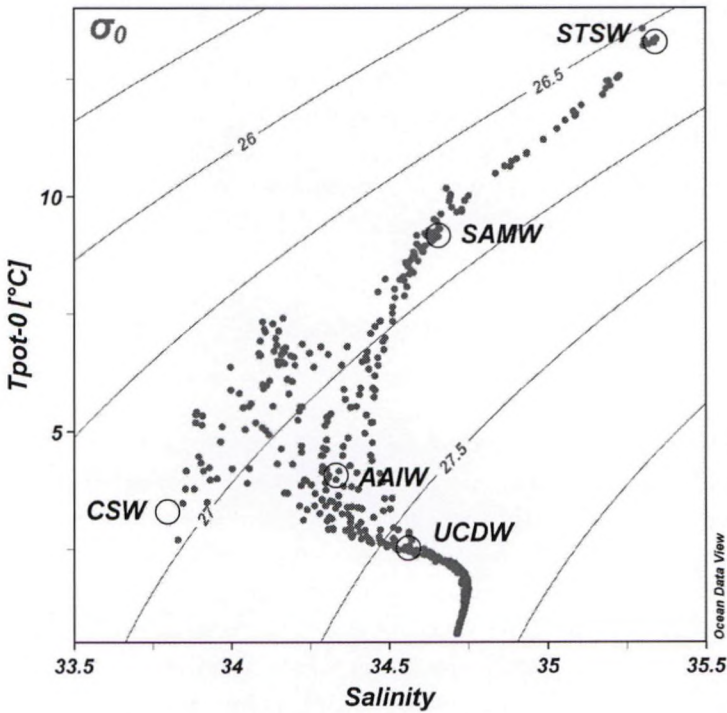


Figure 8.5. Θ -S diagram representing the CLIVAR-SR3 data (grey dots) and the SWTs (open circles). Only the data from the stations used in the analysis (i.e. north of 52.7°S) are shown, but for each station all sampling depths are represented (i.e. also the data deeper than 1000 m). Abbreviated SWT names stand for Antarctic Intermediate Water (AAIW), Subantarctic Mode Water (SAMW), Subtropical Surface Water (STSW), Circumpolar Surface Water (CSW) and Upper Circumpolar Deep Water (UCDW).

Table 8.1. Source water types and weights (derived from the dataset itself). Excluding the last column, these values form the elements of source water matrix **S**. Abbreviated names stand for Antarctic Intermediate Water (AAIW), Subantarctic Mode Water (SAMW), Subtropical Surface Water (STSW), Circumpolar Surface Water (CSW) and Upper Circumpolar Deep Water (UCDW).

	AAIW	SAMW	STSW	CSW	UCDW	weight=1/std
Θ (°C)	3.94	9.17	13.28	3.3	2.52	1/0.07
Salinity	34.355	34.655	35.34	33.79	34.56	1/0.011
PO (μM)	560	434	318	609	555	1/9
NO (μM)	490	399	297	562	486	1/8

Table 8.2. Literature values for the five sources considered in the mixing analysis, taken from the following references: ⁽¹⁾*Orsi et al.* [1995], ⁽²⁾*Rintoul et al.* [1997], ⁽³⁾*Rintoul and Bullister* [1999], ⁽⁴⁾*Rintoul and Sokolov* [2001], ⁽⁵⁾*Sokolov and Rintoul* [2002]. This overview serves as a validation of the values selected in Table 8.1.

	AAIW	SAMW	STSW	CSW	UCDW
Θ (°C)	3.4 – 5.6 ⁽³⁾ 4 – 5 ⁽⁴⁾	± 9 ⁽²⁾ 8.8 – 9.5 ⁽³⁾ 8.75 – 9.25 ⁽⁵⁾	11 – 16.25 ⁽⁵⁾	-	> 0.5 ⁽¹⁾ 2.6 ⁽³⁾
Salinity	< 34.4 ⁽³⁾ 34.35 – 34.44 ⁽⁴⁾	34;58 – 34.73 ⁽³⁾ 34.60 – 34.70 ⁽⁵⁾	35 – 35.50 ⁽⁵⁾	-	34.40 – 34.70 ⁽¹⁾ 34.585 ⁽³⁾
PO (μM)	-	-	-	-	Nutrient max ⁽¹⁾
NO (μM)	-	-	-	-	Nutrient max ⁽¹⁾

The splines' complexity. To parameterize the fraction fields, two-dimensional splines were used. As mentioned above, it is necessary to determine the optimal complexity of the spline function, i.e. the optimal orders (k_l and k_z) and numbers (n_l and n_z) of the splines in latitude, respectively depth direction. A number of model complexities were optimised, and the results are shown in Table 8.3 and Figure 8.6. For any (n_l, n_z) combination (i.e. same number of splines means same number of unknown parameters) the optimal (k_l, k_z) was determined as that combination delivering the lowest cost function. Some models were omitted from this analysis because they did not converge to any reasonable solution, probably due to bad conditioning. To choose the best (n_l, n_z, k_l, k_z) model, a number of strategies were compared.




Figure 8.6. Evolution of residual WLS cost function value for varying number of splines in latitude and depth directions.

- (i) CF test (cf. 6.3.1). This test appeared inadequate, because none of the models delivered acceptable cost function values, i.e. the residual cost functions were systematically too high. This was expected from the simulation tests in section 8.3, due to inequality constraints becoming active and non-optimal splines resolution. The high cost function values can also indicate the presence of other model errors such as nonconservative behaviour. Finally, an incorrect weighting of the WLS cost function may be causing the elevated values too, e.g. if the weighting underestimates the total variance of if the noise is nonnormal.

- (ii) CF evolution (cf. 6.3.2). In Figure 8.6 the evolution of the residual WLS cost function value is shown as a function of n_l and n_z . Graphically, one can see a relatively steep decrease in cost function between $n_l = 8$ and 9, which would suggest that $n_l = 9$ is optimal. Increasing the number of B-splines in the latitude direction further slowly decreases the residual cost function but it should be kept in mind that there are (only) 21 station in this direction and we do not want to overmodel. Besides this, it is arguable whether the change in cost function between $n_z = 4$ and 5 is significant. So from graphical inspection of the cost function evolution, we conclude that the optimal complexity is $(n_l, n_z) = (9, 4)$ or $(9, 5)$.

- (iii) MDL_S and AIC_S (cf. Table 6.1) for each model complexity are shown in Table 8.2. MDL_S returns $(n_l, n_z) = (9, 4)$ as optimal complexity, whereas AIC_S suggests $(13, 4)$. Since it has already been reported that AIC has the tendency to select too complex models, we decided to use $(n_l, n_z) = (9, 4)$. In addition, the mixing patterns associated with both models appear surprisingly similar. In other words, no “external” oceanographical information seems to specifically favour the more complex model. Finally, the fact that the model selected by MDL_S is so close to what was suggested by graphical analysis of the cost function evolution is a supplementary indication in favour of this model complexity. The associated orders are $(k_l, k_z) = (1, 1)$, because all other order combinations for this complexity returned higher cost functions.

Table 8.3. Overview of MDL_S and AIC_S values for the different model complexities. Minimal values are shaded.

MDL _S	n_t										
	8	9	10	11	12	13	14	15	16	17	18
4	64.4	57.4	61.4	59.6	61.9	62.2	67.9	67.5	73.1	73.5	-
5	69.8	61.9	67.6	65.9	69.5	69.1	77.6	77.7	-	85.2	-
n_z 6	78.6	69.1	77.8	77.3	82.7	82.1	93.1	94.2	-	105	-
7	90.2	79.3	90.1	91.9	98.5	99.5	117	120	-	138	-
8	102	91.1	105	106	118	117	-	-	-	168	-

AIC _S	n_t										
	8	9	10	11	12	13	14	15	16	17	18
4	41.9	35.9	37.1	34.8	35.1	34.2	36.3	35.1	37.1	36.4	-
5	42.2	35.9	37.7	35.5	36.2	34.8	38.0	37.0	-	38.5	-
n_z 6	44.5	37.5	40.5	38.8	40.0	38.5	42.3	41.6	-	44.4	-
7	48.2	40.5	44.1	43.3	44.7	43.7	49.9	49.8	-	54.5	-
8	51.5	44.6	48.7	47.2	50.8	48.6	-	-	-	63.0	-

8.4.3. Results from POMP analysis: mixing fractions

The mixing fractions estimated by the POMP method are shown in Figure 8.7. To facilitate interpretation, sampling depths, density isolines and front positions are indicated as well. The five selected water masses are clearly significant in the investigated spatial window.

The Antarctic Intermediate Water (AAIW) appears between 200 and 400 m at 52°S and sinks northward below the SAMW and reaches depths greater than 1000 m north of the STF (~47°S). Its apparent extension at 1000 m in the northern part of the transect (~45°E) could reflect the presence of another variety of AAIW [Rintoul and Bullister, 1999]: a “northern” type of Tasman Sea origin. The fact that the density isolines rise at this position supports this hypothesis. Similarly,

the discontinuity of the AAIW branch between 51 and 52°S probably reflects that the water immediately south of the SAF does not have the same origin as the low salinity core north of the SAF. Instead, the AAIW is supplied by a number of water types upstream and by the time they enter the studied area, these waters have developed distinct signatures [Rintoul and Bullister, 1999]. This could explain why one AAIW definition cannot perfectly account for all intermediate water samples in the investigated dataset.

As reported in *Sokolov and Rintoul* [2002] the Subantarctic Mode Water (SAMW) appears as the principal water mass extending north of the northern branch of the Subantarctic Front (SAF-N; ~49°S) in the upper 700m water column, a zone characterized by extremely homogeneous density. Between 700m and 1000m, SAMW still has a noticeable contribution, but it mixes with AAIW. The SAMW core between the STF and SAF-N reaches the surface, which illustrates the deep winter mixing in this zone, as noticed by *Rintoul and Trull* [2001]. North of the STF, SAMW does not reach the surface anymore, its contribution is concentrated between 300 and 800m. Like the AAIW, the SAMW does not form a single homogeneous water mass [Rintoul and Bullister, 1999]; it results in the large extension of SAMW observed here.

The surface waters north of the STF (~ 47°S) form a homogeneous water mass, the Subtropical Surface Water (STSW). It corresponds to the Subtropical Lower Water identified by *Sokolov and Rintoul* [2002]. STSW appears as a compact water mass up to 250m, lying on top of and slightly mixing with SAMW.

At the surface of the most southern part of the window, a fourth surface water is significant: the Circumpolar Surface Water (CSW). Its core lies at ~52°S and its contribution extends to 500-600m in depth and to SAF-N (~49°S) in latitude. As for AAIW an unexpected contribution in the north is observed, which is probably due to the poor resolution in this zone, as mentioned before.

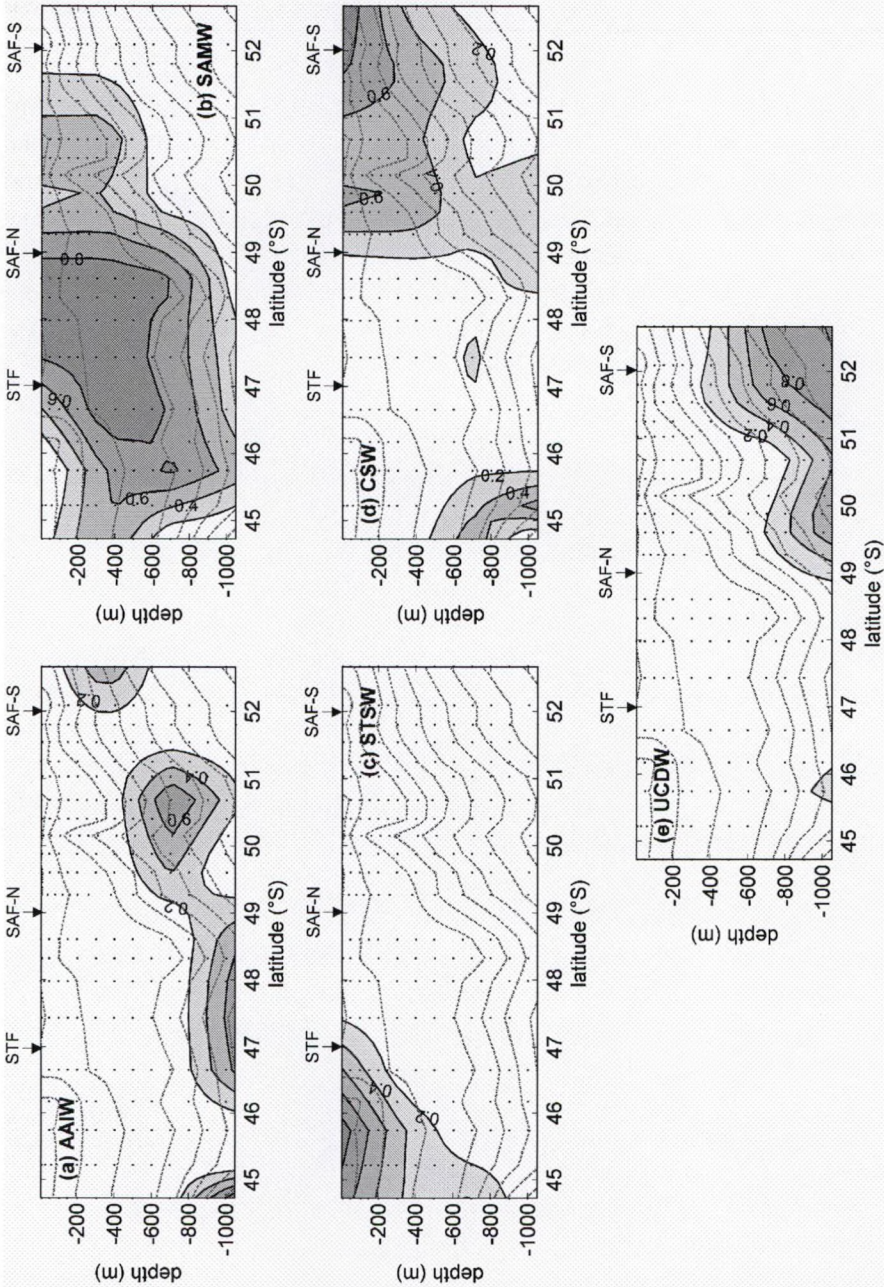


Figure 8.7. Mixing fractions associated with (a) Antarctic Intermediate Water, (b) Subantarctic Mode Water, (c) Subtropical Surface Water, (d) Circumpolar Surface Water and (e) Upper Circumpolar Deep Water, estimated using the POMP method and the CLIVAR-SR3 data. Density isolines are plotted as dotted grey lines and range from 26.6 (upper north corner) to 27.6 (lower south corner) with increments of 0.1. Dots represent the sampling depths.

Finally, the southern part of the section undergoes significant contributions of the Upper Circumpolar Deep Water (UCDW) between 800 and 1000 m. This deep water is normally found between 1500 and 2500m north of the SAF, and rises to depths of 1000 m just south of the SAF-N and 500m south of the SAF-S [Sokolov and Rintoul, 2002], coinciding with the reconstruction shown in Figure 8.7.

8.4.4. Uncertainties

To quantitatively interpret the mixing distributions, it is necessary to quantify the variability of the estimated mixing coefficients due to measurement errors (these latter can be found in Table 8.1). One thousand Monte-Carlo simulations were performed, where in each simulation the original data are perturbed with a random number sampled from the measurement noise distribution [Coatanoan *et al.*, 1999]. The uncertainty of the mixing coefficients is different for each source: 3.5% for AAIW, 1.7% for SAMW, 0.8% for STSW, 1.6% for CSW and 1.9% for UCDW (average values of 1 standard deviation). Note that this uncertainty is much smaller (approximately a factor 10) than the variability if no nonnegativity constraints are applied. This means that the constraints are active and reduce the space of possible mixing fractions, resulting in a smaller variability.

Remember that we chose only to take the output (measurement) uncertainties into account (cf. 8.2.3) in the construction of the weights, because it is so difficult to attribute uncertainties to the source water types (cf. 8.2.3). For the same reason, the above Monte Carlo standard deviations neglect the effect of source water type uncertainty. Therefore, it must be born in mind that these mixing fraction uncertainties are probably underestimations.

8.4.5. Residuals

The reliability of the estimated mixing fraction fields can be further assessed by examining the residuals. Notice that the residuals are defined as the difference between model reconstruction and

measurement (see expression (8.7)). The residuals for this analysis are presented in Figure 8.8. In the figures on the left, showing the residuals as a function of position, it is clear that the residuals are higher in the upper waters, which is probably an indication that the overall uncertainty associated with a property is also dependent of depth. But the histograms on the right show that on the whole the residuals behave as symmetrically distributed variables, centred around zero. This does not guarantee the absence of model errors as seen in the simulation tests. But combined with the spatial plots of the residuals, it is an indication that all major (large-scale) features are captured by the model, and hence that the residuals represent random or small-scale features (including features of ocean physics such as eddies and other products of geostrophic turbulence that can be treated as random).

However, when inspecting the absolute values of the residuals, they appear to be significantly larger than the uncertainties associated with each tracer (cf. Table 8.1). This can either be due to significant model errors present or to an underestimation of the uncertainties. We argue that the first option cannot be of large importance because of the “normality” of the residual histograms, mirrored by the lack of specific large regions in the study area associated with high residuals. On the other hand, it seems likely that the uncertainties are underestimates of the real variability in the system, because (i) the source uncertainties were neglected, and (ii) the measurement uncertainties are so difficult to estimate because no sampling replicates are available. This problem of estimating the measurement uncertainties is further addressed in Chapter 9, where a new estimation algorithm is presented.




Figure 8.8. Residuals associated with the POMP analysis on the CLIVAR-SR3 data for the four variables: potential temperature (first row), salinity (second row), PO (third row) and NO (last row). In the left column they are represented as a function of position, the right column shows them in histograms. Residuals are defined as observation – modelled value.

8.4.6. Comparison with classical OMP analysis

As an illustration of the different performance of the POMP approach, a classical OMP analysis is performed using the same observations and SWT definitions. That is, the fractions are estimated directly from Eq. (7.2) and (8.1), instead of parameterising them first and estimating the spline coefficients from Eq. (8.6). The same weighting scheme and contour visualization is applied as for the POMP results (cf. expression (8.11)), so any differences are due to the parameterization.

In Figure 8.9 the resulting mixing fraction fields are shown. The water masses' average mixing pattern corresponds to the results from the POMP method (see Figure 8.7). However, the distributions clearly are much more "fitful", due to the fact that the mixing fractions are optimized for every point independently. This approach allows the noise, inevitably present in the measurements, to fully propagate in the final mixing fractions (cf. imperfection B). This can lead to locally unrealistic features such as the "hole" present in the SAMW fractions around 49°S. To provide additional evidence to evaluate the POMP vs. OMP results, density isolines were added on the plots in Figures 8.7 and 8.9. Most, although not all, features present in the density pattern are recovered in the POMP fractions. Furthermore, the SAMW nicely fills its characteristic pycnostad, whereas from OMP analysis it would seem that CSW and STSW also contribute in this area. This is probably the result of the pointwise strategy allowing discontinuous mixing distributions. Conversely, the mixing distributions provided by the POMP procedure are modelled as a function of position and smoothed in a natural way because the effect of the noise on the individual points is suppressed by the parameterization. It is important to stress that if the complexity of the spline functions is chosen accurately, all significant features that can be described by linear mixing will be represented in the final mixing fractions. In other words, the fact that the fraction fields calculated by the POMP method are smoother does not mean that significant information has been lost. However, some small-scale features (e.g. eddies) can obviously only be captured if the sampling density is high enough, which is also true for the classical OMP analysis.

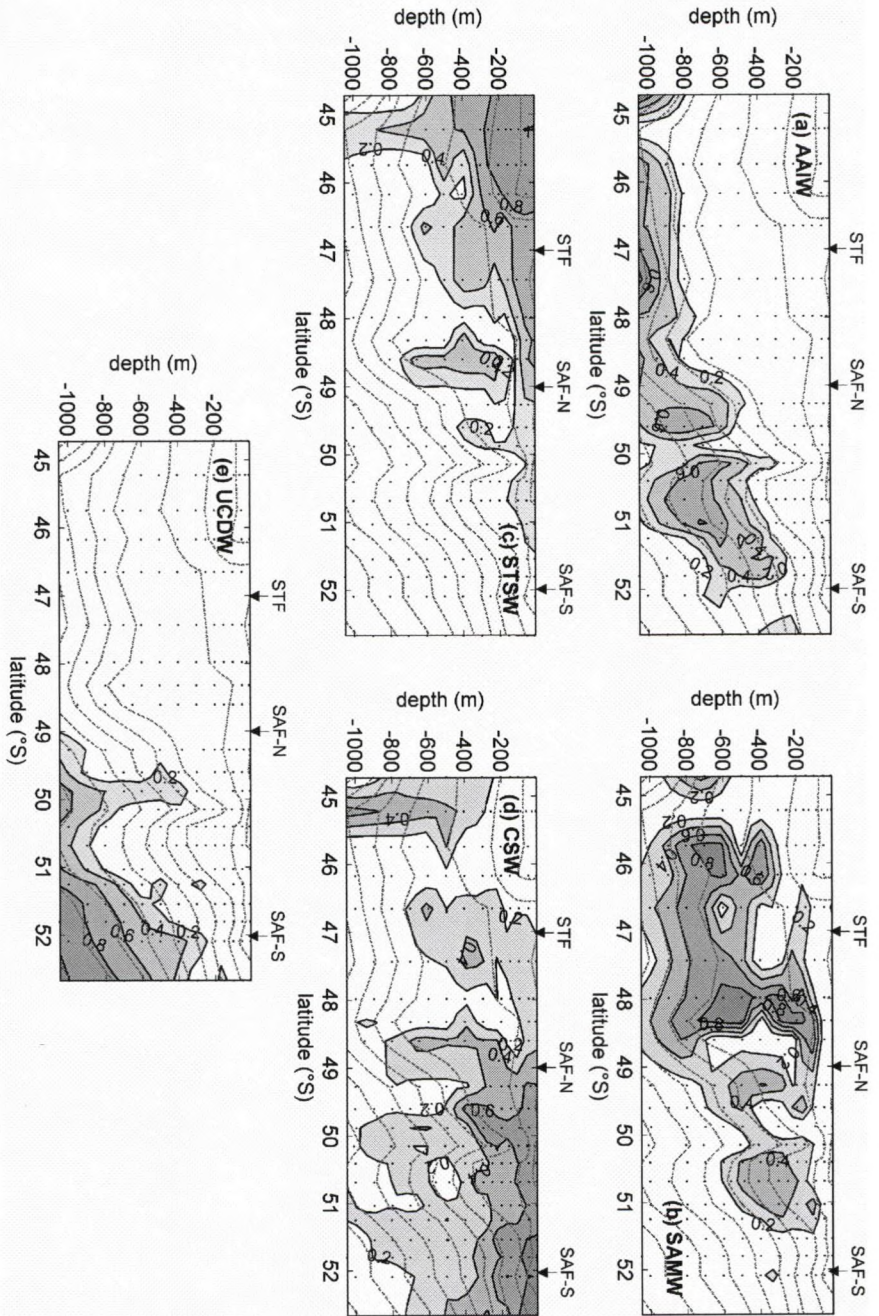
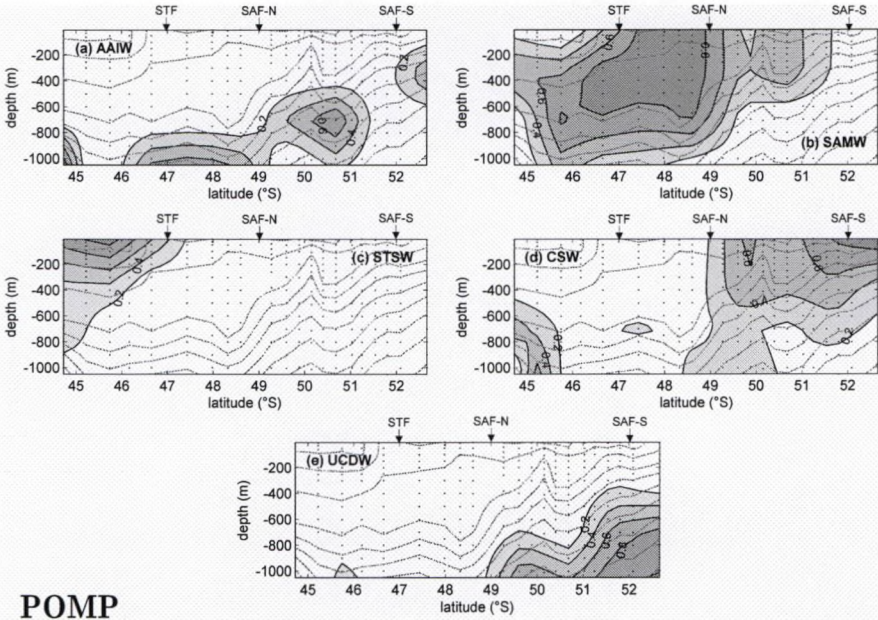


Figure 8.9. Mixing fractions estimated using the classical OMP method and the CLIVAR-SR3 data. Density isolines are plotted as dotted grey lines and range from 26.6 (upper north corner) to 27.6 (lower south corner) with increments of 0.1. Dots represent the sampling depths.



POMP

OMP

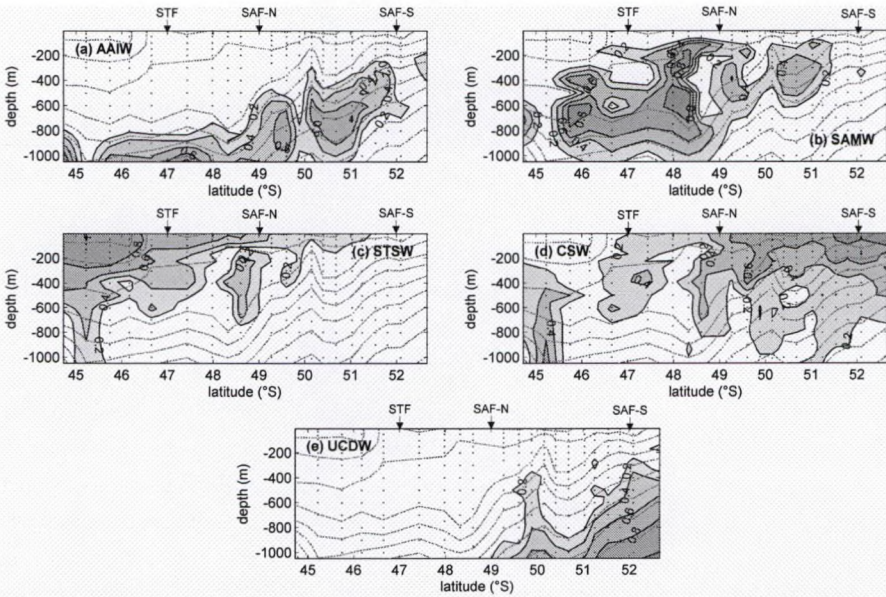


Figure 8.9bis. Copies of Figures 8.7 and 8.9, to facilitate comparison of results found by POMP versus classical OMP analysis.




Figure 8.10. Residuals associated with the classical OMP analysis on the CLIVAR-SR3 data for the four measured variables and the mass balance equation: potential temperature (first row), salinity (second row), PO (third row), NO (fourth row) and mass balance (last row). In the left column they are represented as a function of position, the right column shows them in histograms. Residuals are defined as observation – modelled value.

The residuals for the classical OMP analysis are shown in Figure 8.10; they are much smaller than those shown in Figure 8.8. Small residuals are not necessarily the result of an accurate modelling, but could also be due to overmodelling (cf. linear regression through two data points). To make a real quantitative statement about the residual level which is acceptable, a more profound uncertainty analysis would be required, taking into account all aspects of environmental and SWT variability in addition to mere measurement uncertainties, as already discussed above.

An additional consequence of overmodelling is that the mixing fractions are much more sensitive to stochastic errors. This is illustrated by the uncertainties estimated by Monte-Carlo simulations, as described in section 8.4.4. The uncertainties of the mixing fractions derived by the classical OMP analysis are much higher than with the POMP method, namely: 10 % (vs. 3.5%, see section 8.4.4) for AAIW, 14% (vs. 1.7%) for SAMW, 9% (vs. 0.8%) for STSW, 10% (vs. 1.6%) for CSW and 6% (vs. 1.9%) for UCDW.

8.5. Conclusions

In Chapter 7 the classical OMP was presented and emphasis was put on its attractive simplicity, but also on a number of imperfections:

- A.* No position dependence.
- B.* Few or no residual degrees of freedom.
- C.* Mass balance is not strictly imposed and inconsistent weighting.
- D.* Subjectivity of source water type definition.

In an attempt to address at least some of these imperfections, the Parametric OMP analysis (POMP) was developed in this chapter. By parameterising the mixing fractions as functions of position, the estimated mixing fractions become more robust and, obviously, explicitly dependent of position (imperfection A). The increased robustness is a consequence of the reduction of number of unknowns to be estimated (imperfection B). However, this reduction must be

supported by the data, and hence the determination of the optimal model complexity is an essential part of the POMP procedure. Besides the parameterisation, a weighting scheme compatible with the WLS framework is proposed. This scheme can be applied to all equations in the same way if mass balance is imposed as a strict equality constraint, instead of as an equation to be optimised (imperfection C). However, to achieve a really accurate weighting all sources of uncertainty must be correctly incorporated, which is not straightforward. The measurement (output) uncertainties are difficult to estimate since to replicate samples are available, and the SWT uncertainties are even more cumbersome to include because of their conceptual and dependent nature. Summarising the results regarding the correction of the OMP analysis imperfections, we could state that imperfections A and B are solved by the POMP methodology. Imperfection C is partly improved but a really accurate weighting remains complicated. And imperfection D is not addressed at all. In the next chapters, more attention is direction to these remaining issues.

Besides being an improvement compared to the classical OMP analysis, several other results of POMP are worth summarising. Applying the POMP procedure to a synthetic dataset, revealed a number of noteworthy points. In particular, it is useful to know that in real applications the residual OWLS cost function value will most probably be higher than expected, even if no real model errors are present.

The POMP method has also been applied to the CLIVAR-SR3 dataset. The derived mixing distributions correspond to what is expected from literature and are more reliable (robust) than those estimated using the classical OMP analysis. The residual analysis suggests that the residuals are normally distributed but their standard deviations are significantly larger than expected from the assigned measurement uncertainties, confirming the above anticipation of an inflated residual cost function. Some of the causes of this inflation may not be easy to tackle (e.g. inequality constraints becoming active or the inclusion of the SWT uncertainty in the weighting), but one possible improvement will be considered in more detail in the next chapter: a more accurate estimation of the measurement noise variances.

Chapter 9:

Estimation of heteroscedastic noise variances

“Uncertainty is a term used in subtly different ways in a number of fields, including philosophy, statistics, economics, finance, insurance, psychology, engineering and science. It applies to predictions of future events, to physical measurements already made, or to the unknown.”

From Wikipedia, the free encyclopedia.

9.1. Introduction

When knowledge about their noise levels is available, data can be processed in a much more rigorous way. For instance, model parameters can be optimised taking into account differences in noise level in the dataset (cf. 2.5.2), uncertainties on the optimised parameter values can be computed (cf. Chapter 4), the “goodness” of the model fit can be assessed (cf. Chapter 6), etc. To put it simply, a distinction can be made between what is (more) significant and what is not. Consequently, noise information is crucial to do any quantitative interpretation of data.

Here we are concerned with the general case of noise variances varying with each sample (heteroscedastic noise). The necessity of estimating these variances arose in the previous chapter, when processing oceanographical data, for all the reasons mentioned above. Indeed, it was noticed that when the measurement uncertainties given by the oceanographers were used, the total variance was underestimated, probably due to the fact that sampling uncertainty was not taken into account. Unfortunately, this kind of experiments cannot easily be repeated. A first attempt to include the sampling variability (cf. section 8.4.2) appeared insufficient, so another way had to be found to estimate the total uncertainty associated with the measurements.

This chapter is based on *de Brauwere et al.* [2007b].

9.2. Methods

9.2.1. Literature

One possible strategy to obtain this variance information is by performing replicate measurements. In that case, for example the sample variances can be calculated, although (iterative) methods based on the residuals are proven to be more reliable, especially when the number of replicates is small [*Rao*, 1970; *Jacquez and Norusis*, 1973; *Carroll and Cline*, 1988; *Chen and Shao*, 1993]. Yet, replicate measurements, are often not performed because they are too expensive. An alternative

strategy proposed in literature is to model the noise variances parametrically, for instance as a function of the response [Carroll and Ruppert, 1982; Davidian and Carroll, 1987; Davidian and Haaland, 1990]. For this approach to be useful, the class of the variance function (e.g. polynomial, rational, ...) should be well known, which is not always obvious.

This chapter presents a general procedure to extract heteroscedastic (i.e. varying with position) noise information from observations when neither replicates or a well known variance function are available. With the same objective, Vanlanduit [2001, §4.2] used the absolute value of the residuals as an estimate of the measurement noise standard deviation, but suggested to first locally average the residuals and subtract the local mean value. This procedure increases the accuracy of the estimates, but the price to be paid for this averaging is that information about the local variations of the noise variances is lost. We will use a similar approach, but estimating the noise variances grouped in predefined subsets. In order words, it is assumed that the noise variances vary in a stepwise manner such that within groups of observations the variance can be assumed constant.

The algorithm also starts from residuals, so it necessitates a model to describe the data. However, if model errors are present, the residual is a biased measure of the stochastic noise. Therefore, a first order correction is included in the procedure to suppress possible model errors and hence make the noise variance estimation relatively independent of the specific model used. This is important because the estimated noise variances may be used in a later step exactly to detect model errors. A final difference with the procedure proposed by Vanlanduit [2001] is that it is taken into account that if the noise variances vary over the measurement domain, the model parameters will be more fitted to those measurements associated with low uncertainty. Hence, the residual degrees of freedom will not be uniformly distributed over the domain. This effect can especially become important if the number of parameters is not much larger than the number of data.

9.2.2. No model errors present

Consider observations \mathbf{y} (vector of length N) modelled by a linear model

$$\mathbf{y} = \mathbf{J}\boldsymbol{\theta} + \mathbf{e}, \quad (9.1)$$

where \mathbf{J} stands for the Jacobian matrix associated with the model and \mathbf{e} is the error vector. In the following derivation we assume that \mathbf{y} represents measurements of only one variable. In the case of a multivariate \mathbf{y} (n_v variables), the equations remain valid. Only the size of data-related quantities (\mathbf{y} , \mathbf{J} , \mathbf{r} , ...) changes from N to Nn_v .

For models which are nonlinear in the parameters, Eq. (9.1) and all following equations represent a first order approximation. The model parameters $\boldsymbol{\theta}$ (vector of length n_θ) can be optimised by minimising a least squares cost function. This results in optimised parameter values

$$\hat{\boldsymbol{\theta}} = (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T \mathbf{y}. \quad (9.2)$$

The residuals \mathbf{r} can be written as the remaining difference between data and model, using Equations (9.1) and (9.2)

$$\mathbf{r} = \mathbf{y} - \mathbf{J}\hat{\boldsymbol{\theta}} = \mathbf{y} - \mathbf{J}(\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T \mathbf{y}. \quad (9.3)$$

In order to cope with heteroscedasticity, the complete dataset \mathbf{y} is cut into n_h subsets. For every subset a variance will be estimated. If multivariate data are analysed, each variable is associated with a different variance, i.e. in that case $n_h n_v$ variances are to be estimated. In analogy with Eq. (9.3), the residuals of subset $h \in \{1, 2, \dots, n_h\}$ are given by

$$\mathbf{r}_h = \mathbf{y}_h - \mathbf{J}_h \hat{\boldsymbol{\theta}} = \mathbf{y}_h - \mathbf{J}_h (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T \mathbf{y}, \quad (9.4)$$

where the subscript h indicates the subset considered. Eq. (9.4) expresses that the subset residuals \mathbf{r}_h are not only a function of the subset data \mathbf{y}_h , but depend on the whole dataset \mathbf{y} , via the estimated




Figure 9.1. Schematic representation of the matrices \mathbf{K}_h and \mathbf{K}_{-h} . The total dataset \mathbf{y} has length N . Index h defines the data indices of the subset data \mathbf{y}_h . Subset h counts N_h elements. The general case is depicted where the data indices forming a subset are not subsequent.

model parameters $\hat{\boldsymbol{\theta}}$. Since \mathbf{y} contains the subset data \mathbf{y}_h , Eq. (9.4) can be decomposed in a term depending on \mathbf{y}_h and a term depending on the remaining data, denoted by \mathbf{y}_{-h}

$$\mathbf{r}_h = \mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h}, \quad (9.5)$$

where $\mathbf{P}_h = \mathbf{I}_{N_h} - \mathbf{K}_h$ and the columns of \mathbf{K}_h and \mathbf{K}_{-h} form $\mathbf{J}_h (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T$ as depicted in Figure 9.1.

If no model errors are present, these residuals can be used to estimate the noise variances. The expected value of the residual variance in the h^{th} subset is given by

$$E\{\mathbf{r}_h^T \mathbf{r}_h\} = E\left\{(\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h})^T (\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h})\right\}. \quad (9.6)$$

A scalar is equal to its trace, hence

$$E\{\mathbf{r}_h^T \mathbf{r}_h\} = \text{trace} \left[E\left\{(\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h})^T (\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h})\right\} \right].$$

To further rearrange these equations, trace properties (A.2) – (A.3) are used, such that

$$\begin{aligned} E\{\mathbf{r}_h^T \mathbf{r}_h\} &= \text{trace} \left[E \left\{ (\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h}) (\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h})^T \right\} \right] \\ &= \text{trace} \left[E \left\{ \mathbf{P}_h \mathbf{y}_h \mathbf{y}_h^T \mathbf{P}_h^T \right\} \right] + \text{trace} \left[E \left\{ \mathbf{K}_{-h} \mathbf{y}_{-h} \mathbf{y}_{-h}^T \mathbf{K}_{-h}^T \right\} \right] \\ &\quad - \text{trace} \left[E \left\{ \mathbf{P}_h \mathbf{y}_h \mathbf{y}_{-h}^T \mathbf{K}_{-h}^T \right\} \right] - \text{trace} \left[E \left\{ \mathbf{K}_{-h} \mathbf{y}_{-h} \mathbf{y}_h^T \mathbf{P}_h^T \right\} \right]. \end{aligned}$$

Considering that $\mathbf{P}_{(-)h}$ and $\mathbf{K}_{(-)h}$ do not contain stochastic components, so their expected value equals their value, and further rearranging the trace terms according to property (A.3) gives

$$\begin{aligned} E\{\mathbf{r}_h^T \mathbf{r}_h\} &= \text{trace} \left[E \left\{ \mathbf{y}_h \mathbf{y}_h^T \right\} \mathbf{P}_h^T \mathbf{P}_h \right] + \text{trace} \left[E \left\{ \mathbf{y}_{-h} \mathbf{y}_{-h}^T \right\} \mathbf{K}_{-h}^T \mathbf{K}_{-h} \right] \\ &\quad - \text{trace} \left[E \left\{ \mathbf{y}_h \mathbf{y}_{-h}^T \right\} \mathbf{K}_{-h}^T \mathbf{P}_h \right] - \text{trace} \left[E \left\{ \mathbf{y}_{-h} \mathbf{y}_h^T \right\} \mathbf{P}_h^T \mathbf{K}_{-h} \right]. \end{aligned} \tag{9.7}$$

The expected values in (9.7) can be further reformulated assuming no model errors are present (i.e. $\mathbf{y}_h = \mathbf{y}_h^0 + \mathbf{e}_h$ with \mathbf{y}_h^0 representing the true value and \mathbf{e}_h only stochastic noise; and analogous for \mathbf{y}_{-h}) and white noise with constant variance within each subset (i.e. $E\{\mathbf{e}_h\} = 0$, $E\{\mathbf{e}_h \mathbf{e}_h^T\} = \sigma_h^2$ and analogous for \mathbf{y}_{-h}):

$$\begin{aligned} E\{\mathbf{r}_h^T \mathbf{r}_h\} &= \text{trace} \left[E \left\{ (\mathbf{y}_h^0 + \mathbf{e}_h) (\mathbf{y}_h^0{}^T + \mathbf{e}_h^T) \right\} \mathbf{P}_h^T \mathbf{P}_h \right] \\ &\quad + \text{trace} \left[E \left\{ (\mathbf{y}_{-h}^0 + \mathbf{e}_{-h}) (\mathbf{y}_{-h}^0{}^T + \mathbf{e}_{-h}^T) \right\} \mathbf{K}_{-h}^T \mathbf{K}_{-h} \right] \\ &\quad - \text{trace} \left[E \left\{ (\mathbf{y}_h^0 + \mathbf{e}_h) (\mathbf{y}_{-h}^0{}^T + \mathbf{e}_{-h}^T) \right\} \mathbf{K}_{-h}^T \mathbf{P}_h \right] \\ &\quad - \text{trace} \left[E \left\{ (\mathbf{y}_{-h}^0 + \mathbf{e}_{-h}) (\mathbf{y}_h^0{}^T + \mathbf{e}_h^T) \right\} \mathbf{P}_h^T \mathbf{K}_{-h} \right] \\ &= \text{trace} \left[\begin{array}{c} \underbrace{E\{\mathbf{y}_h^0 \mathbf{y}_h^0{}^T\}}_{\mathbf{y}_h^0 \mathbf{y}_h^0{}^T} + \mathbf{y}_h^0 \underbrace{E\{\mathbf{e}_h^T\}}_0 + \mathbf{y}_h^0{}^T \underbrace{E\{\mathbf{e}_h\}}_0 + \underbrace{E\{\mathbf{e}_h \mathbf{e}_h^T\}}_{\sigma_h^2} \end{array} \mathbf{P}_h^T \mathbf{P}_h \right] \\ &\quad + \dots \end{aligned}$$

and all equations have to be solved simultaneously for the σ_h^2 .

The last term in Eq. (9.11) can be further reorganised in order to isolate the noise variance vector from the second term. Doing this, estimates for the noise variances $\hat{\sigma}_h^2$ ($h = 1, \dots, n_h$) can finally be found as the solution of a set of linear equations of the form

$$\mathbf{v} = (\mathbf{T}_1 + \mathbf{T}_2) \hat{\boldsymbol{\sigma}}^2 = \mathbf{T} \hat{\boldsymbol{\sigma}}^2, \quad (9.12)$$

where \mathbf{v} is a vector containing the residual subset variances, \mathbf{T}_1 is a diagonal matrix (cf. first term in Eq. (9.11)), \mathbf{T}_2 is an off-diagonal matrix (since $\Sigma_{\cdot h}$ refers to all but the h^{th} variance, cf. Eq. (9.9) and (9.11)) and $\hat{\boldsymbol{\sigma}}^2$ is a vector with the noise variance estimates. So, estimating $\hat{\boldsymbol{\sigma}}^2$ now boils down to solving a classical linear inverse problem (9.12).

Several remarks can be made about Eq. (9.12):

- (i) Matrix \mathbf{T} in Eq. (9.12) has size $s \times s$ and expresses how the degrees of freedom are spread throughout the dataset (subsets). The sum of all elements in \mathbf{T} gives the exact number of degrees of freedom of the residuals, i.e. generally $N - n_\theta$. This can easily be checked for the special case of only one subset: residual variance = $(N - n_\theta) \hat{\boldsymbol{\sigma}}^2$. The elements of \mathbf{T}_1 (diagonal elements of \mathbf{T}) express the “local” degrees of freedom of each subset, i.e. $N_h - n_{\theta,h}$, where $n_{\theta,h}$ stands for the number of model degrees of freedom “used” in subset h . If a part of the dataset was measured with much lower uncertainty, the model parameters were mainly fitted to these data. Hence the “loss” of degrees of freedom is largest in this region of the dataset and this will be confirmed by lower elements in \mathbf{T}_1 . If, on the other hand, the noise is homoscedastic, all elements of \mathbf{T}_1 have approximately the same value. \mathbf{T}_1 symbolises that in a first approximation, $\hat{\sigma}_h^2 = v_h / (N_h - n_{\theta,h})$, where v_h is the h^{th} element of \mathbf{v} . \mathbf{T}_2 is a correction on this local view, expressing that the residual variance in a given subset is not only influenced by the noise variances in its own subset, but also by those in all other subsets.

- (ii) Applying the strong law of large numbers [Kendall and Stuart, 1979; Lukacs, 1975], shows that this estimator is consistent. In other words, it (almost surely) converges to the true value with increasing number of data. Alternatively, when the number of data is too limited, the estimation can be unreliable. In section 9.2.4 some practical actions are discussed to improve the general reliability of the estimates.
- (iii) The estimated variances represent all stochastic variations contained in the residuals. It will in general be (much) larger than the precision of the analytical measurement, because it includes the effect of input noise (uncertainty of \mathbf{J}) and all additional variation sources like sampling, sample preparation, etc.
- (iv) Two major assumptions underlying Eq. (9.12) are

Assumptions 9.1 (*Estimation of heteroscedastic noise variances in absence of model errors*)

- A.** the residuals are uncorrelated, i.e. no systematic errors are present and the measurements are mutually independent, and
- B.** the noise variance is constant in each subset.

In the next section these assumptions will be further discussed.

9.2.3. Model errors present

In order to better satisfy assumption A, we suggest to perform a first order correction on the residuals before using them in the variance estimation. In other words, the best linear model (e.g. a line in 2D, a plane in 3D) is fitted to the residuals and then subtracted from them. This preprocessing step is supposed to reduce any slowly varying systematic error present in the residuals and consequently the diminish

importance of the model chosen. Since additional parameters are estimated, the degrees of freedom contained in matrix \mathbf{T} (Eq. (9.12)) should be adjusted. Formally, the corrected residuals \mathbf{r}_h' of the h^{th} subset are found by subtracting a best linear approximation ($\mathbf{L}_h \hat{\boldsymbol{\alpha}}_h$) from the original residuals \mathbf{r}_h (see Eq. (9.5))

$$\begin{aligned} \mathbf{r}_h' &= \mathbf{r}_h - \mathbf{L}_h \hat{\boldsymbol{\alpha}}_h \\ &= \mathbf{r}_h - \mathbf{L}_h (\mathbf{L}_h^T \mathbf{L}_h)^{-1} \mathbf{L}_h^T \mathbf{r}_h \\ &= \mathbf{P}_h' \mathbf{r}_h = \mathbf{P}_h' (\mathbf{P}_h \mathbf{y}_h - \mathbf{K}_{-h} \mathbf{y}_{-h}) \end{aligned} \quad (9.13)$$

This step is comparable to the creation of the original residuals in Eq. (9.3)-(9.5), except here a model fit ($\mathbf{L}_h \hat{\boldsymbol{\alpha}}_h$) is subtracted from the subset residuals, rather than from the real measurements. Analogous to Eq. (9.6)-(9.10) and considering that \mathbf{P}_h' is a symmetric idempotent matrix, the expected residual variance is now given by

$$E\{\mathbf{r}_h'^T \mathbf{r}_h'\} = \sigma_h^2 \text{trace}[\mathbf{P}_h' \mathbf{P}_h \mathbf{P}_h^T] + \text{trace}[\mathbf{P}_h' \mathbf{K}_{-h} \boldsymbol{\Sigma}_{-h} \mathbf{K}_{-h}^T]. \quad (9.14)$$

The unknown variances σ_h^2 can again be estimated by solving a linear set of equations like Eq. (9.12). In this case, matrices \mathbf{T}_1 and \mathbf{T}_2 also depend on \mathbf{P}_h' . \mathbf{T} can still be interpreted as the total degrees of freedom associated with the residuals (cf. remark (i) in 9.2.2), but its sum will generally not be a natural number anymore and hence its interpretation in terms of number of estimated parameters (i.e. n_θ and n_α) is no longer straightforward.

For this correction step to be efficient, the systematic errors present should be quasi-linear in each subset. Therefore, the choice of the subsets is an important factor in the method (cf. remark (iv)B in section 9.2.2). A possible strategy is to form subsets of equal width and to optimise their number. The optimal number of subsets could be determined e.g. as the smallest number giving white corrected residuals (assessed by some whiteness test, *Hoel* [1971]). But the user is free to apply any other procedure, prior knowledge or experience to choose the subsets. A general rule of thumb is to have around $\text{sqrt}(N)$ objects per

subset with a lower limit of approximately 10. This way, both the number of subsets and the number of objects per subset increases with increasing N .

From the derivations given above, a variance estimation algorithm was distilled. One major aspect of the algorithm not yet mentioned is that the variance estimation should be iterated. In a first step residuals from an unweighted least squares fitting are used to obtain first estimates for the noise variances (as in the derivations above). Next, these variances can be used to estimate the model parameters by WLS fitting, resulting in new residuals which can be used to further refine the variance estimates, etc. The algorithm and some practical details which appear to influence the reliability of the results are described in the next section. Subsequently, the algorithm is applied to two simulations as a validation test (section 9.3) and to the same real world CLIVAR-SR3 data (section 9.4), already presented in the previous chapter.

9.2.4. Practical comments and algorithm

In this section some practical aspects of the variance estimation are discussed, concluding with a schematic presentation of the algorithm. This algorithm was used for all results shown further down.

One of the major problems encountered during our preliminary tests is the occurrence of negative variance estimates, which is obviously unacceptable. Based on our experience, these negative estimates arise when too few data are available. In fact, this behaviour is an illustration of the consistency of the estimator (cf. remark (ii) in 9.2.2). In simulations, when the number of data is increased, the estimates become positive. This is usually not possible in real life, but two other rules also seem to increase the reliability of the results.

(i) The weighted Jacobian should be used, i.e. every row of \mathbf{J} is divided by the uncertainty associated with that measurement

$$\mathbf{J}_w[i, j] = \frac{\mathbf{J}[i, j]}{\sigma[i]}. \quad (9.15)$$

Evidently this is not possible in the first stage, since the weights still have to be estimated. Therefore, an iterative procedure is proposed to improve the estimation stagewise. In order to start with good initial estimates, it is recommended to start with an approximate estimation which is guaranteed to provide nonnegative estimates, e.g. by neglecting the off-diagonal terms in \mathbf{T} (i.e. setting $\mathbf{T}_2 = 0$).

(ii) If the model is fitted to the data, simultaneously obeying some equality constraints, it is important to take this into account. In practice, this is done by correcting the Jacobian matrix for this loss in free parameters. If the following constraints should be satisfied

$$\mathbf{A}_{eq}\boldsymbol{\theta} = \mathbf{b}_{eq} \quad (9.16)$$

(cf. Eq. 8.13), the corrected Jacobian is found by [*Rik Pintelon*, personal communication]

$$\mathbf{J}_c = \mathbf{J}\mathbf{H} \quad \text{with} \quad \mathbf{H} = \mathbf{I}_{n_\theta} - \mathbf{A}_{eq}^T \left(\mathbf{A}_{eq} \mathbf{A}_{eq}^T \right)^{-1} \mathbf{A}_{eq}. \quad (9.17)$$

\mathbf{H} can be computed in a numerically more stable way via the singular value decomposition of \mathbf{A}_{eq} ($\mathbf{A}_{eq} = \mathbf{U}_A \boldsymbol{\Sigma}_A \mathbf{V}_A^T$)

$$\mathbf{H} = \mathbf{I}_{n_\theta} - \mathbf{V}_A \mathbf{V}_A^T. \quad (9.18)$$

\mathbf{J}_c will not be of full rank anymore, its rank now expresses the number of parameters left to be freely estimated.

Combining the theory with these practical considerations, the following algorithm is proposed.

1. *initialisation*

- a. No variance information is known. Compute the Jacobian matrix \mathbf{J} , estimate $\hat{\boldsymbol{\theta}}$ using a Least Squares procedure and calculate the residuals (see Eq. (9.3)).

- b.* If linear constraints must be obeyed, compute \mathbf{J}_c from Eq. (9.17). Otherwise $\mathbf{J}_c = \mathbf{J}$.
- c.* Cut grid into predefined subsets.
- d.* Compute the corrected residuals by subtracting the best linear approximation from the residuals of each subset.
- e.* Estimate $\hat{\boldsymbol{\sigma}}^2$ from Eq. (9.12), using the unweighted Jacobian \mathbf{J}_c , and setting $\mathbf{T}_2 = 0$. The latter will avoid negative variance estimates, which could appear because the unweighted Jacobian was used.

2. iteration: for $i = 1, \dots, n_{iter}$

- a.* Now an initial estimate for the weighting is available. Compute the weighted Jacobian matrix \mathbf{J}_w (Eq. (9.15)), estimate $\hat{\boldsymbol{\theta}}$ using a Weighted Least Squares procedure and calculate the residuals (see Eq. (9.3)).
- b.* If linear constraints must be obeyed, compute \mathbf{J}_{wc} using \mathbf{J}_w in Eq. (9.17). Otherwise $\mathbf{J}_{wc} = \mathbf{J}_w$.
- c.* Cut grid into predefined subsets.
- d.* Compute the corrected residuals by subtracting the best linear approximation from the residuals of each subset.
- e.* Estimate $\hat{\boldsymbol{\sigma}}^2$ from Eq. (9.12), using the weighted Jacobian \mathbf{J}_{wc} , and now allowing $\mathbf{T}_2 \neq 0$, i.e. accounting for coupling effects between the subsets.
- f.* If any element in $\hat{\boldsymbol{\sigma}}^2$ is negative, re-estimate it neglecting the coupling. This can happen for those subsets associated with lowest uncertainty.

9.3. Simulation tests

By dividing the data in subsets, the presented method aims to cope with

- (i) heteroscedasticity of the noise variances, which is closely linked to a
- (ii) heterogeneous distribution of degrees of freedom, and
- (iii) (first order) suppression of model errors.

In the following sections the performance of the procedure in achieving these aims is assessed.

9.3.1. Heterogeneous distribution of degrees of freedom

If it is assumed that the degrees of freedom are spread homogeneously over the whole dataset, the expressions greatly simplify. Each subset variance σ_h^2 can then be estimated by dividing the subset residual variance by the subset degrees of freedom

$$\hat{\sigma}_h^2 = \frac{\mathbf{r}_h^T \mathbf{r}_h}{(N - n_\theta) N_h / N} \quad (9.19)$$

(N is the total number of data, n_θ is the total number of estimated model parameters and N_h is the number of data in subset h). However, if the number of parameters is not much smaller than the number of data, this approximation will provide unsatisfactory results.

Consider for instance the following basic simulation example: a number of data simulated with a sixth order polynomial $y = u^6 - u^2 + 0.2$ (Figure 9.2). The first part of the dataset has a lower uncertainty (standard deviation = 0.01) than the last part (standard deviation = 0.1). A sixth order polynomial is fitted to the data. Next, the data are cut in two subsets and the subset variances are estimated assuming

- (A) homogeneous degrees of freedom using Eq. (9.19) and
- (B) heterogeneous degrees of freedom, using the procedure given in section 9.2.




Figure 9.2. Example of data simulated by a sixth order polynomial (see true function) and two different noise levels (1 realisation).

Table 9.1. Mean estimation results over 1000 Monte Carlo simulations. For each simulation, the noise variances of two subsets of data generated with a sixth order polynomial (see Figure 9.2) were estimated. Four different methods were used: with and without linear correction step, and based on two different assumption on the degrees of freedom. (A) homogeneous and (B) heterogeneous distribution of the degrees of freedom. The results are given as standard deviations since these are more easily interpretable. “df” stands for “degrees of freedom”.

Linear correction step (preprocessing)	Assumption on df	$\hat{\sigma}$, subset 1 ($\sigma_{sim}=0.01$)	$\hat{\sigma}$, subset 2 ($\sigma_{sim}=0.10$)
No	A	0.0227	0.0982
	B	0.0100	0.103
Yes	A	0.0229	0.0982
	B	0.0098	0.102




Figure 9.3. Evolution of estimated standard deviation as a function of iteration number for one arbitrary realisation.

The mean results for 1000 Monte-Carlo simulations are shown in Table 9.1. On the whole, method (B) outperforms the more approximate method (A), especially for the smaller variance. The estimation algorithm seems to converge well (Figure 9.3).

In this example no model errors are present, so performing the linear correction preprocessing is not necessary. This is confirmed in the simulation test: when this step is omitted in the procedure of method (B), the estimated variances are accurate. However, in real situations it is generally not possible to guarantee a priori that there are no systematic errors present. Therefore, it is recommended to include this correction step, at least initially. The effect of including this preprocessing step although no model errors are present is also assessed in this test. It does not appear to influence the estimations. Finally, notice that this example also illustrates the performance of method (B) to reconstruct heteroscedastic noise variances.

9.3.2. Performance of model error correction

In this section the performance concerning the reduction of model errors influence on the noise variance estimates is assessed. To do so, the following simple simulation was performed (see Figure 9.4a): a fourth order polynomial was fitted to data that were generated using a sine function ($y = \sin(u)$). Again, the first part of the dataset was disturbed by lower noise levels (standard deviation = 0.03) than the second part (standard deviation = 0.1). Obviously, this is a situation with model errors, as is confirmed by the systematic shape of the residuals (Figure 9.4b and c).

To test the performance of the first order correction, the grid was cut in ten equally wide subsets and the noise variances were estimated with and without this preprocessing (step 1d and 2d of the algorithm). The mean results of 1000 Monte Carlo simulations are given in 9.2. As expected, the outcomes without correction step are overestimations. With the correction step, the estimates are lowered and thus much closer to the true values. This is also illustrated by the shape of the final corrected residuals (Figure 9.4d). These are more random, except for a few subsets where the first order correction is too approximate. The correction could perhaps be improved if the subsets (size, position, number) were optimised.

Table 9.2. Mean estimation results over 1000 Monte Carlo simulations: noise standard deviations are estimated from uncorrected and corrected residuals to illustrate the efficiency of the correction. All simulation results correspond to the outcome after 10 iterations.

	Mean estimated standard deviation for subset...									
	1	2	3	4	5	6	7	8	9	10
Uncorr. residuals	1.80	0.109	0.0610	0.0445	0.126	0.152	0.178	0.174	0.213	0.465
Corr. residuals	0.0883	0.0318	0.0333	0.0305	0.0332	0.100	0.102	0.103	0.107	0.171
True values	0.03	0.03	0.03	0.03	0.03	0.1	0.1	0.1	0.1	0.1




Figure 9.4. (a) Example of data simulated by a sine function and two different noise levels. In addition the final fourth order polynomial is shown that was fitted to these data using the method with the linear correction steps. (b) Uncorrected residuals after first step (unweighted least squares fitting). (c) Uncorrected residuals after last iteration. (d) Residuals in (c) after linear correction.

9.4. Results on field data

9.4.1. CLIVAR-SR3 2001 data

The aim of this section is to use the procedure described in section 9.2 to estimate the noise variances associated with the CLIVAR-SR3 data, already presented in 8.4.1. These data \mathbf{Y} are multivariate: they count 4 variables (Θ , salinity, NO and PO) at every sampling position. The residuals needed in the procedure will be computed by fitting the POMP model Eq. (8.6) to \mathbf{Y} in a constrained optimisation algorithm. Since linear equality constraints must be satisfied (mixing fractions add up to one in each point), the Jacobian matrix used in the variance estimation will be adapted to take these into account (see 9.2.4). In practice, matrix \mathbf{Y} will be vectorised using the *vec* operator, which puts the columns of a matrix on top of each other (cf. Eq. (8.12)) making the model equations equivalent to Eq. (9.1).

The variance estimation procedure is suitable for this problem, because

- (i) it is difficult to perform replicates which really assess the total reproducibility of the measurements, because this would require replicate samplings of water at different depths and at different locations. These samplings are very time-consuming and expensive.
- (ii) the noise variance is surely heteroscedastic. First, some variables are indeed measured with a much higher precision than others. Second, processes influencing the total reproducibility (e.g. internal waves, interactions with atmosphere) change with location and therefore the total noise variances will also be dependent on the position.

9.4.2. Estimated noise variances

The noise variances are estimated for the same set-up as used in Chapter 8 (different from the one in *de Brauwere et al.* [2007b]): 4 variables are used to reconstruct the fractions of 5 sources (given in Table 8.1). A model with 9 first order splines in the latitude direction and 4 in the depth direction is used, as decided from model selection criteria (cf. 8.4.2). The noise variances for all variables are estimated in 20 subsets (see Figure 9.5).




Figure 9.5. Grid with sampled depths for selected stations of CLIVAR-SR3 cruise. The dashed lines separate the subsets considered in the variance estimation algorithm.

Table 9.3. Ranges of estimated standard deviations, compared to typical analytical uncertainties reported in literature (see references in text). The values between brackets are the average standard deviations.

	Standard deviations from...	
	algorithm	literature
Θ ($^{\circ}\text{C}$)	0.012 – 0.62 (0.29)	0.003
Salinity	0.010 – 0.087 (0.038)	0.002
PO (μM)	3.9 – 23 (11)	9
NO (μM)	3.3 – 21 (11)	8

The results are presented in Table 9.3 as ranges and average values for each hydrographic variable, and compared to values usually reported in literature for analytical uncertainties [Mark Rosenberg, personal communication; You, 2002]. The mean standard deviation values estimated using the algorithm are systematically higher than the analytical uncertainties, although the reported analytical precision is included in the estimated range for PO and NO. This is probably due to two factors. First, the analytical precision is usually estimated neglecting the sampling uncertainty. That is to say, the replicates are sampled at the same occasion. Processes like (internal) waves are known to deteriorate the total reproducibility of the measurements, yet with a single sampling this variability is not taken into account. Especially for temperature and salinity this source of uncertainty is probably significant compared to the strict analytical uncertainty. A second factor that could explain the discrepancy seen in Table 9.3, is incomplete correction for model errors. Indeed, a simplified mixing model is used here and some model errors are likely to be present. If they are not completely corrected, this will result in an overestimation of the noise variances.

At least a part of the model errors is already corrected for, as indicated by the higher values of the estimated variances if no linear correction step is included: mean std(Θ) = 0.34, mean std(salinity) = 0.061, mean std(PO) = 50, mean std(NO) = 12.

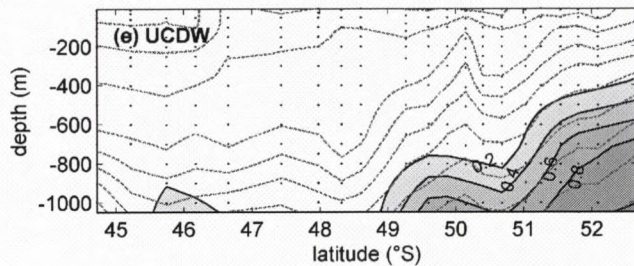




9.4.3. POMP analysis with new noise variance estimates

Having estimated the noise variances, which vary with position (per subset) and per variable, it seems interesting to re-analyse the CLIVAR-SR3 data with the POMP model. Using these new noise variance estimates as weights in the OWLS optimisation, we obtain only slightly different spatial mixing patterns (Figure 9.6, compare with Figure 8.7). However, the residuals (Figure 9.7) are now of the same order of magnitude as the estimated uncertainties, although the first are still larger. Similarly, the OWLS cost function value (1826) has the right order of magnitude (expected value = df = 908). The residuals are thus

still approximately a factor $\sqrt{2}$ too large if they were to be only caused by measurement noise. The remaining variability could be due to source water type variance, which is not taken into account at all, and/or to unmodelled features still present in the data. Remember that even if no model errors are present, the residual cost function is slightly increased because of inequality constraints becoming active (cf. 8.3.4), but this is probably only a minor factor explaining the doubling of the residual cost function.

So, the residuals are still too high, but note that we are able to detect this systematic residual thanks to the estimated noise variances. This confirms the importance of the linear correction step in the variance estimation algorithm. Indeed, without this step, the raw residuals (i.e. including the systematic errors) would have been used to estimate the noise variances, and thus no systematic features could have been detected.

Figure 9.6. Mixing fractions associated with (a) Antarctic Intermediate Water, (b) Subantarctic Mode Water, (c) Subtropical Surface Water, (d) Circumpolar Surface Water and (e) Upper Circumpolar Deep Water, estimated using the POMP method and the CLIVAR-SR3 data. The difference with Figure 8.7 lies in the used weighting for the OWLS cost function: here the weights are inverses of the noise variances, as estimated by the noise variance estimation procedure described in this chapter. Density isolines are plotted as dotted grey lines and range from 26.6 (upper north corner) to 27.6 (lower south corner) with increments of 0.1. Dots represent the sampling depths.



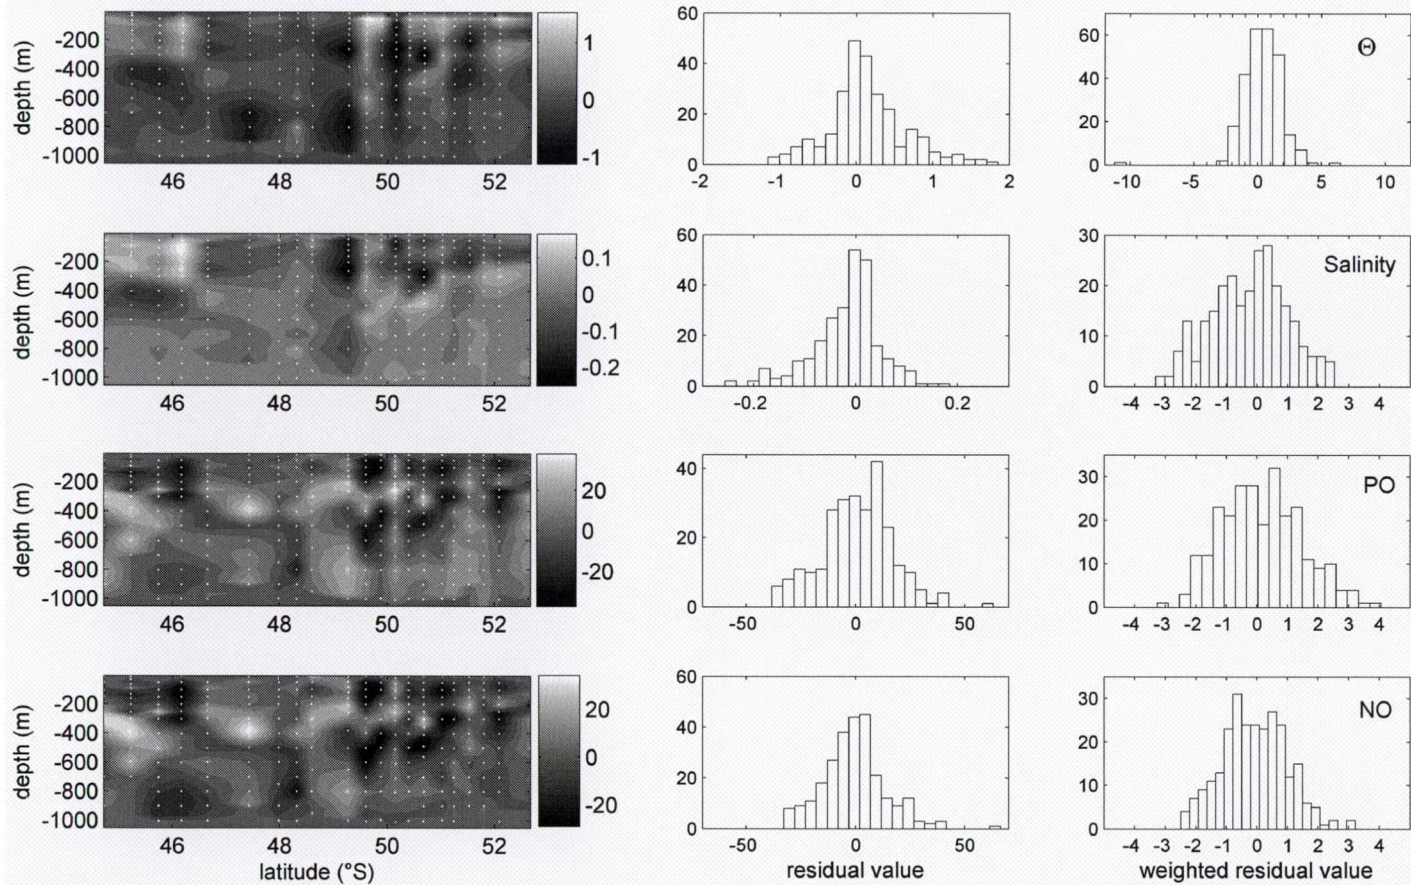


Figure 9.7. Residuals associated with the four variables for the POMP analysis on the CLIVAR-SR3 2001 data and with the improved noise variance estimates for the weighting: potential temperature (first row), salinity (second row), PO (third row) and NO (last row). On the left the raw residuals are represented as a function of position, the central column shows them in histograms (same scale as in Chapter 8), and on the right the weighted residuals are shown, i.e. the raw residuals divided by the priorly estimated standard deviation. Residuals are defined as observation – modelled value.

9.5. Conclusions

A method to estimate heteroscedastic noise variances in the presence of (reasonable) model errors was presented, validated on two simulations and applied to a real dataset. The results show that the variance estimations are accurate even when model errors are present, but can still depend on the subset choice and the number of data available.

An aspect which has not been discussed so far is that the variance estimation procedure also provides an estimate of the subset degrees of freedom. Indeed, each subset variance is equal to the subset cost function (residual sum of squares) divided by the subset degrees of freedom. Due to the subset cutting and to the heterogeneous spreading of uncertainty, these subset degrees of freedom will not be a natural number, nor a divisor of the total number of degrees of freedom.

The estimation of noise variances is often a preprocessing step itself, before a more “essential” calculation. For instance, our primary motivation to find reliable estimates of noise variances was to enable the use of an OWLS cost function for parameter optimisation, model selection and even quality assessment of the results in the POMP application.

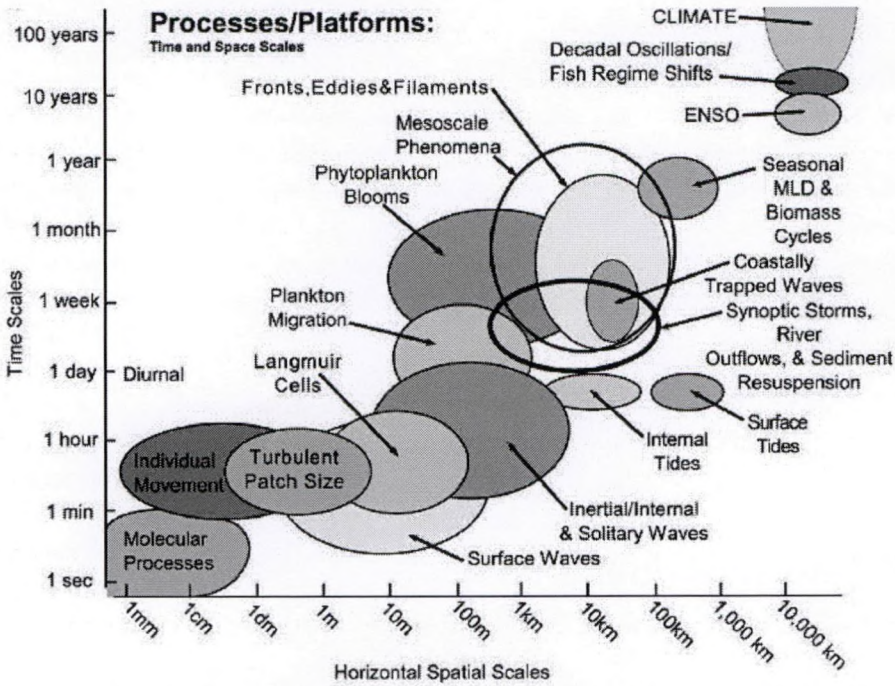


Figure 9.8. Diagram showing the time and horizontal space scale of several physical and biological processes in the ocean. Taken from Dickey [2003].

After these general conclusions, we would like to discuss the POMP application in more detail. First, the *meaning* of the estimated noise variances will be discussed, followed by a further analysis of the results presented in section 9.4. Often, the variances used for the weighting of the cost function are thought to represent the experimental uncertainty associated with the modelled observations. However, in the case of the POMP application, this uncertainty is clearly an underestimation of the real variability related with the observations. Indeed, their value is not only influenced by experimental variations but also by environmental processes which can be treated as random because they occur at a smaller scale than the sampling. This latter effect could be evaluated by performing repeated samplings at the same position. These repeated

observations will be subject to a higher scatter than if they were only perturbed by experimental noise.

This additional error is sometimes called *representation error* and follows from the fact that different scales are present and only some of them can be resolved by the available sampling and modelling. This issue has been extensively discussed in the framework of writing decimal fractions as binary numbers. It is also closely related to rounding and averaging. In the POMP application, a way to view the representation error is as the error due to the fact that one observation is taken to represent a much larger volume of water. In other words, this error is essentially the difference between the point-measurement (taken at a given location and time) and the filtered variable which it is supposed to represent, but is meant to reproduce larger space and time scale variability. In Figure 9.8 the range of variability scales of several oceanographical processes are illustrated, including the meso-scale phenomena, having the scale of interest for this POMP application.

These arguments support the fact that the “noise” variances attributed to the variables in POMP should (or at least could) be larger than the mere experimental precision of the measurements. The variances estimated by the newly proposed algorithm include this type of error as long as its scale is small relative to the sampling density. The fact that these estimated variances are indeed (much) larger than the experimental precisions seem to confirm the importance of representation error as a source of “uncertainty”.

Finally, let us turn to the results themselves and their significance. When the estimated noise variances of the CLIVAR-SR3 2001 data are used for weighting in the POMP model optimisation, the residuals are still significantly higher than the estimated noise levels. At this point it is not possible to determine the cause of this $\sqrt{2}$ discrepancy, but it is probably a combination of the following factors:

- (i) underestimation of the total variance to be expected in the residuals because the SWT uncertainty is ignored,

- (ii) increased residuals due to un- or badly modelled features still in the data, e.g. suboptimal knots positions of the splines, nonconservative processes or additional number of sources, and
- (iii) active inequality constraints will also (slightly) inflate the residuals.

From the simulations in section 8.3.4 it can be assumed that the last factor is probably only of minor importance in explaining the unaccounted residuals. So, the importance of the two remaining factors is at most of the magnitude of half (0.4) the estimated noise level. Further investigation of these factors may be of interest for future work (cf. Chapter 12).

Chapter 10:

Identifiability of the source water matrix

10.1. Introduction

In Chapter 5 the issue of identifiability was already raised in the context of compartmental models. Now it appears that this topic is relevant in the framework of the POMP method as well. One of the advantages of the POMP approach compared to the classical OMP analysis is that the number of unknowns to be estimated is reduced. This increase in degrees of freedom would in theory enable the estimation of additional unknowns, for instance all or part of the elements of the source water matrix \mathbf{S} (cf. remark (ii) in section 8.2.1). This would be an immense improvement as the characterisation of the source water types is still the most subjective part of the method. Purely from the point of view that the number of unknowns must be smaller than (or equal to) the number of observations, this reasoning seems promising. However, it still has to be investigated whether all unknown parameters are identifiable in the current setup. In other words, given the observations \mathbf{Y} and the POMP model (Eq. 8.6), can all elements of \mathbf{S} and \mathbf{C} be uniquely identified?

10.2. Derivation

10.2.1. Method 1: Jacobian rank

Let's recapitulate the situation: the model equations are (cf. Eq. 8.6)

$$\underbrace{\mathbf{Y}}_{n_v \times N} = \underbrace{\mathbf{S}}_{n_v \times n_S} \cdot \underbrace{\mathbf{C}}_{n_S \times n_B} \cdot \underbrace{\mathbf{B}^T}_{n_B \times N},$$

where now both \mathbf{S} and \mathbf{C} are unknowns to be estimated. Identifiability could be verified by inspecting the rank of the Jacobian matrix, as was done in Chapter 5. The Jacobian matrix is defined in Eq. 2.16; translated to the POMP model, this results in

$$\mathbf{J} = \frac{\partial \text{vec}(\mathbf{Y})}{\partial \boldsymbol{\theta}}, \text{ with } \boldsymbol{\theta} = \begin{pmatrix} \text{vec}(\mathbf{S}) \\ \text{vec}(\mathbf{C}) \end{pmatrix}, \quad (10.1)$$

remembering that the *vec* operator stacks the columns of a matrix on top of each other to form a column vector (cf. Appendix). In other words,

$$\mathbf{J} = \left(\begin{array}{c|c} \frac{\partial \text{vec}(\mathbf{Y})}{\partial \text{vec}(\mathbf{S})} & \frac{\partial \text{vec}(\mathbf{Y})}{\partial \text{vec}(\mathbf{C})} \end{array} \right) = \left(\mathbf{B} \otimes \mathbf{S} \middle| \mathbf{B}\mathbf{C}^T \otimes \mathbf{I}_{n_B} \right). \quad (10.2)$$

For a given dataset and model, \mathbf{J} can be computed and hence its rank can be determined. For all elements of \mathbf{S} and \mathbf{C} to be identifiable, the rank of the Jacobian matrix must equal the number of these unknowns minus the number of independent equality constraints imposed by forcing the sum of all mixing fractions at each point to equal one:

$$\text{identifiable} \Leftrightarrow \text{rank}(\mathbf{J}) = n_\theta - n_{eq} = n_v n_S + n_S n_B - n_{eq}. \quad (10.3)$$

The number of independent equality constraints can be determined by the rank of the equality constraint matrix \mathbf{A}_{eq} (cf. Eq. 8.13 and 9.16), defining the equality constraints

$$\mathbf{A}_{eq} \boldsymbol{\theta} = \mathbf{A}_{eq} \begin{pmatrix} \text{vec}(\mathbf{S}) \\ \text{vec}(\mathbf{C}) \end{pmatrix} = \mathbf{b}_{eq}. \quad (10.4)$$

Compared to Eq. (8.14) for the case where only \mathbf{C} was unknown, \mathbf{A}_{eq} will be extended with a block of zeros, because the unknown parameter vector is now enlarged with $\text{vec}(\mathbf{S})$, which does not play a role in the equality constraints. This extension has no influence on the rank of matrix, so the rank of \mathbf{A}_{eq} remains the same as in the situation where only the splines coefficients \mathbf{C} are estimated.

Summarised, Eq. (10.3) reformulates to

$$\text{identifiable} \Leftrightarrow \text{rank}(\mathbf{J}) = n_v n_S + n_S n_B - \text{rank}(\mathbf{A}_{eq}). \quad (10.5)$$

or

$$\text{number of identifiable parameters} = \text{rank}(\mathbf{J}) + \text{rank}(\mathbf{A}_{eq}). \quad (10.6)$$

The drawback of this method to check for identifiability is that the Jacobian and equality constraint matrices have to be computed. A simpler approach is given in the next section.

10.2.2. Method 2

Recall that the model equations are (cf. Eq. 8.6)

$$\underbrace{\mathbf{Y}}_{n_v \times N} = \underbrace{\mathbf{S}}_{n_v \times n_S} \cdot \underbrace{\mathbf{C}}_{n_S \times n_B} \cdot \underbrace{\mathbf{B}^T}_{n_B \times N}, \quad (10.7)$$

where now both \mathbf{S} and \mathbf{C} are unknowns to be estimated. Nothing prevents us to add the following factors in Eq. (10.7)

$$\mathbf{Y} = \mathbf{S} \cdot \mathbf{H}^{-1} \cdot \mathbf{H} \cdot \mathbf{C} \cdot \mathbf{B}^T, \quad (10.8)$$

with \mathbf{H} a regular $n_S \times n_S$ matrix ($\mathbf{H}^{-1} \cdot \mathbf{H}$ is nothing but the identity matrix). This implies that \mathbf{S} and \mathbf{C} cannot be identified closer than to a factor \mathbf{H} . Indeed, different matrices $\mathbf{S}' = \mathbf{S} \cdot \mathbf{H}^{-1}$ and $\mathbf{C}' = \mathbf{H} \cdot \mathbf{C}$ will give equally good solutions as the original \mathbf{S} and \mathbf{C} .

Since the size of \mathbf{H} is $n_S \times n_S$ this indicates that n_S^2 relations exist between the parameters, and thus that all but n_S^2 elements of \mathbf{S} and \mathbf{C} can be identified. However, this is still without taking into account the equality constraints. Due to these constraints \mathbf{H} cannot be chosen completely arbitrarily, but only such that $\mathbf{H} \cdot \mathbf{C}$ satisfies the same equality constraints as \mathbf{C} . So, to determine how many elements of \mathbf{S} and \mathbf{C} are identifiable, we have to find how many elements of \mathbf{H} can still be freely chosen such that the equality constraints (cf. section 8.2.5)

$$\mathbf{1}(1, n_S) \cdot \mathbf{H} \cdot \mathbf{C} \cdot \mathbf{B}^T = \mathbf{1}(1, n_S), \quad (10.9)$$

and

$$\mathbf{1}(1, n_S) \cdot \mathbf{C} \cdot \mathbf{B}^T = \mathbf{1}(1, n_S) \quad (10.10)$$

are satisfied.

Combining both equations one gets

$$\mathbf{1}(1, n_S) \cdot \mathbf{H} \cdot \mathbf{C} \cdot \mathbf{B}^T = \mathbf{1}(1, n_S) \cdot \mathbf{C} \cdot \mathbf{B}^T. \quad (10.11)$$

Applying the *vec* operator to both sides (cf. Appendix for properties), this equation can be reformulated in the form of a classical set of linear equations ($\mathbf{A}\boldsymbol{\theta}=\mathbf{b}$):

$$\left(\mathbf{B}\mathbf{C}^T \otimes \mathbf{1}(1, n_S)\right) \cdot \text{vec}(\mathbf{H}) = \text{vec}\left(\mathbf{1}(1, n_S) \cdot \mathbf{C} \cdot \mathbf{B}^T\right). \quad (10.12)$$

The rank of the first matrix expresses the number of independent relations that \mathbf{H} must satisfy. It can easily be demonstrated that $\text{rank}\left(\mathbf{B}\mathbf{C}^T \otimes \mathbf{1}(1, n_S)\right) = \text{rank}\left(\mathbf{B}\mathbf{C}^T\right) \leq n_S$.

If the splines are well conditioned (which should be the case in all relevant situations), $\text{rank}\left(\mathbf{B}\mathbf{C}^T\right) = n_S$.

Summary

The number of independently identifiable parameters	
= total # unknown parameters	- # relations between parameters
= # elements in \mathbf{S} + # elements in \mathbf{C}	- # free to choose elements in \mathbf{H}
= $n_v n_S + n_S n_B$	- (# elements in \mathbf{H} - # relations involving \mathbf{H})
= $n_v n_S + n_S n_B$	- $n_S^2 + n_S$.

(“#” stands for “number of”)

This result has the advantage to only depend on the a priori (structural) parameters n_v , n_S and n_B . No additional calculation has to be performed, in contrast with the Jacobian rank method.

Interpreting the above results, the following conclusions can be drawn:

- (i) With the current POMP model it is not possible to simultaneously identify all elements of \mathbf{S} and \mathbf{C} ($n_s^2 - n_s \geq 0$ for all $n_s \geq 1$).
- (ii) In the “minimal” experimental setup, i.e. where the minimal number of variables (hydrographic quantities) are measured ($n_v = n_s - 1$), the number of identifiable parameters is $n_s n_B$, this is exactly the number of spline coefficients. So, in this case only \mathbf{C} can be fully identified.
- (iii) If additional variables are used in the analysis, i.e. if $n_v \geq n_s$, some elements of \mathbf{S} can be identified in addition to all elements of \mathbf{C} . It is always good to include additional variables in the analysis because they represent new independent information, but this result represents a supplementary advantage to do so. But no matter how many variables are added, in the current design it will never be possible to simultaneously identify all elements of \mathbf{S} and \mathbf{C} .

10.2.3. Example

As an illustration, the two methods will in this section be applied to a numerical example. The ever-returning CLIVAR-SR3 2001 dataset is taken as a test case. The POMP method is used as described in Chapter 8, i.e. $n_v = 4$, $n_B = 9 \times 4 = 36$ and $n_s = 5$. We can now use the two methods to verify whether it is possible to identify all elements of both \mathbf{S} and \mathbf{C} , i.e. $n_v n_s + n_s n_B = 20 + 180 = 200$ parameters.

Applying method 2, we find that only $n_v n_s + n_s n_B - n_s^2 + n_s = 180$ parameters can be identified, so all elements of \mathbf{C} can be estimated, but then nothing can be said about the source water types.

Applying method 1, from statement (10.5) follows that to verify identifiability the ranks of the \mathbf{J} and \mathbf{A}_{eq} matrices must be computed. This is done by constructing the matrices according to Eq. (10.2) and (8.14) respectively. The ranks of the resulting matrices appear to be: $rank(\mathbf{J}) = 144$ and $rank(\mathbf{A}_{eq}) = 36$. So, according to Eq. (10.6), the number of identifiable parameters = $144 + 36 = 180$, confirming the outcome of method 2.

10.3. Discussion and conclusions

The main conclusion from this chapter is that it is impossible to simultaneously estimate the spline coefficients (i.e. the mixing fractions) and the SWTs in the current POMP design. Only if enough variables are included, some elements of \mathbf{S} may be identifiable besides a full identification of \mathbf{C} . This conclusion is not trivial, because at first sight enough data are available to estimate all unknowns, in contrast with the classical OMP analysis.

The problem of estimating the SWTs is of current interest because a nonlinear optimisation procedure has recently been proposed to estimate the SWTs simultaneously with the mixing fractions in the classical OMP analysis framework [Henry-Edwards and Tomczak, 2006]. As they say, this “*means that the minimisation function becomes non-linear and highly under-determined. Such systems have an infinite number of solutions, and it is necessary to impose additional constraints upon the minimisation in order to achieve a viable result.*” Indeed, they were only able to find “realistic” results because they strongly constrained the possible solutions. For instance, fixed upper and lower limits were defined for each SWT and for the mixing fractions. Moreover, not all SWTs are optimised at once, only those “*that can reasonably be expected to have undergone significant variations*” [Henry-Edwards and Tomczak, 2006]. The main problem with this procedure is that it can hardly be called objective, and in fact is designed to find more or less what is expected.

Nevertheless, this recent development illustrates that one is aware of the importance of a better definition of the SWTs in OMP-like analyses. The choice of the SWTs remains the most subjective component of the whole procedure. In this chapter it has been demonstrated that even within the current POMP setup the SWT problem remains unsolved. More objective source selections can only be achieved if additional information can be included in the analysis, for example if repeated transects are available, preferably from the same year and season.

Chapter 11:

General conclusions

The aim of the study presented herein was to improve two existing environmental modelling problems such that their results are more reliable. The focus was thus on the development and application of suitable estimation algorithms, and no new models or new measurements come out of the study. Indeed, the merit of this work is *not* to have enabled the estimation of flux rates or of mixing fractions, but rather to have enabled their *more accurate and precise* estimation, together with an estimate of the associated uncertainty. In order to achieve this objective, a number of new methodological developments were necessary. The major novel contributions are summarised below.

1. Introduction of IOWLS – LCM combination for the estimation of flux rates and their uncertainty (Chapter 4).

For the first time in the field of tracer experiments it was proposed to include the input uncertainty in the weighting of a WLS cost function. Although the input effect is linearised, the approach seems satisfactory in the case of the simple silicon model. Most importantly, the residual cost function values appear to be very close to their expected values, which enables their use for checking the quality of the results. Combining this input-output weighting with the LCM expression provides a satisfactory method to estimate the uncertainties associated with the flux rates. The estimated variances are again approximations but they can be computed quickly and automatically, thus shortcutting the need for time-consuming Monte-Carlo simulations.

To illustrate the relevance of this section, it may be mentioned here that this estimation procedure has been adopted as the default routine for all other compartmental modelling work performed in the ANCH group.

2. Model selection based on the statistical interpretation of the residual WLS cost function (Chapter 6).

Compared to other existing model selection methods, the strategy proposed herein has the advantage to be very fast and reliable even if few observations are available. Joined with the IOWLS formulation, it is the preferred method for model selection for tracer experiments modelled by compartmental models. However, to increase the insight in a giving modelling exercise, it has recently been suggested to combine the overall quality information given by the residual cost function with a more detailed inspection of the individual residuals [Elskens *et al.*, submitted]. For instance, an acceptable cost function may hide the fact that one variable is systematically associated with a very high residual.

3. Improvement of OMP analysis for large-scale applications (Chapter 8).

As an attempt to improve some of the deficiencies of the classical OMP analysis, the POMP method was proposed. Its main development consists of describing the sources' mixing contributions by 2D-functions. This approach effectively reduces the number of unknowns to be estimated, resulting in more robust mixing patterns. Obviously, this increased robustness should not be paid by a lowered accuracy. Therefore, it is crucial to perform a thorough model selection to determine the optimal complexity of the mixing functions. This could be achieved thanks to the tools presented in Chapter 6. The second development included in POMP regards the more technical aspect of weighting. In order to allow a weighting scheme consistent with the WLS framework, it was necessary to impose mass balance as an equality constraint rather than to optimise it as a model equation. This was easily implemented, and the main difficulty remaining is the accurate quantification of the weights. This last point is important to enable the statistical interpretation of the residuals (or the residual cost function), i.e. to identify possible systematic errors. To achieve this accurate weighting the uncertainties associated with both the observations (output) and the SWTs (input) should be known. As a first step, a method to objectively and accurately estimate the first has been

developed in Chapter 9, but the SWT variances remains beyond the scope of this study.

4. Algorithm to estimate noise variances (Chapter 9).

Originally intended to improve the weighting in POMP applications, an algorithm was designed which can be applied to estimate noise variances in most modelling applications. Using this algorithm in the POMP application enabled position dependent variance estimates and resulted in much higher uncertainties (especially for temperature and salinity) than implied by the instrumental precision. This outcome confirms the suggestion that the uncertainties to include in the weighting should include the sampling variability in addition the mere instrumental precision. However, with these improved estimates of the output variances the weighting still appears to be unsatisfactory, indicating the need to include the SWT variances.

5. Identifiability studies (Chapters 5 and 10).

The issue of identifiability lies slightly apart from the rest of the problems. Indeed, it is not concerned with observations and unavoidably related uncertainties. Instead, it assesses whether a planned model-experiment can provide the requested information under the form estimated parameters. This topic was addressed in both parts of this thesis. In the first part, the methodology using the Jacobian rank as a measure of the amount of available information was extensively described and applied to the complete set of two-compartments models. In Part B, the Jacobian rank was again used, this time to determine how many parameters can be independently determined in a typical POMP application. The results are confirmed by a formal derivation, demonstrating that it is never possible to estimate all splines coefficients simultaneously with all SWT elements. In the typical situation where there is one source more than the number of hydrographic variables, only the spline coefficients can be estimated. It thus remains impossible to objectively define the SWTs, only based on the data.

Chapter 12:

Perspectives for future research

“De quoi demain sera-t-il fait?”

Victor Hugo, *Les chants du crépuscule*, Napoléon II (1835).

12.1. Introduction

As the research is never finished at the end of a PhD, a thesis always seems to contain the obligate chapter on future perspectives. It is no different for this thesis.

The focus will be mainly on work in progress, as I find it difficult to enclose the greatness of (all) possible future developments.

All developments proposed in this chapter are situated within the scope of research priorities of the ANCH and ELEC groups. The collaboration between these two groups may have started as an uncertain experiment, it now seems to have become a fact. More and more ANCH topics are studied using tools from the system identification, sometimes even requiring the development of new specialised tools, as illustrated in several chapters of this thesis. The promising results so far seem to convince at least some people that this kind of interdisciplinary work can be meaningful, as implied by some recent granted projects. On the other hand, it is the communication aspect that appears to be the most troublesome in practice. The introduction of new methods and especially new (according to us more objective) evaluation criteria is not always welcomed. I hope that my work will be a contribution to a better communication between scientific communities and that I have convinced you of its significance. Especially because the work is not yet finished...

12.2. Future developments in system identification?

One of the conclusions of this thesis (and other simultaneous work) is that the branch of system identification considering small datasets is underdeveloped. Asymptotic properties are not relevant anymore in this context, so new fundamental research on the properties of estimators for small datasets is needed. How to proceed is not obvious and it seems a long term project in which the link with the applications must remain active.

One step has been made in this thesis with the proposed model selection procedure (Chapter 6), and its applicability can be further enlarged by combining it with the noise variance estimation of Chapter 9.

12.3. Future developments in modelling tracer experiments?

The tracer models have been further developed since my work in Part A, the main progress being the introduction of small kinetic experiments and related first order and saturation models. The model development is mainly constrained by the complexity of the experiments which must still be feasible to perform e.g. on board of a ship. For instance, the number of samples (and hence reagents) necessary for kinetic experiments is beyond the possibilities during long range cruises as those crossing the Southern Ocean. Therefore, the effort is mostly directed to designing the most efficient experiment to answer the question of interest, taking into account the experimental constraints. This means that today a range of experimental designs exists, all linked to their specific model to answer their specific question. As the main constraint is the experimental complexity, it appears that today the tracer models cannot be much further improved [*Marc Elskens*, personal communication].

Instead, a focus for future research should be the combination of different available experiments to retrieve even more information about the system. This is more specifically of interest in the context of estimating export fluxes of organic material through the water column. Nitrogen experiments provide information about the fraction of the total production that can be potentially exported out of the surface layer (i.e. new production). Silicon observations offer some insight in the contribution of diatoms in this flux. In addition, barium and thorium tracers are currently used to make inferences about the remineralisation of the exported organic matter in deeper layers of the ocean. Finally, sediments and sediment traps contain information about the fraction reaching the ocean floor. Until now these pieces of information are used separately, but it seems promising to combine them to get a bigger

picture of the fate of organic material leaving the upper ocean. This approach would take advantage of existing experimental expertise but would require the construction of a new model combining the different inputs, and thus consists of a real modelling development. Of course, concepts like identifiability and model selection will help to create this model, at least for it to be numerically useful.

12.4. Future developments of POMP analysis?

A number of developments further improving the POMP analysis seem useful. From our experience so far, the major problem for interpreting the results is the residual cost function which is too high. As argued above, at this point it is difficult to say which factor(s) is (are) the primary cause. Two developments may help to solve this problem. First, including the possibility to change the knots' positions of the splines would certainly improve the fit that can be achieved with a spline of a given complexity. Indeed, this would allow the spline to be even more flexible and to have different "resolutions" in different areas of the data grid (now equidistant knots are used, thus the "resolution" of the spline is equal everywhere). More flexible splines would enable a closer fit of the data and hence would reduce the residuals, as wanted. This development is not so much technically complicated, but it is difficult to *automatically* optimise the knots positions. This is actually the reason why it was not implemented in the POMP method so far. So, optimising the knots positions will probably reduce the model errors but at the same time it will introduce a spark of subjectivity.

A second, in my opinion much more tricky, development bringing the residuals closer to the expected value, is including the SWT uncertainties in the estimation of the total variance (thus hopefully achieving the expected residual value). The difficulty here lies on two levels. First, an estimate must be found for these SWT uncertainties. If the definition of the uncertainty associated with the measurements was already subtle (cf. discussion in 9.5), it is even more so for the SWT uncertainties. It also depends on the way the sources themselves are

defined: based only on the investigated dataset (as we do) or as more general, average sources which can be found in literature or by comparing a number of datasets of the same region. It is obvious that some choices have to be made by the user. But as soon as these SWT uncertainties are quantified, the second problem arises: how to incorporate them in an estimate of the total variance, or, which is similar, in a total weighting matrix? It is not so much a problem to include the input variances as long as they are independent from sample to sample (see 2.5.4 and Chapter 4), but since the influence of a source extends over a number of samples this is clearly not the case here. And since the source variances are not independent, covariances should be estimated too. In short, estimation of the SWT variances (and covariances) seems tricky but it would allow the estimation of a *real* total variance, to which the residuals could be compared to detect systematic errors.

Another progress improving the POMP method involves the definition of the source water types. If this definition can be performed more objectively, this would increase the reliability of the POMP results, and it would also help to solve the previously raised problem of quantifying the SWT covariance matrix. It was demonstrated in Chapter 10 that in the current POMP setup it is not possible to estimate the SWTs simultaneously with the spline coefficients. Therefore, this problem can only be solved by including additional information in the analysis. Maybe this information can be found under the form of repeated observations. For instance, in some regions the cruise frequency may be high enough to enable the combination of a number of datasets which are similar in terms of SWTs. It is important to carefully select these repeated datasets since the SWTs can differ according to season and also between years [e.g. *de Brauwere et al.*, 2007a].

A final “development” of the POMP analysis is more of a modification, but it could be useful to improve the fit and interpretability in some applications. It consists of performing the spline parameterisation as a function of other variables than geographical position (latitude and depth). For instance, instead of depth, neutral density [*Jackett and McDougall*, 1997] may be used as the vertical axis. An example of how

the mixing fractions for the CLIVAR-SR3 2001 data and SWTs look like with this axis transformation is shown in Figure 12.1. Alternatively, the latitude axis could be replaced by dynamic height, a variable that would rearrange the stations such that their temperature and salinity values are (more) monotonically ordered. This could have an advantage if eddies are present in the transect. However, so far, these axis transformations do not appear to significantly improve the description of the data, but it seems worthwhile to remember this possibility for future applications.

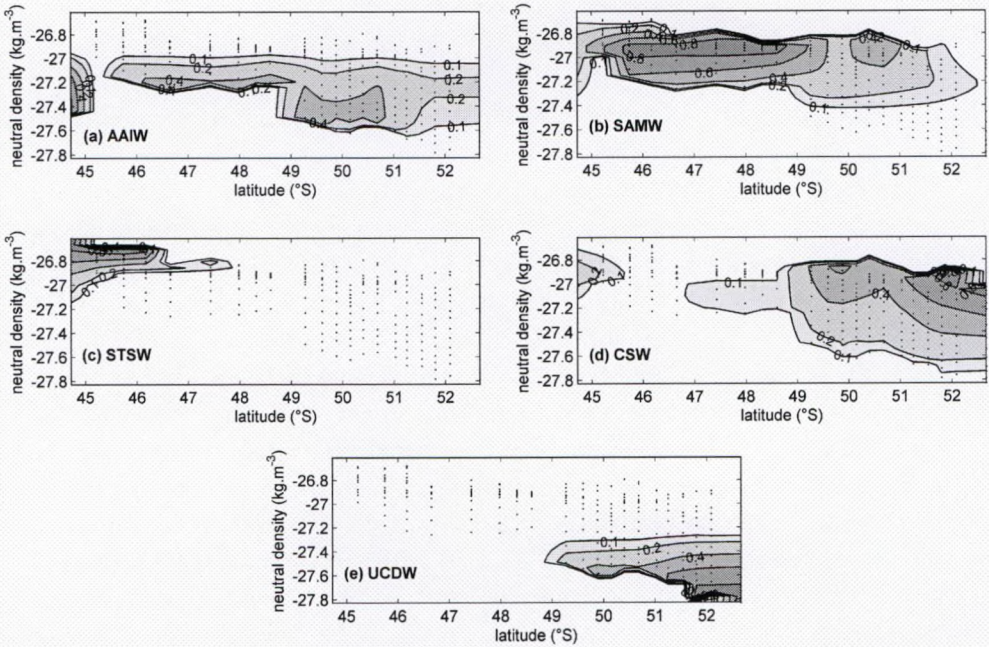


Figure 12.1. Mixing fractions derived for the same data and SWTs as in Chapter 8, but parameterised as function of latitude and neutral density.

12.5. Future application of POMP: detecting biogeochemical processes

12.5.1. Introduction

Although a number of further improvements have been proposed in the previous section, the POMP method is far from worthless as it is now. Its performance in reconstructing mixing fields has been shown in the previous chapters, especially compared with the classical OMP analysis which is used regularly in literature. In some applications these mixing fields may be the product of interest. In this section, however, we propose to use the mixing fractions as an intermediate result and use it in a subsequent calculation to predict the behaviour of dissolved barium. In fact, this application was the incentive of the whole POMP development (cf. Part B).

Why study barium? Barium in the oceans has been suggested to be involved in biogeochemical processes such as the microbial remineralisation of sinking organic matter. Dissolved barium profiles consistently show depletion in the surface layers, implying a consumption of biological nature. As barium is not known to be an essential micronutrient, another mechanism has been suggested to explain the observed profiles. Dissolved barium would be used to form barite crystals within microenvironments present in decaying phytoplankton cells and fecal pellets [Dehairs *et al.*, 1980]. Biological mediation of barite formation is supported by correlations between barite concentrations in intermediate waters and biological productivity in the overlying water column [Dehairs *et al.*, 1980]. In addition, enrichments of dissolved barium are often observed in intermediate to deep waters, pointing toward a regeneration of dissolved barium during microbial remineralisation of the organic aggregates containing barite. This apparent link between productivity and barite formation, and between barite transport to intermediate depths and export production has led to the use of barium as a tracer for past oceanic productivity and export production, crucial parameters in the understanding of the carbon cycle. Nevertheless, the quantitative importance of these biogeochemical processes for the dissolved barium distribution is not well known. Surely its overall distribution in the oceans is primarily

determined by physical processes at all scales, but it is possible that “sometimes” it is significantly influenced by biogeochemical processes as well. When, where and how much barium is subject to these influences is difficult to predict. Therefore, it would be useful to have a tool to determine the location and importance of these nonconservative processes, from a given dataset. A way to do this involves the use of the POMP analysis in a first step. In the next section, the actual method to estimate the importance of mixing on the dissolved barium observations is described.

12.5.2. Method

The basic idea of the method is that dissolved barium in the ocean is composed of a part due to conservative processes and a part influenced by biologically mediated processes. The conservative processes are simplified to be only (large scale) mixing, such that the observed dissolved barium can be decomposed as follows:

$$[\text{Ba}]_{\text{obs}} = [\text{Ba}]_{\text{mix}} + [\text{Ba}]_{\text{bio}}. \quad (12.1)$$

$[\text{Ba}]_{\text{bio}}$ can be quantified if $[\text{Ba}]_{\text{mix}}$ can be estimated. This can be achieved if in the same study area measurements of strictly conservative tracers are available. If this is the case, they can be used to “calibrate” a mixing model (e.g. POMP) which would then fully describe the variations in the study area due to mixing. Using the POMP model, this means that the spatial distribution of the mixing coefficients for the relevant sources is fixed (“calibrated”). With these mixing coefficients, $[\text{Ba}]_{\text{mix}}$ can be reconstructed if barium concentrations can be attributed to the SWTs.

Summarising, a three step procedure is followed to estimate $[\text{Ba}]_{\text{bio}}$:

Step 1: determine the mixing fields using the POMP method with only conservative tracers.

Step 2: Compute the “mixing-only” barium at each point:

$$[\text{Ba}]_{\text{mix}} = \sum_{i=1}^{n_s} x_i [\text{Ba}]_i, \quad (12.2)$$

with x_i the mixing fraction of SWT i (derived in step 1) and $[\text{Ba}]_i$ the barium content of SWT i .

Step 3: Compute the biogeochemical barium component:

$$[\text{Ba}]_{\text{bio}} = [\text{Ba}]_{\text{obs}} - [\text{Ba}]_{\text{mix}}. \quad (12.3)$$

The main weakness of this procedure is that the uncertainty associated with the final values for $[\text{Ba}]_{\text{bio}}$ is relatively high, because it is computed as a difference of two uncertain terms of comparable magnitude. Note that if the “mixing-only” barium was computed using mixing fractions derived from the classical OMP, the uncertainties would be even higher, and probably too high to make any statements, which is one of the justifications for executing the work presented in Part B. Even using the POMP, the resulting $[\text{Ba}]_{\text{bio}}$ distributions are mostly interesting as qualitative information, i.e. to indicate subregions associated with nonconservative consumption ($[\text{Ba}]_{\text{bio}}$ is negative) or rather with production ($[\text{Ba}]_{\text{bio}}$ is positive). This information can then be related to independent indications from other data, e.g. particulate barium, to make inferences about the role of barium in the biological cycling of organic material.

12.5.3. Results and possible future applications

This method has already been applied to two Southern Ocean datasets, among which the ever revisited CLIVAR-SR3 2001 dataset. This work was done in cooperation with Stéphanie Jacquet and has been materialised in a manuscript submitted to Geophysical Research Letters [Jacquet *et al.*, submitted]. In Figure 12.2 an example for the CLIVAR-SR3 2001 section is shown. A region of consistent production and a region of consumption becomes apparent, which can now be further interpreted, e.g. by comparing with other available data. For instance, in the subantarctic zone where dissolved barium seems to be depleted, elevated levels of particulate barium have been reported, which is

consistent with a subtraction of dissolved barium in ambient water during barite formation [*Jacquet et al.*, submitted].

Similar studies, using a classical OMP model, were executed by *Kumar and Li* [1995] to evaluate the amount of Si and ^{226}Ra added to the water column by decomposing particles, and by *Schneider et al.* [2005] to approximate changes in DIC and nutrients due to remineralisation. The advantage of using the POMP model instead of the classical OMP equations is that the mixing fractions estimated first are associated with (much) lower uncertainties, as was shown in Chapter 8. This aspect is particularly important in this application because the mixing fractions are only an intermediate result, and thus their uncertainty is further propagated through the subsequent steps.

In analogy with *Schneider et al.* [2005] the proposed procedure can be repeated for nutrients instead of barium, provided that enough conservative tracers are available (since NO and PO cannot be used if NO_3 and PO_4 are analysed). Quantifying the biogeochemical component of nutrients has the advantage that it can be linked to primary production. An interesting variation of the procedure would therefore be to investigate the upper layer of a particular ocean region, and thus to perform the POMP analysis in the (latitude, longitude) dimensions. From the biogeochemical nutrient component, the local production can be estimated, and by integrating this value over a representative volume of water around the sample, large scale maps of production can be derived. In many regions, local production estimates are available, which would allow a cross-check of the derived production estimates. The advantages of this three-step procedure compared to an “extended OMP analysis” (cf. section 7.4) are

- (i) a reduction of the noise influence due to the parameterisation, and
- (ii) no assumption about Redfield ratios must be made.

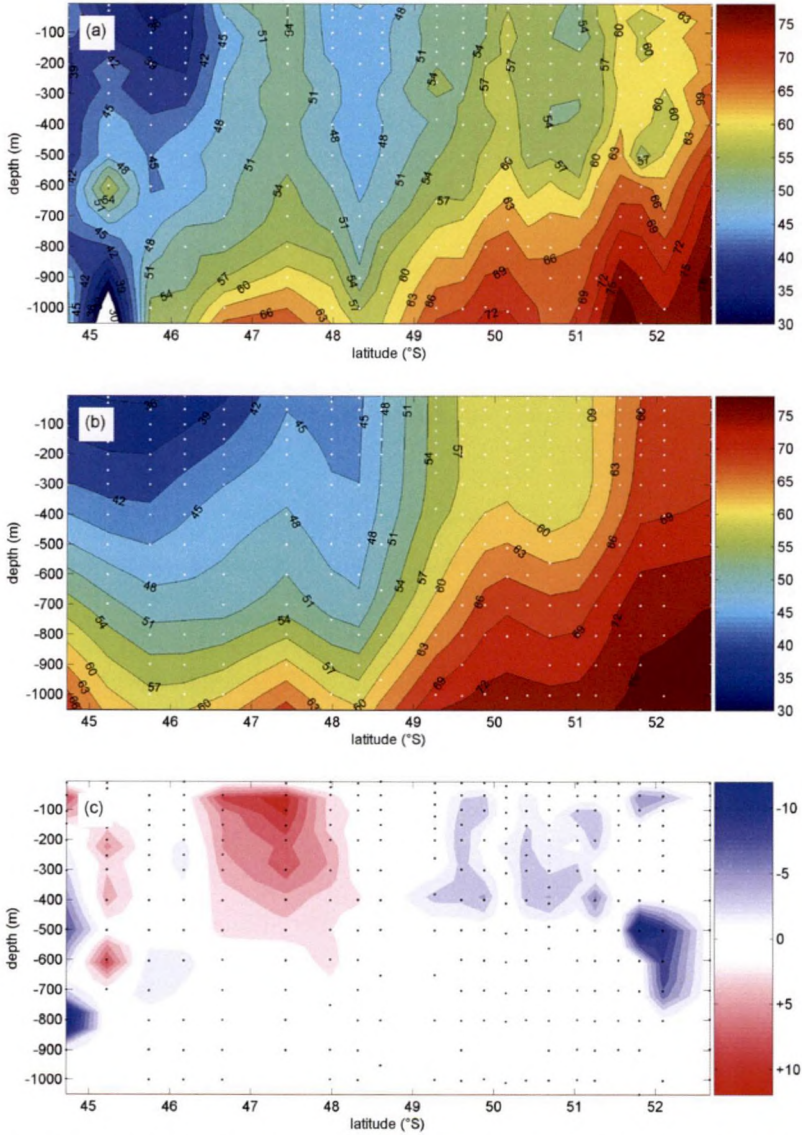


Figure 12.2. (a) Measured barium profile (in nM) along the CLIVAR-SR3 2001 section. (b) Example of reconstructed “mixing-only” barium profile. (c) Difference between (a) and (b), giving an estimate of excess and missing barium. Here only those values are shown which are above the estimated uncertainty (1 standard deviation).

12.6. Future application of POMP: inclusion of more tracers

A final suggestion for future application of POMP concerns the inclusion of more conservative tracers than we did so far, i.e. potential temperature, salinity, NO and PO. This is useful since more tracers will increase the reliability of the results. Besides, if one wants to use NO₃ and PO₄ to make statements about biological production (vide supra), other tracers will be needed too. In literature, concentrations of inert compounds such as CFCs have been used in OMP analyses [e.g. *Olsson et al.*, 2005]. Measurement of CFCs have been made during several cruises, including the CLIVAR-SR3 2001 campaign. At the same time, some new tracers may appear as useful conservative variables. For instance, aluminium is currently investigated as a potential tracer for water masses [*Clare Johnston and Toby Sherwin*, personal communication]. Following these developments and checking which additional tracers are already available now, is something that must certainly be done in future applications of POMP.

Appendix

Matrix algebra

Consider the following matrices

$$\mathbf{A} \in \mathbb{R}^{p \times q}, \mathbf{B} \in \mathbb{R}^{p \times q}, \mathbf{C} \in \mathbb{R}^{q \times p}, \mathbf{D} \in \mathbb{R}^{s \times t}, \mathbf{E} \in \mathbb{R}^{q \times s}.$$

A1. Trace operator

$$\text{Definition: } \text{trace}[\mathbf{A}] = \sum_{i=1}^{\min(p,q)} \mathbf{A}_{(i,i)} \quad (\text{A.1})$$

$$\text{Properties: } \text{trace}[\mathbf{A} + \mathbf{B}] = \text{trace}[\mathbf{A}] + \text{trace}[\mathbf{B}] \quad (\text{A.2})$$

$$\text{trace}[\mathbf{AC}] = \text{trace}[\mathbf{CA}] \quad (\text{A.3})$$

A2. Vec operator

$$\text{Definition: } \text{vec}(\mathbf{A}) = \begin{pmatrix} \mathbf{A}_{(:,1)} \\ \mathbf{A}_{(:,2)} \\ \vdots \\ \mathbf{A}_{(:,q)} \end{pmatrix} \in \mathbb{R}^{pq \times 1} \quad (\text{A.4})$$

A3. Kronecker product

$$\text{Definition: } \mathbf{A} \otimes \mathbf{D} = \begin{pmatrix} \mathbf{A}_{(1,1)}\mathbf{D} & \mathbf{A}_{(1,2)}\mathbf{D} & \cdots & \mathbf{A}_{(1,q)}\mathbf{D} \\ \mathbf{A}_{(2,1)}\mathbf{D} & \mathbf{A}_{(2,2)}\mathbf{D} & \cdots & \mathbf{A}_{(2,q)}\mathbf{D} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{(p,1)}\mathbf{D} & \mathbf{A}_{(p,2)}\mathbf{D} & \cdots & \mathbf{A}_{(p,q)}\mathbf{D} \end{pmatrix} \in \mathbb{R}^{ps \times qt} \quad (\text{A.5})$$

Properties involving the *vec* operator:

$$\text{vec}(\mathbf{AED}) = (\mathbf{D}^T \otimes \mathbf{A}) \text{vec}(\mathbf{E}) \quad (\text{A.6})$$

$$\text{vec}(\mathbf{AE}) = (\mathbf{I}_s \otimes \mathbf{A}) \text{vec}(\mathbf{E}) \quad (\text{A.7})$$

$$= (\mathbf{E}^T \otimes \mathbf{I}_p) \text{vec}(\mathbf{A}) \quad (\text{A.8})$$

$$= (\mathbf{E}^T \otimes \mathbf{A}) \text{vec}(\mathbf{I}_q) \quad (\text{A.8})$$

Reference List

- Akaike H. (1974), New look at statistical-model identification, *IEEE Trans. Autom. Control*, *AC19*(6), 716-723.
- Analytical Methods Committee (M. Thompson, W.H. Evans, M.J. Gardner, E.J. Greenhow, R. Howarth, E.J. Newman, B.D. Ripley, K. Swan, R. Wood and J.J. Wilson) (1992), Proficiency testing of analytical laboratories: organization and statistical assessment, *Analyst*, *117*(1), 97-104.
- Andersson J. H., J. J. Middelburg and K. Soetaert (2006), Identifiability and uncertainty analysis of bio-irrigation rates, *J. Mar. Res.*, *64*, 407-429.
- Aoki S. N. L. Bindoff and J. A. Church (2005a), Interdecadal water mass changes in the Southern Ocean between 30°E and 160°E, *Geophys. Res. Lett.*, *32*, L07607, doi:10.1029/2004GL022220.
- Bard Y. (1974), *Nonlinear Parameter Estimation*, Academic Press, New York.
- Beck M. B. (1987), Water-quality modeling - a review of the analysis of uncertainty, *Water Resour. Res.*, *23*(8), 1393-1442.
- Beck M. B. (1999), Coping with ever larger problems, models, and data bases, *Water Sci. Technol.*, *39*(4), 1-11.
- Beucher C., P. Tréguer, R. Corvaisier, A. M. Hapette and M. Elskens (2004a), Production and dissolution of biosilica, and changing microphytoplankton dominance in the Bay of Brest (France), *Mar. Ecol. Prog. Ser.* *267*, 57-69.
- Beucher C., P. Tréguer, R. Corvaisier, A. M. Hapette, J.-J. Pichon and N. Metzl (2004b), Intense summer biosilica recycling in the Southern Ocean, *Geophys. Res. Lett.*, *31*, doi: 10.1029/2003GL.018998.
- Box M. J. (1970), Improved parameter estimation, *Technometrics*, *12*(2), 219-229.
- Bozdogan H. and D. M. A. Haughton (1998), Informational complexity criteria for regression models, *Comput. Stat. Data Anal.*, *28*(1), 51-76.

- Brewer J. W. (1978), Kronecker products and matrix calculus in System Theory, *IEEE Trans. Circuits and Systems*, CAS-25(9), 772-781.
- Broecker W. S. (1974), "NO", a conservative water-mass tracer, *Earth Planet. Sci. Lett.*, 23(1), 100-107.
- Browne M. W. (2000), Cross-validation methods, *J. Math. Psychol.*, 44(1), 108-132.
- Brun R., P. Reichert and H. R. Künch (2001), Practical identifiability analysis of large environmental simulation models, *Water Resour. Res.*, 37(4), 1015-1030.
- Brzezinski M. A., D. M. Nelson, V. M. Franck and D. E. Sigmon (2001), Silicon dynamics within an intense open-ocean diatom bloom in the Pacific sector of the Southern Ocean, *Deep-Sea Res II*, 48, 3997-4018.
- Capron X., B. Walczak, O. E. de Noord and D. L. Massart (2005), A modification of the ICOMP criterion for estimation of optimum complexity of PCR models, *J. Chemometr.*, 19(5-7), 308-316.
- Carroll R. J. and D. Ruppert (1982), A comparison between maximum-likelihood and generalized least-squares in a heteroscedastic linear-model, *J. Am. Stat. Assoc.*, 77(380), 878-882.
- Carroll R. J. and D. B. H. Cline (1988), An asymptotic theory for weighted least-squares with weights estimated by replication, *Biometrika*, 75(1), 35-43.
- Castro C. G., F. F. Perez, S. E. Holley and A. F. Rios (1998), Chemical characterisation and modelling of water masses in the Northeast Atlantic, *Prog. Oceanogr.*, 41(3), 249-279.
- Chaigneau A. and R. Morrow (2002), Surface temperature and salinity variations between Tasmania and Antarctica, 1993-1999, *J. Geophys. Res.-Oceans*, 107, (C12), 8020, doi:10.1029/2001JC000808.
- Chandler J. P. (1972), Iterative procedure for estimating functions when both variables are subject to error, *Technometrics* 14(1), 71-76.
- Chatfield C. (1995), Model uncertainty, data mining and statistical inference, *J. R. Stat. Soc. Ser. A*, 158(3), 419-466.
- Chen J. H. and J. Shao (1993), Iterative weighted least-squares estimators, *Ann. Stat.*, 21(2), 1071-1092.
- Clutton-Brock M. (1967), Likelihood distributions for estimating functions when both variables are subject to error, *Technometrics* 9(2), 261-269.

- Coatanoan C., N. Metzl, M. Fieux and B. Coste (1999), Seasonal water mass distribution in the Indonesian throughflow entering the Indian Ocean, *J. Geophys. Res.-Oceans*, 104(C9), 20801-20826.
- Cobelli C., A. Lepschy and G. Romanin Jacur (1979), Identification experiments and identifiability criteria for compartmental systems, *Compartmental analysis of ecosystem models*, 10, 99-123.
- Cobelli C. and J. J. DiStefano (1980), Parameter and structural identifiability concepts and ambiguities - a critical-review and analysis, *Am. J. Physiol.*, 239(1), R7-R24.
- Cobelli C., A. Lepschy and G. R. Jacur (1980), Identifiability in tracer experiments, *Federation Proceedings*, 39(1), 91-96.
- Cobelli C. and A. Caumo (1998), Using what is accessible to measure that which is not: Necessity of model of system, *Metab.-Clin. Exp.*, 47(8), 1009-1035.
- Davidian M. and R. J. Carroll (1987), Variance function estimation, *J. Am. Stat. Assoc.*, 82(400), 1079-1091.
- Davidian M. and P. D. Haaland (1990), Regression and calibration with nonconstant error variance, *Chemometrics Intell. Lab. Syst.*, 9(3), 231-248.
- de Boer C. J. and H. M. van Aken (1995), A study of objective methods for water mass analysis, applied to the Iceland Basin, *Deutsche Hydrographische Zeitschrift*, 47, 5-22.
- de Brauwere A., F. De Ridder, R. Pintelon, M. Elskens, J. Schoukens and W. Baeyens (2005a), Model selection through a statistical analysis of the minimum of a weighted least squares cost function, *Chemometrics Intell. Lab. Syst.*, 76(2), 163-173.
- de Brauwere A., F. De Ridder, M. Elskens, J. Schoukens, R. Pintelon and W. Baeyens (2005b), Refined parameter and uncertainty estimation when both variables are subject to error. Case study: Estimation of Si consumption and regeneration rates in a marine environment, *J. Mar. Syst.*, 55(3-4), 205-221.
- de Brauwere A., S. H. M. Jacquet, F. De Ridder, F. Dehairs, R. Pintelon, J. Schoukens and W. Baeyens (2007a), Water mass distributions in the Southern Ocean derived from a parametric analysis of mixing water masses, *J. Geophys. Res.* 112, C02021, doi:10.1029/2006JC003742.

- de Brauwere A., R. Pintelon, F. De Ridder, J. Schoukens and W. Baeyens (2007b), Estimation of heteroscedastic measurement noise variances, *Chemometrics and Intelligent Laboratory Systems*, 86, 130-138.
- Dehairs F., R. Chesselet and J. Jedwab (1980), Discrete suspended particles of barite and the barium cycle in the open ocean, *Earth Planet. Sci. Lett.*, 49(2), 528-550.
- De La Rocha C. L., M. A. Brzezinski and M. J. De Niro (1997), Fractionation of silicon isotopes by marine diatoms during biogenic silica formation, *Geochemica and Cosmochimica Acta*, 61(23), 5051-5056.
- Deleersnijder E., B. Tartinville and J. Rancher (1997), A simple model of the tracer flux from the Mururoa Lagoon to the Pacific, *Appl. Math. Lett.*, 10(5), 13-17.
- Deleersnijder E., J. Wang and C. N. K. Mooers (1998), A two-compartment model for understanding the simulated three-dimensional circulation in Prince William Sound, Alaska, *Continental Shelf Research*, 18, 279-287.
- de Luna X. and K. Skouras (2003), Choosing a model selection strategy, *Scand. J. Stat.*, 30(1), 113-128.
- De Ridder F., R. Pintelon, J. Schoukens and D. P. Gillikin (2005), Modified AIC and MDL model selection criteria for short data records, *IEEE Trans. Instr. Meas.*, 54(1), 144-150.
- Dickey T. D. (2003), Emerging ocean observations for interdisciplinary data assimilation systems, *J. Mar. Syst.*, 40-41, 5-48.
- Dierckx P. (1995), *Curve and surface fitting with splines*, Oxford University Press, Oxford.
- Donaldson J. R. and R. B. Schnabel (1987), Computational experience with confidence-regions and confidence-intervals for nonlinear least-squares, *Technometrics*, 29(1), 67-82.
- Dugdale R. C. and J. J. Goering (1967), Uptake of new and regenerated forms of nitrogen in primary productivity, *Limnol. Oceanogr.*, 12, 196-206.
- Elskens M., W. Baeyens and L. Goeyens (1997), Contribution of nitrate to the uptake of nitrogen by phytoplankton in an ocean margin environment, *Hydrobiologia*, 353, 139-152.

- Elskens M. (1999), Modelling new and regenerated production from ^{15}N -tracer experiments: a case study of the European continental/Ocean margin in the North Atlantic. PhD in Science, Vrije Universiteit Brussel. 122 p.
- Elskens M., W. Baeyens, T. Cattaldo, F. Dehairs and B. Griffiths (2002), N uptake conditions during summer in the subantarctic and polar frontal zones of the Australian sector of the southern ocean, *J. Geophys. Res.-Oceans*, 107(C11), Art. No. 3182.
- Elskens M., W. Baeyens, N. Brion, S. De Galan, L. Goeyens and A. de Brauwere (2005), Reliability of N flux rates estimated from ^{15}N enrichment and dilution experiments in aquatic systems, *Global Biogeochemical Cycles*, 19, GB4028, doi:10.1029/2004GB002332.
- Elskens, M., A. de Brauwere, C. Beucher, R. Corvaisier, N. Savoye, P. Tréguer and W. Baeyens, Statistical process control in assessing production and dissolution rates of biogenic silica in marine environments, *Marine Chemistry*, in press.
- Fogelqvist E., J. Blindheim, T. Tanhua, S. Osterhus, E. Buch and F. Rey (2003), Greenland-Scotland overflow studied by hydro-chemical multivariate analysis, *Deep-Sea Res. Part I-Oceanogr. Res. Pap.*, 50(1), 73-102.
- Forster M. R. (2001), The new science of simplicity, *Simplicity, Inference and Modelling*, 83-119.
- Friedman E., D. W. Massaro, S. N. Kitzis and M. M. Cohen (1995), A comparison of learning models, *J. Math. Psychol.*, 39, 164-178.
- Glibert P. M., F. Lipschultz, J. J. McCarthy and M. A. Altabet (1982), Isotope-dilution models of uptake and remineralization of ammonium by marine plankton, *Limnol. Oceanogr.*, 27(4), 639-650.
- Glibert P. M. and D. G. Capone (1993), *Mineralization and assimilation in aquatic, sediment, and wetland systems*, in: Nitrogen Isotope Techniques, Academic Press, San Diego, 243-272.
- Grünwald P. (2000), Model selection based on minimum description length, *J. Math. Psychol.*, 44(1), 133-152.
- Harrison W. G. (1983), *Use of isotopes*, in: Nitrogen in the marine environment, Academic Press, New York, 763-807.
- Henry-Edwards A. and M. Tomczak (2006), Remote detection of water property changes from a time series of oceanographic data, *Ocean Sci.*, 2, 11-18.

- Hoel P. G. (1971), Introduction to mathematical statistics, Wiley, New York.
- Jackett D. R. and T. J. McDougall (1997), A neutral density variable for the world's oceans, *J. Phys. Oceanogr.*, 27(2), 237-263.
- Jacquet S. H. M., F. Dehairs and S. Rintoul (2004), A high resolution transect of dissolved barium in the Southern Ocean, *Geophys. Res. Lett.*, 31, L14301, doi:10.1029/2004GL020016.
- Jacquet S. H. M., A. de Brauwere, F. De Ridder and F. Dehairs, An investigation of dissolved barium behavior in the Southern Ocean using a parametric mixing model, submitted to *Geophys. Res. Lett.*
- Jacquez J. A. and M. Norusis (1973), Sampling experiments on the estimation of parameters in heteroscedastic linear regression, *Biometrics*, 29, 771-779.
- Kadane J. B. and N. A. Lazar (2004), Methods and criteria for model selection, *J. Am. Stat. Assoc.*, 99(465), 279-290.
- Karstensen J. and M. Tomczak (1998), Age determination of mixed water masses using CFC and oxygen data, *J. Geophys. Res.-Oceans*, 103(C9), 18599-18609.
- Karstensen J. and M. Tomczak (1999), Manual of the OMP Analysis Package for MATLAB Version 2.0, *OMP User Group*, http://www.ldeo.columbia.edu/~jkarsten/omp_std/README.html.
- Kendall M. and A. Stuart (1979), *The advanced theory of statistics*, Charles Griffin & Company Limited, London.
- Kieseppä I. A. (2003), AIC and large samples, *Philos. Sci.*, 70(5), 1265-1276.
- Li B., J. Morris and E. B. Martin (2002), Model selection for partial least squares regression, *Chemometrics Intell. Lab. Syst.*, 64(1), 79-89.
- Limpert E., W. A. Stahel and M. Abbt (2001), Log-normal distributions across the sciences: keys and clues, *BioScience*, 51(5), 341-352.
- Lo Monaco C., N. Metzl, A. Poisson, C. Brunet and B. Schauer (2005), Anthropogenic CO₂ in the Southern Ocean: Distribution and inventory at the Indian-Atlantic boundary (world ocean circulation experiment line I6), *J. Geophys. Res.-Oceans*, 110, C06010, doi:10.1029/2004JC002643.
- Lukacs E. (1975), *Stochastic convergence*, Academic Press, New York.

- Mackas D. L., K. L. Denman and A. F. Bennett (1987), Least-squares multiple tracer analysis of water mass composition, *J. Geophys. Res.-Oceans*, 92(C3), 2907-2918.
- Mallows, C. L. (1973), Some comments on C_p , *Technometrics*, 15(4), 661-675.
- Maltrud M. E., R. D. Smith, A. J. Semtner and R. C. Malone (1998), Global eddy-resolving ocean simulations driven by 1985-1995 atmospheric winds, *J. Geophys. Res.-Oceans*, 103(C13), 30825-30853.
- Mandel J. (1984), Fitting straight-lines when both variables are subject to error, *J. Qual. Technol.* 16(1), 1-14.
- Mantel N. (1970), Why stepdown procedures in variable selection, *Technometrics*, 12(3), 621-625.
- Massart D. L., B. G. Vandeginste, L. M. Buydens et al. (1997), *Handbook of Chemometrics and Qualimetrics*, Elsevier, Amsterdam.
- MATLAB, version 7 (2005), The Mathworks Inc., Massachusetts.
- Meinrath G., C. Ekberg, A. Landgren and J. O. Liljenzin (2000), Assessment of uncertainty in parameter evaluation and prediction, *Talanta* 51(2), 231-246.
- Minasny B. and A. B. McBratney (2002), Uncertainty analysis for pedotransfer functions, *Eur. J. Soil Sci.*, 53(3), 417-429.
- Myung I. J. (2000), The importance of complexity in model selection, *J. Math. Psychol.*, 44(1), 190-204.
- Nihoul J. C. J. (1975), *Modelling of marine systems*, Elsevier Oceanography Series, Volume 10, Elsevier, Amsterdam.
- Olsson K. A., E. Jeansson. T. Tanhua and J.-C. Gascard (2005), The East Greenland Current studied with CFCs and released sulfur hexafluoride, *J. Mar. Syst.* 55, 77-95
- Perez F. F., L. Mintrop, O. Llinas, M. Glez-Davila, C. G. Castro, M. Alvarez, A. Kortzinger, M. Santana-Casiano, M. J. Rueda and A. F. Rios (2001), Mixing analysis of nutrients, oxygen and inorganic carbon in the Canary islands region, *J. Mar. Syst.* 28(3-4), 183-201.
- Pintelon R., J. Schoukens and G. Vandersteen (1997), Model selection through a statistical analysis of the global minimum of a weighted nonlinear least squares cost function, *IEEE Trans. Signal Process.*, 45(3), 686-693.

- Pintelon R. and J. Schoukens (2001), *System Identification: A Frequency Domain Approach*, IEEE, New York.
- Pitt M. A., W. Kim and I. J. Myung (2003), Flexibility versus generalizability in model selection, *Psychon. Bull. Rev.*, 10(1), 29-44.
- Powell D. R. and J. R. Macdonald (1972), Rapidly convergent iterative method for solution of generalized nonlinear least-squares problem, *Comput. J.* 15(2), 148-155.
- Rao C. R. (1970), Estimation of heteroscedastic variances in linear models, *J. Amer. Stat. Ass.*, 65, 161-172.
- Reichert, P. and M. Omlin (1997), On the usefulness of overparameterized ecological models, *Ecol. Model.*, 95(2-3), 289-299.
- Rintoul S. R. and J. L. Bullister (1999), A late winter hydrographic section from Tasmania to Antarctica, *Deep-Sea Res. Part I-Oceanogr. Res. Pap.*, 46(8), 1417-1454.
- Rintoul S. R. and S. Sokolov (2001), Baroclinic transport variability of the Antarctic Circumpolar Current south of Australia (woce repeat section SR3), *J. Geophys. Res.-Oceans*, 106(C2), 2815-2832.
- Rintoul S. R., J. R. Donguy and D. H. Roemmich (1997), Seasonal evolution of upper ocean thermal structure between Tasmania and Antarctica, *Deep-Sea Res. Part I-Oceanogr. Res. Pap.*, 44, (7):1185-1202.
- Rissanen J. (1978), Modeling by shortest data description, *Automatica*, 14(5), 465-471.
- Rod V. and V. Hancil (1980), Iterative estimation of model parameters when measurements of all variables are subject to error, *Comput. Chem. Eng.*, 4(1), 33-38.
- Rosenberg M., S. Rintoul, S. Bray, C. Curran and N. Johnston. Aurora Australis Marine Science Cruise AU0103, CLIVAR-SR3 Transect - Oceanographic Field Measurements and Analysis. ACE-CRC, Hobart, unpublished report.
- Saccomani M. P., S. Audoly, L. D'Angio, R. Sattier and C. Cobelli (1994), PRIDE: a program to test a priori identifiability of linear compartmental models, *SYSID 1994, 10th IFAC Symposium on System Identification*, 3, 25-30.
- Schoukens J., Y. Rolain and R. Pintelon (2002), Modified AIC rule for model selection in combination with prior estimated noise models, *Automatica*, 38(5), 903-906.

- Schwarz G. (1978), Estimating the dimension of a model, *Ann. Stat.*, 6(2), 461-464.
- Shao J. (1993), Linear model selection by cross-validation, *J. Am. Stat. Assoc.*, 88(422), 486-494.
- Siedler G., J. Church and J. Gould (2001), *Ocean Circulation and Climate*, Academic Press, San Diego.
- Skoog D. A., D. M. West, F. J. Holler and S. R. Crouch (2000), *Analytical Chemistry. An introduction*. Seventh edition, Harcourt College Publishers.
- Soetaert K. and P. Herman (2001), *Ecological modelling. Lecture notes*. NIOO-CEME.
- Sokolov S. and S. R. Rintoul (2002), Structure of Southern Ocean fronts at 140°E, *J. Mar. Syst.*, 37(1-3), 151-184.
- Sokolov S. and S. R. Rintoul (2003), Subsurface structure of interannual temperature anomalies in the Australian sector of the Southern Ocean, *J. Geophys. Res.-Oceans*, 108(C9), 3285, doi:10.1029/2002JC001494.
- Stone M. (1974), Cross-validators choice and assessment of statistical predictions, *J. R. Stat. Soc. Ser. B-Methodol.*, 36(2), 111-147.
- Stone M. (1977), Asymptotic equivalence of choice of model by cross-validation and Akaike's criterion, *J. R. Stat. Soc. Ser. B-Methodol.*, 39(1), 44-47.
- Thompson R. O. R. Y. and R. J. Edwards (1981), Mixing and water-mass formation in the Australian Sub-Antarctic, *J. Phys. Oceanogr.*, 11(10), 1399-1406.
- Tomczak M. (1981), A multi-parameter extension of temperature/salinity diagram techniques for the analysis of non-isopycnal mixing, *Prog. Oceanog.*, 10, 147-171.
- Tomczak M. and D. G. B. Large (1989), Optimum multiparameter analysis of mixing in the thermocline of the Eastern Indian Ocean, *J. Geophys. Res.-Oceans*, 94(C11), 16141-16149.
- Tomczak M. (1999), Some historical, theoretical and applied aspects of quantitative water mass analysis, *J. Mar. Res.*, 57(2), 275-303.
- Valsami G., A. Iliadis and P. Macheras (2000), Non-linear regression analysis with errors in both variables: Estimation of co-operative binding parameters, *Biopharm. Drug Dispos.* 21(1), 7-14.

- Vanlanduit S. (2001), High spatial resolution experimental modal analysis, PhD thesis, 304 pp., Vrije Universiteit Brussel, Brussels.
- Varah J. M. (1990), Relative sizes of the Hessian terms in nonlinear parameter-estimation. *SIAM J. Sci. Stat. Comput.* 11(1), 174-179.
- Wagenmakers, E. J. and S. Farrell (2004), AIC model selection using Akaike weights, *Psychon. Bull. Rev.*, 11(1), 192-196.
- Wisnowski, J. W., J. R. Simpson, D. C. Montgomery and G. C. Runger (2003), Resampling methods for variable selection in robust regression, *Comput. Stat. Data Anal.*, 43, 341-355.
- You Y. and M. Tomczak (1993), Thermocline circulation and ventilation in the Indian-Ocean derived from water mass analysis, *Deep-Sea Res. Part I-Oceanogr. Res. Pap.*, 40(1), 13-56.
- You Y. Z. (1997), Seasonal variations of thermocline circulation and ventilation in the Indian Ocean, *J. Geophys. Res.-Oceans*, 102(C5), 10391-10422.
- You Y. (2002), Quantitative estimate of Antarctic Intermediate Water contributions from the Drake Passage and the Southwest Indian Ocean to the South Atlantic, *J. Geophys. Res.*, 107, doi:10.1029/2001JC000880.
- You Y., J. R. M. Lutjeharms, O. Boebel and W. P. M. de Ruijter (2003), Quantification of the interocean exchange of intermediate water masses around southern Africa, *Deep-Sea Res. Part II*, 50(1), 197-228.



Vrije Universiteit Brussel

PhD thesis
Anouk Debrauwere
April 2007