

Nonlinear model evaluation: ι -shadowing, probabilistic prediction and weather forecasting



Isla Gilmour
Exeter College
University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas Term 1998

Abstract

Physical processes are often modelled using nonlinear dynamical systems. If such models are relevant then they should be capable of demonstrating behaviour observed in the physical process. In this thesis a new measure of model optimality is introduced: the distribution of ι -shadowing times defines the durations over which there exists a model trajectory consistent with the observations. By recognising the uncertainty present in *every* observation, including the initial condition, ι -shadowing distinguishes model sensitivity from model error; a perfect model will always be accepted as optimal. The traditional root mean square measure may confuse sensitivity and error, and rank an imperfect model over a perfect one. In a perfect model scenario a good variational assimilation technique will yield an ι -shadowing trajectory but this is not the case given an imperfect model; the inability of the model to ι -shadow provides information on model error, facilitating the definition of an alternative assimilation technique and enabling model improvement.

While the ι -shadowing time of a model defines a limit of predictability, it does not validate the model as a predictor. Ensemble forecasting provides the preferred approach for evaluating the uncertainty in predictions, yet questions remain as to how best to construct ensembles. The formation of ensembles is contrasted in perfect and imperfect model scenarios in systems ranging from the analytically tractable to the Earth's atmosphere, thereby addressing the question of whether the apparent simplicity often observed in very high-dimensional weather models fails 'even in or only in' low-dimensional chaotic systems. Simple tests of the consistency between constrained ensembles and their methods of formulation are proposed and illustrated. Specifically, the commonly held belief that initial uncertainties in the state of the atmosphere of realistic amplitude behave linearly for two days is tested in operational numerical weather prediction models and found wanting: nonlinear effects are often important on time scales of 24 hours. Through the kind consideration of the European Centre for Medium-range Weather Forecasting, the modifications suggested by this are tested in an operational model.

Acknowledgements

This thesis would not have come to fruition without the assistance and encouragement of a number of people. In particular I would like to thank my supervisor, Lenny Smith, for providing guidance, encouragement and enthusiasm over the past three years. Special thanks for advocating my attendance at a variety of conferences and workshops; the interaction with so many kind and encouraging people, in often glorious surroundings, meant that I always returned a little more enlightened and with renewed enthusiasm. Thanks also to Andrew Fowler who, along with Guy Kember, gave me an unforgettable and enjoyable introduction to research, and who has continued to be supportive.

The NWP results presented here could not have been obtained without the assistance of Tim Palmer and Zoltan Toth in discussion of the operational ensembles and in obtaining the operational data. The kind consideration of the European Centre for Medium-range Weather Forecasting and the assistance and unlimited patience of Roberto Buizza enabled the construction and evaluation of 24h SV ensembles, for me the ‘icing on the cake’.

I gratefully acknowledge the assistance of Christine Ziehmann, Jan Barkmeijer, Martin Ehrendorfer, Alastair Mees and Jurgen Kurths in discussions; of Peter Read in providing data from the annulus; and of Wesley Ebisuzaki with implementing code to read NWP data. The assistance of my examiners, Dave Broomhead and Myles Allen, in the thorough discussion during my ‘viva’ is also gratefully acknowledged.

In addition, I am grateful to the Engineering and Physical Sciences Research Council for a D. Phil. studentship and spring school funding; to the Amelia Jackson memorial fund, the Alan Tayler memorial fund, Exeter College, the University of Oxford, Data Connection and the Mathematical Institute for (numerous!) travel grants; and to the European Union, the ‘Enrico Fermi’ School of Physics and the Newton Institute for grants.

Finally, I would like to thank my parents for their endless support; those who have taught me and facilitated my enjoyment of mathematics over the years; and all my friends and colleagues throughout my time in Oxford, with special thanks to Clare Harraway, Katie Kelsey and Taryn Malcolm for their enduring friendship.

Contents

1	Introduction	1
2	Reconstructing dynamical systems	7
2.1	Introduction to dynamical systems	8
2.1.1	Flows and maps	8
2.1.2	Onset of chaos	10
2.1.3	Attractors, Lyapunov exponents and chaos	12
2.2	Modelling data	15
2.2.1	Building a model	15
2.2.1.1	Polynomial approximations	15
2.2.1.2	Radial Basis Functions	16
2.2.1.3	Models using physical principles	17
2.2.2	Determining the optimal order of a model	18
2.2.2.1	Exact data	20
2.2.2.2	Noisy data	23
2.3	Other systems: lasers, annuli and the atmosphere	26
2.3.1	Reconstructing a state space	26
2.3.2	Analytical systems	28
2.3.2.1	Logistic map	28
2.3.2.2	Ikeda map	29
2.3.2.3	Sinai map	30
2.3.2.4	Lorenz equations	31
2.3.2.5	Asymmetric Marzec-Spiegel system	32
2.3.3	Laboratory systems: the thermally driven rotating annulus . .	33
2.3.4	Real world systems: the Earth's atmosphere	35
2.4	Summary	36

3	Evaluating erroneous models from uncertain data	38
3.1	Complications of chaos	39
3.1.1	Anosov-Bowen ϵ -shadowing	40
3.2	Uncertainty of observation	42
3.2.1	Uncertainty in the data	42
3.2.1.1	Quantisation uncertainty	42
3.2.1.2	Gaussian uncertainty	43
3.2.1.3	Non-uniqueness of observed trajectories	43
3.2.2	Uncertainty of the model	43
3.3	Quantifying model optimality	44
3.3.1	Discrepancies of r.m.s.	45
3.3.2	ι -shadowing	46
3.4	Demonstration of the ι -shadowing algorithm	49
3.5	Summary	54
4	Evaluating models: enlightening shadows	56
4.1	Quantifying model optimality	56
4.1.1	Logistic map	57
4.1.2	Moran-Ricker map	61
4.1.3	Lorenz equations	62
4.2	ι -shadowing failure, model error and system sensitivity	66
4.3	Improving a RBF model of the annulus	70
4.4	From evaluation to prediction	73
4.4.1	Limitations of ι -shadowing	73
4.4.2	ι -shadowing as an assimilation technique	74
4.5	Finite improvements: the need for ensembles	75
4.6	Summary	76
5	Probabilistic predictions: ensemble construction and performance evaluation	78
5.1	Ensemble construction	79
5.1.1	Unconstrained ensembles	80
5.1.2	Perfect ensembles	87
5.1.3	Constrained vector subspaces	89
5.1.3.1	Singular vectors	89
5.1.3.2	Lyapunov vectors	92
5.1.3.3	Dream perturbations	93

5.1.3.4	Breeding vectors	94
5.2	Performance evaluation	96
5.2.1	CV ensemble evaluation measures	97
5.2.1.1	Evaluation of SV ensembles	100
5.2.1.2	Evaluation of BV ensembles	100
5.2.1.3	Other measures of evaluation	101
5.2.2	Marzec-Spiegel system ensemble evaluation	102
5.2.2.1	Similarity between CV bases	102
5.2.2.2	Spread	103
5.2.2.3	Projection of dream perturbations onto CV subspaces	106
5.2.2.4	Projection of CV onto local manifold	109
5.2.2.5	Summary of CV performance	109
5.2.3	Annulus ensemble evaluation	110
5.2.3.1	Similarity between CV bases	110
5.2.3.2	Spread	111
5.2.3.3	Projection of dream perturbations onto CV subspaces	113
5.2.3.4	Projection of CV onto local manifold	115
5.2.3.5	Summary of CV performance	117
5.3	Summary	117
6	Testing the internal consistency of Ensemble Prediction Schemes	120
6.1	Review of literature on the validity of linear approximations in numerical weather models	121
6.2	Quantifying the relevance of the linear assumption	124
6.3	Linearity evaluation of the Marzec-Spiegel system	127
6.3.1	Evaluation of linearity measures	127
6.3.2	Doubling times	130
6.3.3	Summary of linearity evaluation measures	131
6.4	Linearity evaluation of the annulus	131
6.4.1	Evaluation of linearity measures	132
6.4.2	Doubling times	134
6.4.3	Summary of linearity evaluation measures	136
6.5	Details of ensemble employment in NWP	136
6.6	Linearity evaluation of operational NWP ensembles	139
6.6.1	Evaluation of linearity measures	139
6.6.2	Doubling times	145

6.6.3	Summary of linearity evaluation measures	149
6.7	Comparing 48h SV and 24h SV ensembles	151
6.7.1	Evaluation of linearity measures for one day	152
6.7.2	Amplification factors for one day	156
6.7.3	Evaluation of linearity measures, averaged over 7 days	160
6.8	Summary	162
7	Conclusions	165
A	Anosov-Bowen ϵ-shadowing	171
B	ι-shadowing: the algorithm	173
B.1	Method of interval convergence	174
B.2	Random search method	174
B.3	Localised linear feedback	176
B.3.1	Maps	178
B.3.2	Flows	179
B.3.3	Inversion	180
C	Singular Value Decomposition and its applications	182
C.1	SVD: a matrix decomposition	182
C.2	Singular systems analysis	183
	Bibliography	I

List of Figures

2.1	A trajectory of the Lorenz system.	8
2.2	Non-invertibility of the Moran-Ricker map.	9
2.3	Second iterates of the Moran-Ricker map, illustrating period-doubling.	10
2.4	Bifurcation diagram for the Moran-Ricker map.	11
2.5	An approximation to the invariant measure of the Moran-Ricker map.	13
2.6	Predictions of Gaussian RBF models of the Moran-Ricker map constructed using exact data and up to 12 centres.	21
2.7	Singular values for a Gaussian RBF model, with 12 centres, of the Moran-Ricker map.	22
2.8	Errors for Gaussian RBF models of the Moran-Ricker map, constructed using exact data and up to 12 centres.	22
2.9	Singular values for a Gaussian RBF model, with 64 centres, of the Moran-Ricker map.	24
2.10	Predictions of Gaussian RBF models of the Moran-Ricker map, constructed using up to 64 centres and both clean and noisy data.	24
2.11	Errors for Gaussian RBF models of the Moran-Ricker map, comparing models constructed using exact and noisy data.	25
2.12	Logistic map.	28
2.13	Invariant measure of the logistic map.	29
2.14	Attractor of the Ikeda map.	30
2.15	Attractor of the Sinai map	31
2.16	A section of a Marzec-Spiegel attractor.	33
2.17	The thermally driven rotating annulus as a laboratory analogy to the mid-latitude circulation of the Earth's atmosphere.	34
3.1	Limitations of finite precision computations.	40
3.2	ϵ -shadowing a pseudo-trajectory.	41
3.3	Comparing r.m.s. and RBF models of the Ikeda map.	45

3.4	Typical trajectories for the saddle system, showing stable and unstable directions.	50
3.5	Illustration of localised linear feedback implemented for the saddle system.	51
3.6	A succession of trial initial conditions with increasing ι -shadowing times for the saddle system.	53
4.1	ι -shadowing results for the logistic map using direct linear interpolation models with equally spaced base points.	58
4.2	ι -shadowing results for the logistic map using linear interpolation models, comparing direct and iterative models.	59
4.3	ι -shadowing results for the logistic map using linear interpolation models, comparing distributions of base points.	60
4.4	ι -shadowing results for the logistic map comparing direct and iterative, linear and quadratic models.	60
4.5	ι -shadowing results for the Moran-Ricker map, comparing direct and iterative, linear and quadratic models.	61
4.6	ι -shadowing results for the Lorenz equations comparing the random search and combined methods for finding the optimal initial condition.	63
4.7	ι -shadowing results for an imperfect model of the Lorenz equations.	64
4.8	ι -shadowing results using the Lorenz equations to model observations generated with stochastic dynamic noise.	65
4.9	Distribution of the final points of the model to ι -shadow observations from the Lorenz equations.	67
4.10	Distribution of the error doubling times on the Lorenz attractor.	67
4.11	Comparing model error, one and five step observed prediction errors and the error of shadowing failure points as a function of y' for the Sinai map.	68
4.12	Regions of large error for the Sinai map, for different errors.	70
4.13	Comparison of ι -shadowing time distributions for RBF models of the annulus.	72
5.1	Evolution of an unconstrained ensemble, for one time step, under a perfect Marzec-Spiegel model, with associated final time p.d.f.	81
5.2	Comparison of unconstrained and perfect ensembles evolved under a perfect Marzec-Spiegel model.	82

5.3	Evolution of unconstrained ensembles evolved under a RBF model of the annulus.	84
5.4	Comparing system observations, ι -shadowing trajectory and ensemble predictions for the annulus.	86
5.5	Schematic illustration to explain the physical meaning of singular vectors and Lyapunov vectors.	90
5.6	Schematic illustration to show some of the assumptions made, in a perfect model scenario, in using constrained vectors defined using a model as approximations to those of the system.	92
5.7	Imperfect model scenario schematic illustration of dream perturbations and breeding vectors.	94
5.8	Schematics of breeding: alternative definitions of breeding vectors. . .	95
5.9	Illustration showing constrained vectors, analysis, observation, system value and local manifold points for the Marzec-Spiegel system.	98
5.10	Histogram comparing ratios of constrained vector spread to real spread using a perfect model of the Marzec-Spiegel system.	103
5.11	Histograms comparing constrained vector spread to real spread for singular vector and breeding vector perturbations of varying magnitudes, using a perfect model of the Marzec-Spiegel system.	104
5.12	Scatterplot of constrained vector spread against real spread, using a perfect model of the Marzec-Spiegel system.	105
5.13	Scatterplot of projection indices of dream perturbations onto 1d constrained vector subspaces, using a perfect model of the Marzec-Spiegel system.	107
5.14	Scatterplot of projection indices of dream perturbations onto 2d constrained vector subspaces, using a perfect model of the Marzec-Spiegel system.	108
5.15	Histogram of the projection indices of constrained vector subspaces onto the leading local manifold directions, using a perfect model of the Marzec-Spiegel system.	109
5.16	Histograms of the ratios of constrained vector spread to real spread, comparing both constrained vectors and vector magnitude, using a RBF model of the annulus.	112
5.17	Scatterplot of the ratios of constrained vector spread to real spread using a RBF model of the annulus.	112

5.18	Scatterplot of projection indices of dream perturbations onto 3d constrained vector subspaces using a RBF model of the annulus.	113
5.19	Scatterplot of the projection indices of dream perturbations onto 1d constrained vector subspaces using a RBF model of the annulus. . . .	114
5.20	Scatterplot of projection indices of dream perturbations onto subspaces for method A and method B breeding vectors, using a RBF model of the annulus.	116
5.21	Histogram of the projection indices of constrained vector subspace onto leading local manifold directions, using a RBF model of the annulus.	116
6.1	Schematic explaining the definition of Θ	125
6.2	Linearity results of SV for the Marzec-Spiegel system.	128
6.3	Linearity results of individual SV for the Marzec-Spiegel system. . . .	128
6.4	Linearity results of BV for the Marzec-Spiegel system.	129
6.5	q -pling times of SV perturbations for a perfect model of the Marzec-Spiegel system.	130
6.6	q -pling times of BV perturbations for a perfect model of the Marzec-Spiegel system.	131
6.7	The fraction of initial conditions for which SV perturbations exceed a given linearity threshold as a function of time for a RBF model of the annulus.	132
6.8	Linearity results of SV for a RBF model of the annulus.	133
6.9	Linearity results of BV perturbations, generated using method A, for a RBF model of the annulus.	134
6.10	Linearity results of individual SV for a RBF model of the annulus. . .	135
6.11	q -pling times of SV perturbations for a RBF model of the annulus. . .	135
6.12	q -pling times of BV perturbations for a RBF model of the annulus. .	136
6.13	ECMWF ‘stamp map’ showing 500hPa geopotential height fields of an evolved SV ensemble	140
6.14	Linearity results for ECMWF 48h SV twin perturbations for 500hPa geopotential height over the Northern Hemisphere, averaged over 25 days.	141
6.15	Linearity results for NCEP BV twin perturbations for 500hPa geopotential height over the Northern Hemisphere, averaged over 25 days. .	142
6.16	Linearity results for individual ECMWF 48h SV twin perturbations for 500hPa geopotential height over the Northern Hemisphere for 19.xii.96.	144

6.17	Comparison of the degree of nonlinearity of a 36h ECMWF SV ensemble prediction with the corresponding synoptic structures of the control forecast.	146
6.18	Comparison of linearity results for NCEP BV twin perturbations for 500hPa geopotential height, averaged over 25 days, calculated over different regions of the world.	147
6.19	Comparison of linearity results calculated over the Northern Hemisphere and averaged over 25 days, using different norms.	148
6.20	q -pling times of ECMWF SV perturbations.	149
6.21	q -pling times of NCEP BV perturbations.	150
6.22	Singular value spectra for ECMWF 48h SV and 24h SV for 1.i.98. . .	152
6.23	Average total energy fields of ECMWF 48h SV and 24h SV ensembles constructed for 1.i.98.	153
6.24	Linearity results for ECMWF twin (rotated and non-rotated) 48h SV perturbations calculated in the total energy norm over the Northern Hemisphere for 1.i.98, and evolved at different resolutions.	154
6.25	Linearity results for ECMWF twin (rotated and non-rotated) 24h SV perturbations calculated in the total energy norm over the Northern Hemisphere for 1.i.98, and evolved at different resolutions.	155
6.26	Linearity results for ECMWF twin (rotated and non-rotated) 48h SV and 24h SV perturbations calculated in the 500hPa geopotential height norm over the Northern Hemisphere for 1.i.98.	157
6.27	Amplification of ECMWF (rotated and non-rotated) first 48h SV and 24h SV perturbations with respect to singular values, calculated in the 500hPa geopotential height norm over the Northern Hemisphere for 1.i.98.	158
6.28	Linearity results for ECMWF twin (rotated) 48h SV and increased amplitude 24h SV perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere and averaged over 7 days.	161
B.1	Illustration of the random search method for finding the ι -shadowing trajectory.	175
B.2	Illustration of the linear feedback method for finding the ι -shadowing trajectory.	177

Nomenclature

The following is a list of most of the symbols used within this thesis, with multiple uses listed. The list is alphabetical with the following precedences: Roman before Greek, lower case before upper case, scalars before vectors and symbols without sub- or super-scripts before those with.

\mathbf{a}_e	Approximated analysis error
$a_j, j = 1, \dots, k$	Parameters used to define a k^{th} order model
appcos_k	k term Taylor series expansion of cosine
\mathbf{A}_t	Analysis, or best guess of the system state, at time t
α	Name of model of Ikeda map; also constant used in ι -shadowing algorithm
BV	Breeding vector
β	Name of model of Ikeda map; also constant used in ι -shadowing algorithm
$\mathbf{c}_j, j = 1, \dots, n_c$	n_c centres of RBF model
CV, (CV)	Constrained vector (perturbation)
CV_{max}	Maximum number of CV used, <i>i.e.</i> maximum dimension of CV subspace
γ_n	Error in observation, with respect to system value, at time $n\tau_s$
d	Attractor dimension
δ	Analysis or observation error
Δt	Time step
ΔT	Temperature difference between cylinders in annulus
Δ_n	Evolved perturbation at time $t = n$; also model error at time $n\tau_s$
e	One step prediction error
\hat{e}	Average one step prediction error
$E_{\mathbf{x}}^s$	Stable subspace at \mathbf{x}
$E_{\mathbf{x}}^u$	Unstable subspace at \mathbf{x}
ϵ_c	Machine precision
ϵ_n	Model error with respect to observation at time $n\tau_s$
ϵ	Perturbation, also model error with respect to observation
ϵ_{CV}	Magnitude of initial CV perturbation
$f^{(t)}, \mathbf{f}^{(t)}$	Function corresponding to evolution under the system (for time t)
$\tilde{f}^{(t)}, \tilde{\mathbf{f}}^{(t)}$	Function corresponding to a model (for time t)
ζ	Quantisational observational uncertainty (<i>i.e.</i> a measurement is $\mathbf{x} \pm \zeta$)
g	Acceleration due to gravity taken as a constant, g_{38}
g_ϕ	Acceleration due to gravity at latitude ϕ
$\boldsymbol{\eta}$	Vector of distances from centre to surface of hyper-cuboid

\mathcal{H}	Hessian
Θ	Linearity measure; also region defined in ι -shadowing algorithm
\mathcal{J}	Jacobian
ℓ	(Anti-)correlation
L	log likelihood function
LV (LV_∞)	Finite time (or global) Lyapunov vector
Λ	(First) Lyapunov exponent
m	Dimension of system
$\mathcal{M}(\mathbf{x}, \Delta t)$	Linear propagator of \mathbf{x} over time Δt
\mathbf{M}_t	Model representation of the system state at time t
n	Number of base points used in defining polynomial interpolation model, also embedding dimension
n_c	Number of RBF centres
n_e	Number of ensemble members
N	Number of data points
$\mathcal{N}_a(0, 0.01)$	Normal, or Gaussian, distribution with mean 0 and standard deviation 0.01 in each of a dimensions
ξ	Constant used in ι -shadowing algorithm
Ξ_t	Neighbourhood of observation at time t
\mathbf{p}	Rotated SV perturbation
P	Projection matrix
P_n	Polynomial of degree n
\mathcal{P}	Probability
$\mathbf{P} = \{\mathbf{p}_t\}_{t=a}^b$	Pseudo-orbit for time $t = a$ to time $t = b$
\mathbf{P}_δ	δ pseudo-orbit
$r, (r_c)$	(Critical value of) system parameter
r_{lin}	Radius of the linear range
$\mathcal{R}^{(m)}$	Set of (m -dimensional) real numbers
R_n	Scaling parameter used in defining ECMWF SV ensembles
s_j	Image associated with centre j in RBF models
SV, (SSV)	(System) Singular vector
$\mathbf{S} = \{\mathbf{s}_t\}_{t=a}^b$	True (system) orbit for time $t = a$ to time $t = b$
\mathbf{S}_t	System representation at time t
Σ	Diagonal matrix of singular values
σ_i	i^{th} singular value
t	Time

t_c	Critical time at which r.m.s. model switches to mean system value
$t = a \downarrow b$	Time $t = a$, with initiation (of ensemble) at $t = b$
T	Temperature
$\mathcal{T}_{\mathbf{x}}$	Linear tangent space at \mathbf{x}
TaLb	Numerical weather model of spectral resolution a with b vertical levels of resolution, using a Gaussian grid
T _l aLb	Numerical weather model of spectral resolution a with b vertical levels of resolution, using a linear grid
τ_a	Time for which trajectory shadows to give analysis value
$\tau_c, \tau_{\text{BVC}}$	Breeding vector cycle time
τ_d	Delay time of embedding
τ_e	Estimated time for machine precision to grow to magnitude of observational uncertainty
τ_{BVinit}	Pre-breeding time, or initialisation time, for breeding vectors
$\tau_{\text{opt}} (\tau_{\text{LVopt}}, \tau_{\text{SVopt}})$	Optimisation time (specifically of Lyapunov or singular vectors)
τ_{q^n}	Time taken for a perturbation to grow by a factor of q^n
τ_s	Time interval between samples
$\tau_l^{(k)}$	l -shadowing time (of k^{th} ensemble member)
uv500	500hPa horizontal wind velocity
$U(U^T)$	Left singular vector matrix (transposed)
$\mathcal{U}_a(0.01)$	Uniform distribution over an interval of 0.01 in each of a dimensions
$V(V^T)$	Right singular vector matrix (transposed)
ϕ	Latitude
$\phi(r)$	Radial basis function
Φ	Geopotential
x, \mathbf{x}	Point or model representation of system state
\hat{x}	Approximation to system state
$(\mathbf{x})_l$	l^{th} component of \mathbf{x}
\mathbf{x}_0	Initial condition
$\{\mathbf{x}_0^k\}, k = 1, \dots, n_e$	Ensemble of initial conditions with n_e members
x^*	Steady state or fixed point
\mathbf{y}_t	Observation at time t
Ω	Rotation rate
z500	500hPa geopotential height
$\mathcal{Z}_{(0)}^+$	Set of positive integers (including zero)

Truth forever on the scaffold,
wrong forever on the throne -
Yet that scaffold sways the future, and,
behind the dim unknown,
Standeth God within the shadow,
keeping watch above his own.

James Lowell, 1844.

Chapter 1

Introduction

Predicting physical systems is the aspiration of many; predictions about physical systems is of interest to the masses, as demonstrated by the publishing of sunrise and sunset times alongside weather forecasts in many newspapers. While the times of sunrise and sunset are accepted as fact, weather forecasts are often viewed with some scepticism arising from inaccuracies in spite of plentiful endeavour¹. In this thesis we consider the evaluation of models and their predictions of physical systems from uncertain observations. Much research has been done into the construction of both models and ensemble prediction schemes used to produce probabilistic predictions. It is not the purpose of this thesis to discuss or contribute to this work directly, but rather to evaluate and compare the performances of competing models and of different ensemble prediction schemes. Traditional model evaluation measures do exist, but the uncertainty of data, combined with the complications due to the chaotic nature of the systems (and models) considered, is shown to advocate the construction of a new measure, the *ν -shadowing time*, presented here. Applications of this measure, specifically in localising model error and as an assimilation technique, are discussed and demonstrated.

Improvements in the accuracy of estimated initial conditions from uncertain data using erroneous models are however, finite, suggesting the use of ensembles of initial conditions to provide probabilistic predictions. We consider the assumptions made in defining ensembles and highlight the need to interpret ensemble predictions in light of the formation scheme used; particular attention is given to the ensemble prediction schemes used in operational weather forecasting. Given the computational constraints of running the high-dimensional Numerical Weather Prediction models, ensemble evaluation is practically limited to the manipulation of data produced operationally.

¹Both are, of course, chaotic.

A new statistic is developed for which this data is sufficient to enable the validity of an important assumption (namely linearity) to be evaluated. The operational ensemble prediction schemes are found wanting. Kind consideration of the European Centre for Medium-range Weather Forecasting (ECMWF) allowed a modification to the operational ensemble suggested by these results to be implemented; the evaluation of this modified ensemble is also presented. All of the main results presented in this thesis are summarised in Chapter 7, as are the conclusions which may be drawn from them.

Outline of research Ultimately we wish to predict observed physical systems. Given observations, we need to construct a model of the system, which requires some understanding of the underlying dynamics of the system. The behaviours which may be exhibited by a relatively simple dynamical system are briefly considered and several methods of model formation are then introduced. Given a specific model structure there is still the question of what is the optimal number of degrees of freedom, or order, to use? At some model order the information required to describe a further increase in order will exceed the resulting gain in information about the system (from increased accuracy of predictions). If exact data are available, then models of different order may be compared using either prediction errors or information criteria (which are designed so that their minimisation indicates the optimal balance between model order and model accuracy). New results are presented to compare these measures, calculated using a variety of radial basis function (RBF) models of the Moran-Ricker map.

While methods used to build models from exact data are often applied when using observed data, high order models may now fit the noise in the data rather than the dynamics of the system. RBF models of the Moran-Ricker map are again used to illustrate this phenomenon and to compare the different measures (calculated with respect to the exact data), but now the models are constructed from noisy data. It is important to note that the measures of model optimality used here require knowledge of the exact data. Physical systems are imperfectly observed; data is uncertain and an alternative measure of model optimality must be employed. Further, we wish to compare models with different structures as well as models of different orders. Throughout this thesis one dimensional analytical systems will be used to illustrate the questions to be addressed wherever possible. Later discussions require consideration of a variety of systems with different properties and of varying complexity however, and so an overview of the systems and models used in this thesis is presented in Chapter 2.

We require a measure of model optimality to compare erroneous models using uncertain data, but first we discuss how to represent the uncertainty in the observations and note the implications of imperfection in the models: imperfect models will have model error, but this should not be confused with system sensitivity which is an inherent property of the system (and hopefully also of the model). In 1963 Lorenz found a manifestation of this sensitivity: small changes in the initial state of a vastly simplified weather model [45] rapidly led to disparate evolutions. ‘Weather patterns’ which initially differed by small uncertainties quickly evolved to very different weather scenarios; this phenomenon is often quoted as ‘The Butterfly Effect’². If our predictions are to be accurate then we need to use models which are able to capture sensitivity exhibited by the system (and observed). Further, measures of model optimality should reflect the similarity, or otherwise, of model sensitivity and system sensitivity.

The root mean square (r.m.s.) error is often used to quantify optimality as many model construction techniques aim to minimise one step prediction error in the r.m.s. sense. If the model state space is equivalent to that for the system and the data with which the errors are calculated is exact, then the r.m.s. error will show the system equations, *i.e.* the perfect model, to be optimal. Using models of the Ikeda map it is shown that when calculated using inexact observations, the r.m.s. error may reject the system in preference for an inexact model. A new measure, the ι -shadowing time, is therefore constructed in light of the uncertainties present and data available. In Chapter 3 ι -shadowing is demonstrated to exploit model sensitivity while penalising model error; it will always rank a perfect model as optimal.

ι -shadowing time distributions are then calculated and compared for various low-dimensional models of the logistic and Moran-Ricker maps and for the Lorenz system in Chapter 4. The ability of this new measure to evaluate competing models is discussed in light of the information available, which is considerable since the perfect models are known. When an imperfect model is used to try and ι -shadow observations of the Lorenz system, the regions in which shadowing fails are found to be correlated to regions of rapid error growth. A systematic study of the relation between model error and regions of shadowing failure is presented for the Sinai map, and the results applied to improve a radial basis function (RBF) model of the rotating fluid annulus. Other applications of ι -shadowing, arising through its use as an assimilation technique, are discussed briefly. While an analysis, or best guess of the initial condition, may be

²It is a popular abbreviation of Lorenz’s question ‘*Predictability: does the flap of a butterfly’s wings in Brazil set off a tornado in Texas?*’ which he addressed at a meeting in 1979 [30].

computed using assimilation techniques, the improvements possible are finite, and the focus therefore moves to the use of ensembles to provide probabilistic predictions.

Probabilistic predictions aim to quantify the potential error in a prediction due to the analysis error (or observational uncertainty) by evolving ensembles of initial conditions. Various methods of ensemble construction are introduced in Chapter 5. Perfect and unconstrained ensembles are contrasted using the Marzec-Spiegel system, whose system equations are known, while a RBF model of the thermally driven rotating annulus enables the effects of model error on ensemble formation to be investigated. Constrained vector (CV) subspaces used in the construction of ensembles in Numerical Weather Prediction are defined. The theoretical justification that a constrained vector ensemble will achieve a given aim is shown to depend upon a number of assumptions. Performance evaluation measures of Ensemble Prediction Schemes (EPS) used in Numerical Weather Prediction (NWP) are usually limited to a set of scores, standard within the NWP community, which quantify forecasting ability of the ensemble [12, 100]. Evaluation of ensemble performance should assess both the success of the ensembles at fulfilling their aims and the validity of any assumptions made in defining the ensemble.

Novel performance measures, quantifying both the ability of an ensemble to achieve its aims and the validity of assumptions used in definition of the ensemble subspace, are evaluated for a perfect model of the Marzec-Spiegel system and a RBF model of the annulus where the true state of the system is either known or may be well approximated. Due to the expense of running the high-dimensional operational Numerical Weather Prediction models, any validation of the Ensemble Prediction Schemes employed must be procured using values operationally available (*i.e.* those which are produced routinely). The true system values are not known for these models, and analysis values can not be considered to approximate the truth well enough for the same performance measures to be applied. The employment of twin (equal and opposite) pairs of perturbations in operational ensembles does however, enable the relevance, or otherwise, of the linear approximation to be assessed for operational Ensemble Prediction Schemes. The linear approximation is assumed to be relevant for operational perturbations in the definition of singular vector (SV) ensembles employed by the European Centre for Medium-range Weather Forecasting (ECMWF), and is (indirectly) assumed in formation of breeding vector (BV) ensembles employed by the American National Centre for Environmental Prediction (NCEP).

In Chapter 6 a new statistic, calculated from the nonlinear evolution of twin pair perturbations, is constructed in order to quantify the error made in assuming that

linearity holds. Such a measure allows a verification, or otherwise, of the internal consistency of SV and BV ensembles with respect to the linearity assumption made in defining the subspaces. In order to gain insight into how nonlinear processes affect perturbation evolution, two other measures are introduced: q -pling times which describe the time taken for a perturbation to amplify by a factor of q ; and nonlinear amplification rates which, for SV perturbations, may be compared with their linear equivalents given by the singular values. Linearity results are presented for CV ensembles evolved under a perfect model of the Marzec-Spiegel system, and under a RBF model of the annulus. The relatively inexpensive running costs of these low-dimensional models enables several sets of parameters defining the CV perturbations to be assessed, and the results obtained motivated the calculation of statistics for operational NWP ensembles.

Operational perturbations are commonly held to evolve approximately linearly for at least two days; the results presented here show this *not* to be true. Implications of the NWP linearity results are discussed and, together with those of the Marzec-Spiegel system and the annulus results, they suggest a modification to the operational ECMWF ensemble. Collaboration with ECMWF has enabled such a modification to be implemented using the operational NWP model and its effects investigated; in light of the results, which are presented here, ECMWF is currently considering modification of its operational ensemble.

Determining the distribution of times for which the linear approximation is relevant for current atmospheric models is one of a number of questions to be addressed on a list compiled by those gathered at the Newton Institute in Cambridge to discuss atmospheric predictability in 1996; see the ‘Summary of Addressable Questions’ in the overview of the meeting [81]. The results obtained using twin pair perturbations in operational models provide an answer to this question. Consideration of linearity results for individual SV perturbation pairs, of q -pling times and of nonlinear amplification rates enables discussion of two other questions listed, namely:

- Does the linear approximation break down in a systematic manner? In particular, is the image of the direction corresponding to the largest singular value (for a relevant optimisation time) the first in which the linear approximation fails?
- Many linear-dynamics-based intuitions are violated even in low-dimensional nonlinear systems, like the Lorenz 1963 model; yet these NWP models appear to behave consistently with these intuitions. Is there some principle which

indicated that there are pathologies which happen only in low order systems. Do these occur “Even In, Or Only In” (EIOOI) low order systems?

The latter is an iteration of the question of whether results from low-dimensional systems ‘really apply to the atmosphere’, in the words of *Lorenz* [45]; all are discussed in light of the results presented.

Chapter 2

Reconstructing dynamical systems

In this chapter some dynamical system terminology is introduced and the different behaviours which may be exhibited by a relatively simple dynamical system, whose equations are known, are then investigated. Having gained some understanding of the system, we consider how to build a model of the system from data if the equations are unknown; several methods of model formation are introduced. If exact data is available, then models may be compared using a variety of measures. For models constructed using exact data, errors are shown generally to decrease with increasing model order. While the initial increase of model order precipitates a rapid reduction of error, further increases of order are found to have less impact on the error. At some model order the information required to describe a further increase in order will exceed the resulting gain in information about the system (from increased accuracy of predictions). Information criteria are designed so that their minimisation indicates the optimal balance between model order and model accuracy, in terms of the amount of information required to describe the model and corresponding errors. New results are presented to compare optimality defined by information criteria with that defined by prediction errors for RBF models of the Moran-Ricker map.

Physical systems are imperfectly observed, so the data is uncertain. While methods used to build models from exact data are still applicable when using observed data, high order models may now fit the noise in the data rather than the dynamics of the system. This is illustrated by calculation of prediction errors (with respect to exact data) for models constructed from noisy data. As the model order increases, the prediction error decreases to a minimum before increasing again as the model becomes over-fit. An alternative measure of model optimality is introduced in the next chapter.

In this thesis a one dimensional analytical system will be used to illustrate the questions to be addressed wherever possible; later discussions require consideration

of a variety of systems with different properties and of varying complexity. Reconstruction of system state spaces from observations is discussed in section 2.3.1, before an overview of the systems and models used in this thesis is presented.

2.1 Introduction to dynamical systems

2.1.1 Flows and maps

Finite dimensional dynamical systems may have time either as a continuous or discrete variable. If time is represented as being continuous then the system is a *flow* and can be represented as a set of m first order autonomous ordinary differential equations¹ (o.d.e.s), represented as

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{f}(\mathbf{x}) \quad (2.1)$$

where $\mathbf{x} = (x^{(1)}, x^{(2)}, \dots, x^{(m)})$ and is called the state of the system. The evolution of the system from a given state at time t_0 traces a continuous path, or *trajectory*, in $(x^{(1)}, x^{(2)}, \dots, x^{(m)})$ space, a *state space* of the system, as illustrated for the Lorenz system (as described in Lorenz’s 1963 paper [45]) in figure 2.1 (see section 2.3.2.4 on page 31 for details).

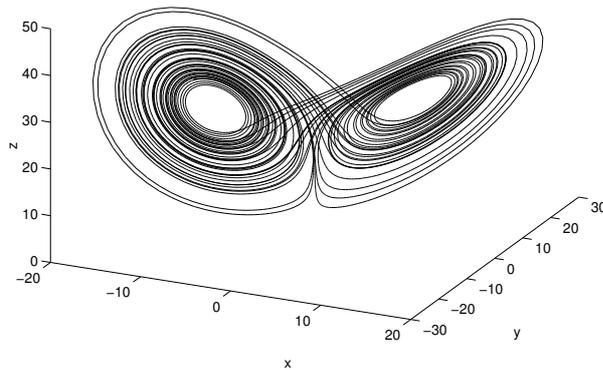


Figure 2.1: A trajectory of the Lorenz system in the (x, y, z) state space. Starting from an initial condition in the left half of the state space ($x < 0$), the trajectory spends some time investigating the left ‘wing’ before crossing over to the right half of the state space ($x > 0$) to investigate the right ‘wing’, etc.; the trajectory may return arbitrarily close to itself, but never crosses.

¹If the original system has explicit time dependence then representation as a set of autonomous equations requires the state space to be expanded, and the solution may be unbounded [64].

A discrete time system is called a *map*, defined by

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t), \quad \mathbf{x} \in \mathcal{R}^m \quad (2.2)$$

where t denotes the number of discrete time steps. Given an initial condition \mathbf{x}_0 , iteration of the system forwards generates an ordered series of points, a discrete path or an *orbit*. If there are multiple values which evolve under the map to the same image, *i.e.* multiple \mathbf{x}_t for which $\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t)$ are identical, then the map is said to be *non-invertible*. Consider the Moran-Ricker map [51, 69], constructed to describe the evolution of fish population dynamics and given by

$$x_{t+1} = x_t \exp\{r(1 - x_t)\}. \quad (2.3)$$

For $r = 3.7$, values near to 0.112 and 0.533 are both mapped to 3.0, as illustrated in figure 2.2; the Moran-Ricker map is non-invertible.

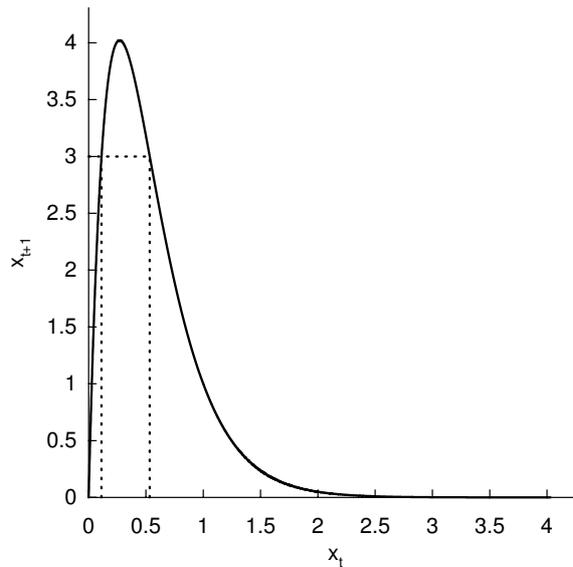


Figure 2.2: Non-invertibility of the Moran-Ricker map (solid line) with $r=3.7$: the map is two to one for a range of values of x , *e.g.* both $x \simeq 0.112, 0.553$ are mapped to 3.0 (as indicated by the dotted lines).

Numerous maps may be created by sampling a continuous flow in different ways. Sampling the flow at equal time intervals gives an md map which depends on the sampling interval. Alternatively a $(m - 1)d$ surface may be chosen and the sequence of states $x_{t-1}, x_t, x_{t+1}, \dots$, recorded when the flow's trajectory crosses the surface in a given direction, thereby defining a map, often called a Poincaré map [24], which will, of course, depend on the choice of surface. All maps generated by sampling continuous flows, and the continuous flows themselves, are *invertible* since the differential

equations from which they arise, integrated backwards in time, give a unique prior state [57].

2.1.2 Onset of chaos

The flows and maps used to describe dynamical systems usually contain parameters which, when varied, change the behaviour of the system. Let us consider the Moran-Ricker map, with $r = 1$; there are two fixed points of the map, $x^* = 0, 1$. Local (linear) stability analysis [54] shows that a small perturbation to the fixed point $x^* = 0$ of the map will initially grow exponentially and the fixed point is said to be locally unstable to perturbations. The other fixed point of the map, $x^* = 1$, is locally stable for $0 < r < 2$. At $r = 2$, local stability of $x^* = 1$ is lost and the system is said to exhibit a *bifurcation*.

Figure 2.3 shows the second iterate of the Moran-Ricker map and the line $x_{t+1} = x_t$. For $r = 1.75$ there is only one positive, non-zero point of intersection, $x = 1$, which is

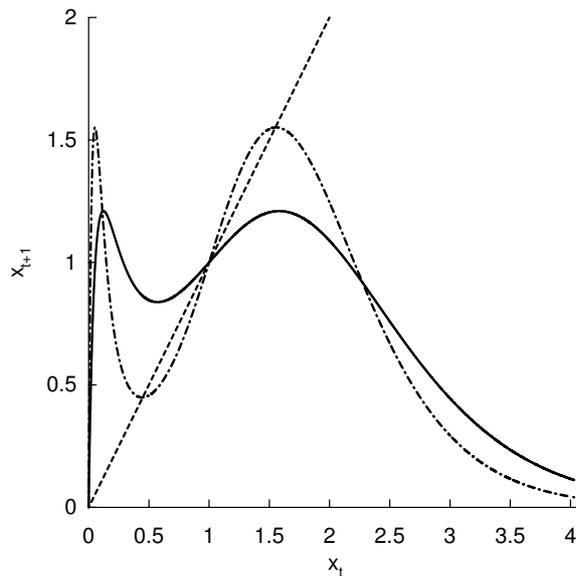


Figure 2.3: The second iterate of the Moran-Ricker map for $r = 1.75$ (solid) and $r = 2.25$ (dot-dashed). The change in the number of intersections with the line $x_{t+1} = x_t$ (dashed) is due to the period-doubling bifurcation as r passes through $r = 2$: $r = 1.75$ (solid) crosses $x_{t+1} = x_t$ only once, while $r = 2.25$ (dot-dashed) has 3 non-zero steady states.

the period-1 fixed point. As r increases through 2, the curve $f^2(x)$ becomes tangent to the line $x_{t+1} = x_t$ and a new, stable, period-2 orbit is formed, while the period-1 orbit at $x = 1$ becomes unstable; the Moran-Ricker map undergoes a period-doubling bifurcation at $r = 2$. Stability analysis may be repeated for the fourth, eighth, etc. iterates in order to find the sequence of bifurcation values as r increases further.

Alternatively, we may construct a bifurcation diagram which summarises the stable orbits of the map [3], as shown in figure 2.4.

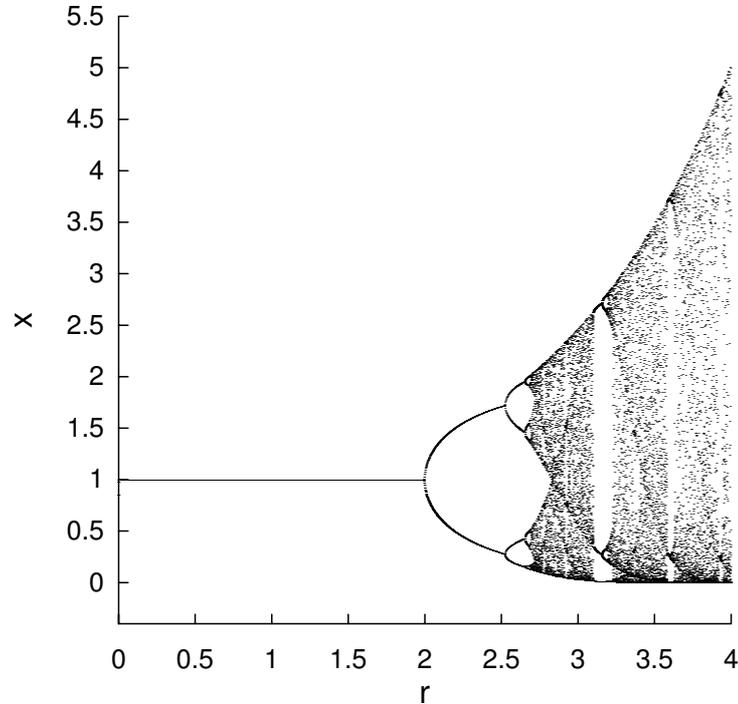


Figure 2.4: Bifurcation diagram for the Moran-Ricker map. The values of x on periodic orbits are shown as a function of parameter value, r . Note the appearance of a period-3 orbit at $r \simeq 3.2$.

To construct a bifurcation diagram, an arbitrary initial condition is taken for each parameter value, iterated forwards through any transient behaviour and its subsequent orbit plotted. In figure 2.4, the period-doubling bifurcation from a fixed point to a period-2 orbit at $r = 2$ is evident, as are the subsequent period-doubling bifurcations to a period-4 orbit at $r \simeq 2.53$ and to a period-8 orbit at $r \simeq 2.65$. The distance (in terms of parameter value) between successive bifurcations decreases and a cascade of period- 2^n , $\forall n \in \mathcal{Z}^+$, solutions is formed, until a critical parameter value, r_c , is reached, at which point all solutions of period- 2^n become unstable. For a value of $r > r_c$ a period-3 orbit emerges, which similarly undergoes a series of period-doubling bifurcations. For a 1d map, *Sarkovskii* [74] proved that for a continuous map on an interval, if an odd period (≥ 3) orbit exists for r_c then orbits of all periods exist; chaotic solutions (in the sense of *Li & Yorke* [44]) are said to exist for $r > r_c$. Note that a period-3 orbit at $r \simeq 3.2$ is clear in figure 2.4, indicating that the 1d Moran-Ricker map has chaotic solutions for some $r \gtrsim 3.2$. An alternative definition of chaos is given in the next section.

A bifurcation diagram provides detailed information about the behaviour of the system as the parameter values are varied. If the system equations are known, then construction is limited only by computational constraints. For an observed system, construction of such a diagram requires observations of the system in physically realisable states corresponding to different parameter regimes. Exceptionally, sufficient data is available, *e.g.* for the differentially heated rotating annulus [29] where the laboratory apparatus may be altered. Often such an extensive range of data is not available, complicating the determination of whether or not the system is in a chaotic state, as discussed below.

2.1.3 Attractors, Lyapunov exponents and chaos

A *conservative* system is volume preserving for some choice of state space variables [57], *e.g.* position or momentum; if no such choice of state space variables exists then the system is *dissipative*. The terminology derives from physical systems, for which preservation of volume in state space corresponds to conservation of energy (for some choices of variable). A dissipative system will characteristically have one or more *attractors*, with dimension $d \leq m$. For each attractor there will be a region of initial conditions, known as a *basin of attraction*, which will evolve asymptotically to the attractor under the system. In many attractors, different regions of state space will be visited with varying frequency, information given by an *invariant measure* of the system [21]. The invariant measure is rarely known analytically, but can be approximated by evolving the system forwards over a long period of time and recording the frequency with which the system visits different regions. Such an approximation to the invariant measure of the Moran-Ricker map is shown in figure 2.5. *Ergodicity* of a system implies asymptotic equality of a space average (weighted by the invariant measure) and a time average (for almost every initial condition, except possibly on a set of null measure) [5].

We will define *chaotic* behaviour of the system through sensitive dependence on initial conditions: consider the growth of an infinitesimally small perturbation, ϵ_0 , to a given initial condition \mathbf{x}_0 . For an infinitesimally small perturbation, the dynamics are given by the linearised flow [24]. Setting $\mathbf{x}(t) = \mathbf{x}_0 + \epsilon(t)$ and Taylor expanding about \mathbf{x}_0 , we obtain the linearisation

$$\dot{\epsilon} = \mathcal{J}(\mathbf{x}_0)\epsilon; \tag{2.4}$$

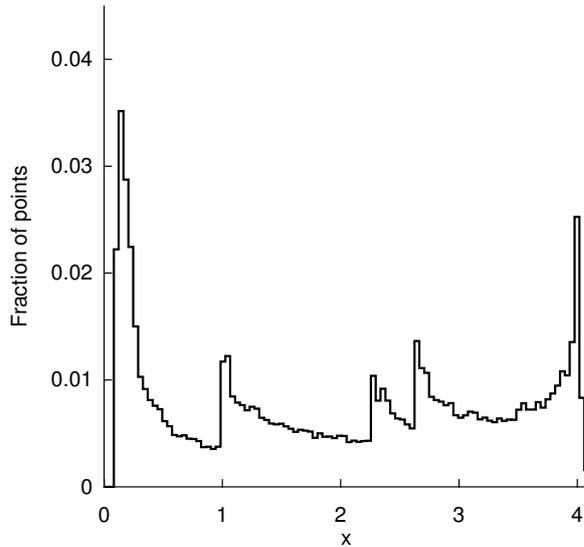


Figure 2.5: An approximation to the invariant measure of the Moran-Ricker map, for $r = 3.7$, given by a histogram of the frequency of occurrence of values of x .

$\mathcal{J}(\mathbf{x})$ is the *Jacobian* of \mathbf{f} at \mathbf{x} , defined as

$$\mathcal{J}_{ij}(\mathbf{x}) = \left. \frac{\partial f_i}{\partial x_j} \right|_{\mathbf{x}}. \quad (2.5)$$

Differentiating equation 2.1 with respect to \mathbf{x}_0 and defining the linear operator (also called the linear or tangent propagator), by

$$\mathcal{M}(\mathbf{x}_0, t) = \frac{d\mathbf{x}(t)}{d\mathbf{x}_0} \quad (2.6)$$

gives

$$\dot{\mathcal{M}}(\mathbf{x}_0, t) = \mathcal{J}(\mathbf{x}(t))\mathcal{M}(\mathbf{x}_0, t). \quad (2.7)$$

It follows that ϵ evolves according to

$$\dot{\epsilon}(t) = \mathcal{M}(\mathbf{x}_0, t)\epsilon_0. \quad (2.8)$$

The linearisation thereby provides a description of the initial local behaviour of ‘nearby’ initial conditions. No information is provided about the global characteristics of the system, and, strictly speaking, information is only valid locally while the perturbation remains infinitesimal.

Lyapunov exponents aim to give a measure of the averaged stretching and contracting characteristic of an attractor through the exponential average rate of divergence of points initially nearby. Consider a 1d system, f , and two points initially separated by an infinitesimal distance ϵ . After time T the separation between the two points,

initially at x_0 and $x_0 + \epsilon$, will be $|f^T(x_0 + \epsilon) - f^T(x_0)|$; the Lyapunov exponent, Λ , calculated in the limit as $t \rightarrow \infty$ and $\epsilon \rightarrow 0$, is given by [21, 90]

$$\Lambda(x_0) = \lim_{T \rightarrow \infty, \epsilon \rightarrow 0} \frac{1}{T} \log \left| \frac{f^T(x_0 + \epsilon) - f^T(x_0)}{\epsilon} \right| \quad (2.9)$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T} \log \left| \frac{df^T(x_0)}{dx} \right|. \quad (2.10)$$

It should be emphasised that definition of the Lyapunov exponent assumes that the initial conditions lie on the attractor. The above definition of the Lyapunov exponent also depends on the trajectory of the system on which the initial condition lies, but due to the ergodicity of the system Λ is independent of x_0 almost everywhere [21].

For higher dimensional systems a theorem of Oseledec applies [55], enabling a *spectrum of Lyapunov exponents*, $\mathbf{\Lambda}$ to be similarly defined (see *e.g.* Drazin [19]). If the sum of this global spectrum is negative then the system has an attractor, while at least one positive exponent indicates *sensitive dependence to initial conditions*, *i.e.* the average divergence of two points initially arbitrarily near to one another is exponential; the system is then *chaotic* [52]. For a chaotic system to have an attractor, nearby trajectories must remain in a finite region of state space without self-intersecting. If the system is linear then divergence between initially nearby points will be unbounded, so nonlinearity is required if a chaotic system is to have an attractor. In a map, orbits which are initially nearby may have on average exponential divergence while remaining in a finite line interval; since the orbits are a sequence of points and there are an infinite number of points contained in any 1d interval, 1d maps may be chaotic. For a flow, trajectories often avoid self-intersecting by stretching and folding of the state space under evolution, requiring at least a 3d state space. This folding and stretching of the state space may create attractors with intricate fractal structure [56], known as *strange attractors*, such as the Lorenz attractor illustrated in figure 2.1.

Methods for estimating Lyapunov exponents from experimental time series often only work for limited noise in the data (see *e.g.* Eckmann and Ruelle [21]). While more complex methods have been developed to overcome this problem [17], simple techniques to approximate the largest Lyapunov exponent are often used [102], since this allows classification of the system as chaotic or otherwise. Evaluating the accuracy of such estimates remains problematic.

2.2 Modelling data

Often the only information which we have about a system is a sequence of observations of one or more physical variables. In order to better understand, and predict, the system we often wish to construct a model of the system. If a system is to be modelled from observational data, the data must contain sufficient information to fully describe the system. The position of an oscillating pendulum is an intrinsic variable, since on differentiation it yields the velocity and acceleration of the pendulum, together sufficient to fully describe the system. Physical variables need not be intrinsic, although it is often assumed. One method of reconstructing a state space of the observed system is to use a variable and its derivatives, *e.g.* position, velocity and acceleration as above. Manipulations of observed data to reconstruct a state space of the system are discussed later in section 2.3.1. Here we assume that such a state space exists, and that the dynamics of the system may be represented as

$$\mathbf{x}(t + T) = \mathbf{f}_T(\mathbf{x}(t))$$

where $\mathbf{x}(t)$ is the state of the system at time t and we wish to predict the future state in T time steps.

Provided that sufficient observations are available, estimation of the future state by $\tilde{\mathbf{x}}(t+T) = \tilde{\mathbf{f}}_T(\mathbf{x}(t))$ may be considered as a problem of constructing an interpolation function $\tilde{\mathbf{f}}$ to model the system dynamics. A brief description of some techniques used to formulate models follows; some are *direct* methods which involve constructing a new approximating function $\tilde{\mathbf{f}}_T(\mathbf{x}(t))$ for each T , while others are *iterative* methods which construct $\tilde{\mathbf{f}}_1(\mathbf{x}(t))$ and then apply it iteratively T times. There is also a choice of whether to construct a *local* or *global* approximation, discussions of which are given by *Casdagli* [15] and *Farmer & Sidorowich* [26]. Local approximations have the advantage that they need only fit the local, usually simpler, dynamics, but may be discontinuous and tend to require more information to specify them. In order to compare and evaluate models the observational data available should be divided into a learning set and a test set; the model may then be formulated from the learning set and its performance evaluated out-of-sample using the test set.

2.2.1 Building a model

2.2.1.1 Polynomial approximations

If \mathbf{f} is smooth, then by partitioning the state space into sufficiently small regions, \mathbf{f} may be well approximated locally by a polynomial approximation. An approximation

formulated using data from within a given distance would give a model with varying accuracy since the data will be distributed uniformly according to the invariant measure, and not uniformly in space. Further, the contribution of higher order effects will vary with location. The method of nearest neighbours aims to overcome this problem; it determines the *n-point nearest neighbourhood* of $\mathbf{x}(t)$ as the set of observed points $\{\mathbf{y}_i(t)\}$, $i = 1, \dots, n$, which minimise

$$\sum_i \|\mathbf{x}(t) - \mathbf{y}_i(t)\|$$

for some given metric $\|\cdot\|$. This also gives some overlap, and hence continuity, between regions of state space.

A linear approximation of $\mathbf{x}(t + T)$ is given by fitting a function to the pairs $(\mathbf{y}_i(t), \mathbf{y}_i(t + T))$ from the $(m + 1)$ -point nearest neighbourhood of $\mathbf{x}(t)$ chosen from the learning set. Success of such an approximation depends on there being enough data for the neighbourhood to be sufficiently small for the dynamics to be well approximated by a linear function. Higher order polynomial approximations are defined similarly, with success depending on the polynomial well approximating dynamics in the neighbourhoods given the local data density.

Nearest neighbour polynomial approximations may be implemented as direct or iterative methods. *Farmer and Sidorowich* [26] compare the rate of error growth for direct and iterative linear approximations analytically. Their calculations make several assumptions, including the assumption that the model is sufficiently accurate for the Lyapunov exponents of the model to be similar to those of the system. While the assumptions hold, iterative methods are superior to direct methods; they include however, an experimental example where the opposite is true.

2.2.1.2 Radial Basis Functions

Radial basis function (RBF) models are a global interpolation scheme based on a set of n_c centres in the md state space, denoted as

$$\mathbf{c}_j, \quad j = 1, \dots, n_c \quad \mathbf{c}_j \in \mathcal{R}^m.$$

In the simplest case the centres are chosen to be points in the learning set, and each centre has an associated observation, s_j , which is some future value of the system.

The predictor of \mathbf{x} to estimate s is given by the function

$$\tilde{f}(\mathbf{x}) = \sum_{j=1}^{n_c} \lambda_j \phi(\|\mathbf{x} - \mathbf{c}_j\|) + \sum_{n=1}^d \mu_n P_n(x),$$

where $\phi(r)$ is the chosen *radial basis function* ($\phi(r) = r^3$ is a typical (and physically very curious) choice, see [15]), $\|\cdot\|$ is the Euclidean norm and $\{P_n; n = 1, 2, \dots, d\}$ is a basis for polynomials of degree at most d . Frequently the polynomial term is omitted; using $\tilde{f}(\mathbf{x}_i) \simeq s_i$, the λ_j may be determined by solving

$$\mathbf{b} = A\lambda \tag{2.11}$$

where

$$A_{ij} = \omega_i \phi(\|\mathbf{x}_i - \mathbf{c}_j\|), \quad b_i = \omega_i s_i,$$

for $i = 1, \dots, n_L$, the number of points in the learning set, and $j = 1, \dots, n_c$; where the ω_i are chosen weightings, for details see [78].

There are various options for solving equation 2.11. If the centres are chosen to be distinct data points, and the s_j are their images in evolution for some time T , A is square and one can either use a standard QR factorisation algorithm, see *Casdagli* [15] for details, or alternatively singular value decomposition (SVD) using a Moore-Penrose pseudo-inverse (as described in Appendix C, or see *Smith* [78] for details). In order to utilise all the available information we shall follow *Broomhead & Lowe* [11] and use the entire learning set for \mathbf{b} , with a limited number of linear combinations of centres; A is no longer square and instead of solving equation 2.11 we seek to minimise $\psi^2 = \|\mathbf{b} - A\lambda\|^2$. If SVD is used, then the order of the model is determined by the number of singular values ‘retained’ (*i.e.* the number of values not set to zero when calculating the Moore-Penrose pseudo-inverse). These singular values give an ordering of the orthogonal vectors, obtained by application of SVD, which are linear combinations of the radial basis functions. This ordering gives no information about the importance of the orthogonal vectors for prediction [83].

Using radial basis functions as a global interpolation technique provides a smooth function with fixed coefficients, making it a technique suited to processing large data sets [79]. The length of learning sets available even from laboratory experiments are often too short to expect to achieve good results for systems with many active degrees of freedom; work is being done to incorporate known underlying physics of the system in the predictor [78].

2.2.1.3 Models using physical principles

The above methods attempt to reconstruct the dynamics of the system by fitting a set structure to the observed data. Alternatively, it is possible to combine physical principles in an attempt to simulate the system itself. Construction of a set of governing physical equations is unique to the system and usually requires including a

large number of variables and parameters which results in a complex set of equations; the observational data may then be used to fit the parameters. Numerical weather prediction methods often employ such a model, as described below in section 2.3.4.

2.2.2 Determining the optimal order of a model

Having chosen a method by which to construct a model, there is a choice of how many degrees of freedom to allow the model; how this is manifested depends on the model structure, *e.g.* it may be determined by the order of polynomial used, the number of nearest neighbours considered, the number of centres used or the number of physical properties included. In order to compare models of different orders, we need some measure of optimality.

If the observations are noise free, then the parameters in the models are ‘fit’ using the data, *e.g.* by minimising the root mean square distance between the one step model prediction, $\hat{x}_{t+1} = \tilde{f}(x_t)$, and the corresponding system value, x_{t+1} , for the data in the learning set. While initial increases in model order usually result in a marked decrease of this *in-sample* error, further increases in order have a less notable effect on error; at what point is the small improvement in accuracy not a worthwhile return for the increase in complexity of the model? *Information criteria* are designed to indicate an optimal balance between model order and model accuracy, by measuring the amount of information required to describe both the model and corresponding errors.

Given a model $\tilde{f}_{a_1, a_2, \dots, a_k}$, which depends on k parameters, a_j ; $j = 1, \dots, k$, and a learning set of data $\{y_t\}$, then maximising the probability $\mathcal{P}(\{y_t\}|\{a_j\})$ of observing the sequence $\{y_t\}$ given the model \tilde{f} is equivalent to maximising the log likelihood function, $L = -\ln \mathcal{P}(\{y_t\}|\{a_j\})$. If the one step prediction errors, $e_{t+1} = \tilde{f}(y_t) - y_{t+1}$, have Gaussian distribution, then L is given by:

$$L = -\frac{N}{2} [\ln \hat{e}^2 + \ln 2\pi + 1],$$

where \hat{e} is the average one step prediction error over N data points; the log likelihood function is maximised by minimising the average mean square error. The *Akaike information criteria* [1] incorporates the complexity of the model by adding a term $-k$ to the log likelihood function; this is equivalent to multiplying the mean square error by a factor of $[\exp(\frac{k}{N})]$. If two models have comparable predictions, the Akaike information criteria will prefer the simpler of the models. *Rissanen* [70] considers the optimal model to be that which minimises the length of a description of the system,

considered as the amount of information required to specify the model parameters and the residual errors. It can be shown (see *Kantz & Schreiber* [41] for details) that this equates to adding a term $-\frac{k}{2} \ln N$ to the log likelihood function, or equivalently multiplying the mean square error by a factor of $[\exp(\frac{k}{2N} \ln N)]$. Obviously such information criteria are not unique, and penalising model complexity only according to the number of parameters it contains is an issue of contention.

The above discussion focuses on the average in-sample error; we wish an optimal model to be the best representation of (or to require the least information to describe) the system in general, and not just for the data sample from which the model was constructed. This suggests calculating the one step prediction error averaged over a set of points distinct from the learning set to obtain the out-of-sample error. This may either be done using a test set (which will be distributed according to the invariant measure) or alternatively, if the system is known analytically, the *function fit* of the model to the system may be measured as the average error between numerical integrations of the model and the system over the domain of the system.

Data is rarely exact and is usually contaminated with noise. If the model is considered as a function expansion, then increasing the order of the model is equivalent to extending the span of the model space. While an increase in order will hopefully increase the intersection of the model space with the system space, it also allows additional freedom to fit the particular realisation of noise in the data. While initial increases in order may allow the principal dynamics of the system to be captured, we expect some limit beyond which the additional span of the model fits a particular realisation of the noise distribution rather than improving the approximation to the system dynamics. In practice the best order represents a balance between capturing the dynamics of the system without fitting the noise, and may be considered to be that which minimises the out-of-sample one step prediction error or function fit error.

In order to compare the different error measures introduced above, and to examine the effect of noisy observations on model construction, we need models with varying numbers of parameters fit using both accurate and noisy data. The Moran-Ricker map is used as it is easy to iterate and yet has a structure which is not well fit by standard techniques; radial basis function models do provide a reasonable fit, and the order of the model is determined by the numbers of centres, the choice of which is a standard procedure when singular value decomposition is used. The decreasing return of improved accuracy with increasing model order is demonstrated before investigating over-fitting of models due to the effects of noisy data.

2.2.2.1 Exact data

Radial basis function models of the Moran-Ricker map are constructed using up to 12 centres chosen from 12 exact data points, which are a segment of an orbit and hence distributed according as the invariant measure. A Gaussian radial basis function was chosen so that each centre will have fairly localised influence. A reasonable choice of variance is the average distance between the centres, which for a map with range 4.02 and 12 centres would be 0.34. As can be seen from figure 2.6 however, the centres are not equally spaced; many are much closer together and a variance of 0.063 was found to give models with much smaller errors, so results are presented using this variance. The weightings ω_j were set to unity. Models were constructed using between 1 and 12 centres, and these were then used to predict the values for 128 equally spaced points to give a function fit, illustrated in figure 2.6. As compared to the system map (see figure 2.2 on page 9), it can be seen that the initial increase in order (from using, say, 1 centre to 4 centres) drastically improves the model's representation of the system. Further increases in model order have less obvious impact, although the introduction of a minima and maxima not seen in the system for the model with 12 centres is noticeable.

As mentioned above, the use of singular value decomposition provides *an* ordering of the linear combinations of the centres by the magnitude of their associated singular values. Figure 2.7 gives the singular values corresponding to a model using upto 12 centres; model k uses the (linear combinations of) centres with the k largest corresponding singular values. Noting the logarithmic scale, the first 9 singular values are of a similar order of magnitude, the tenth is slightly smaller and the last three are each several more orders of magnitude smaller.

In order to evaluate which order of model is optimal, the errors discussed above are calculated. In order to facilitate comparison between models a factor of the average deviation of the images (under the system, with the average taken over points distributed with respect to the invariant measure) is included in all the errors. The average in-sample one step prediction error is calculated as the root mean square of the different models' predictions over the 12 data points from which the centres were chosen, the function fit is the root mean square error of predictions for 128 equally spaced data points on the interval of the map, and the out-of-sample one step prediction error is averaged over 256 points of an orbit (hence they are distributed according to the invariant measure). The Akaike and Rissanen information criteria errors are also evaluated.

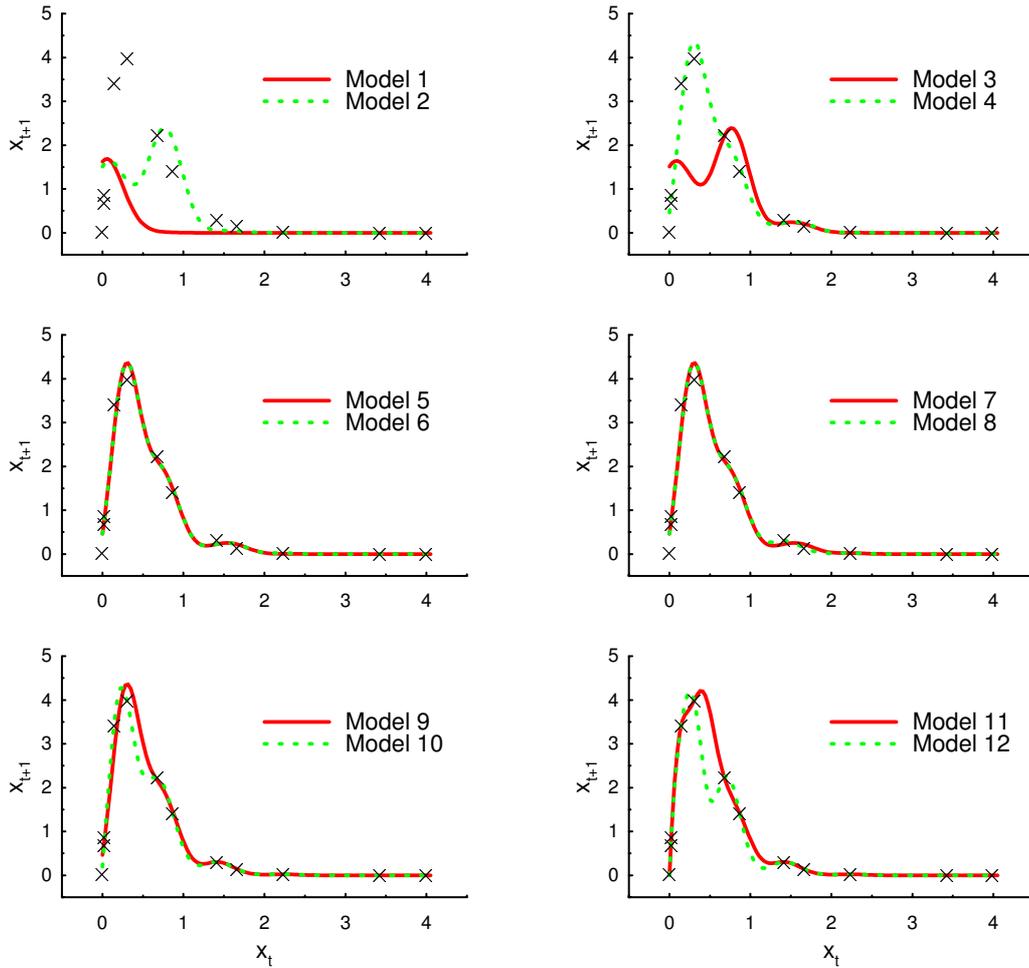


Figure 2.6: Predictions of Gaussian RBF models of the Moran-Ricker map constructed using up to 12 centres chosen from 12 accurate data points. Each diagram shows the data points (crosses) from which the centres may be chosen, and the predictions of the 12 models which use between 1 centre (Model 1) and all 12 (Model 12).

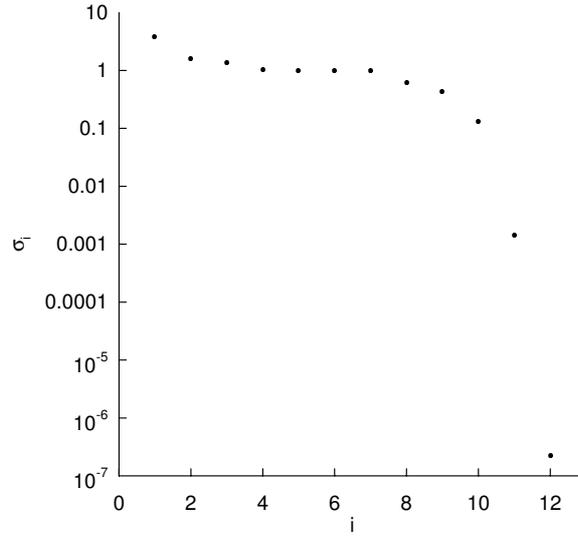


Figure 2.7: Singular values, σ_i , for Gaussian RBF models of the Moran-Ricker map with up to 12 centres chosen from an exact 12 point trajectory.

Figure 2.8 enables comparison of the different errors, with the minimum of each error indicating the optimal order of model using that error measure. As expected,

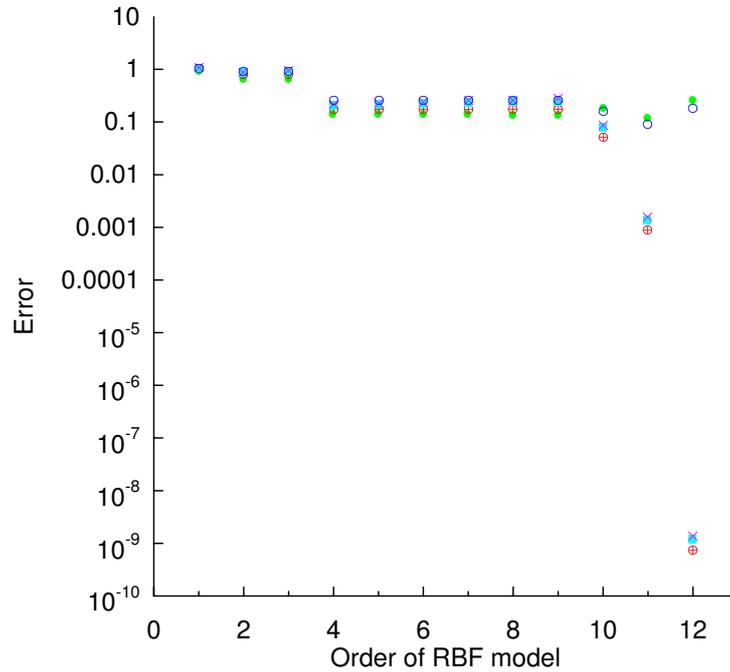


Figure 2.8: Errors for Gaussian RBF models of the Moran-Ricker map, the order of which is the number of centres used. The average root mean square error is calculated in-sample (red plus sign in circle), out-of-sample (blue open circle), and the function fit is determined (green filled circle). The Akaike (cyan filled square) and Rissanen (magenta cross) information criteria errors are also shown. Minima denote the optimal model order using the various error measures.

the in-sample error is monotonically decreasing since the minimisation of this error is used to construct the RBF model. The out-of-sample prediction error decreases monotonically to a minima at an order 11 model; including a twelfth centre makes the out-of-sample one step predictions worse, on average. This is unsurprising given the very small corresponding singular value; with limited accuracy of model it may not be optimal (in terms of out-of-sample prediction error) to overly constrain the model to the limited data. The function fit gives similar results, with any differences due to the equal spacing of points at which the error is evaluated. The improvement in in-sample error is greater than the penalty given by the information criteria for changes in model order from order 1 to order 2, from order 3 to order 4, and for each increase in order from order 9 to order 12. Local minima of the information criteria can however, be seen for order 2 and order 4 models.

These results demonstrate the importance of testing a model out-of-sample. The global minima of both the out-of-sample one step prediction error and function fit are at order 11, while the in-sample one step prediction error and the information criteria have global minima at order 12.

2.2.2.2 Noisy data

Given only the 12 (exact) data points above, the best model uses 11 of the points as centres. Increasing the amount of data available, and the number of centres used, enables a fuller examination of the effects of noisy data on optimal model order; below 64 centres are chosen from 1024 data points. The centres are determined using exact data (so as to simplify the comparison between models); corresponding singular values decrease gradually, as shown in figure 2.9. Models are then constructed using up to 64 of these centres, with parameters determined using either the exact data or one of 16 realisations of noisy data sets, which are contaminated with noise which has Gaussian distribution and a standard deviation which is 1% of the range of the system, *i.e.* with $\mathcal{N}(0, 0.0402)$ distribution. Predictions of some of the models are shown in figure 2.10. A model of order 16 (*i.e.* with 16 centres) differs little whether it is constructed from exact or noisy data, and neither is a good fit to the system. At order 30, both models are a better representation of the system but exhibit very small oscillations *not* exhibited by the system. Increasing the model order to 38 or further, to order 55, rids the true model (constructed from exact data) of the oscillations, but exaggerates those found in the noisy model.

The errors are calculated (with respect to the system values) as above, and are shown in figure 2.11. The errors for the exact model are easy to discern; the in-

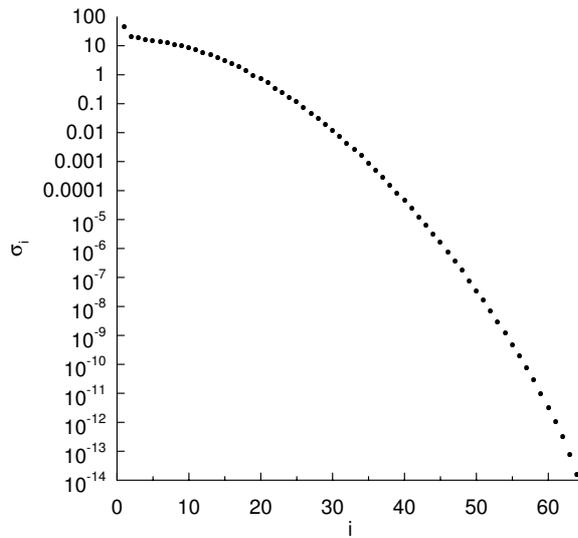


Figure 2.9: Singular values, σ_i , for a Gaussian RBF model of the Moran-Ricker map with 64 centres chosen from a 1024 point trajectory.

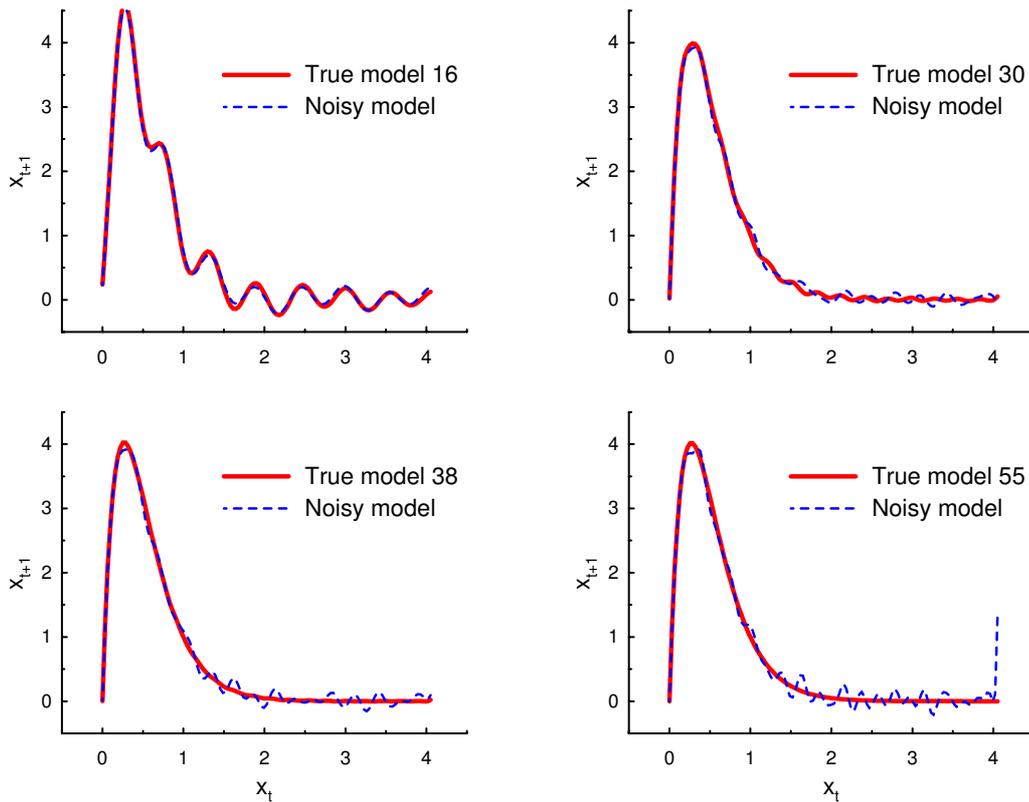


Figure 2.10: Predictions of Gaussian RBF models of the Moran-Ricker map, constructed using up to 64 centres chosen from a 1024 data point trajectory. Predictions for models constructed using 16 (upper left), 30 (upper right), 38 (lower left) and 55 (lower right) centres and clean images (red solid lines), or images contaminated with $\mathcal{N}(0, 0.0402)$ noise (blue dotted), are shown.

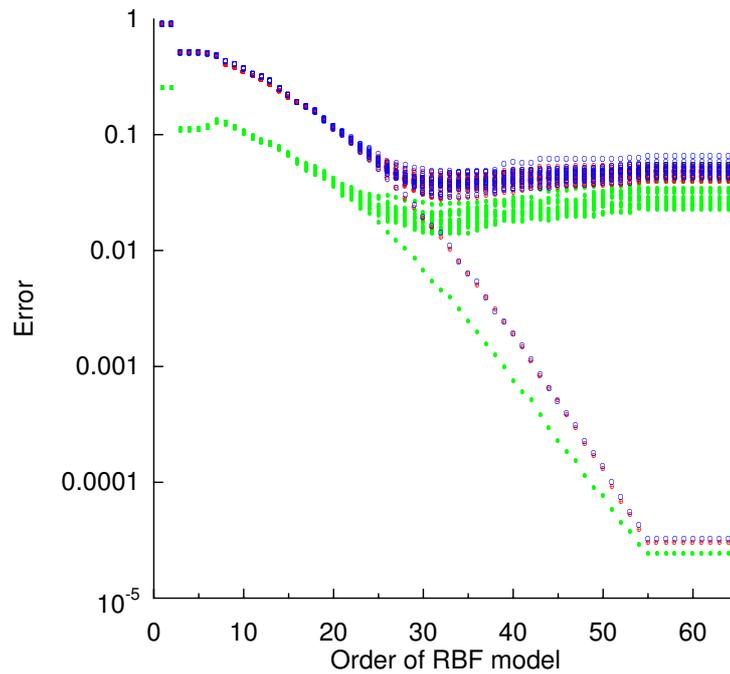


Figure 2.11: Errors for Gaussian RBF models of the Moran-Ricker map, the order of which is determined by how many of the 64 possible centres, chosen from a 1024 point trajectory are used. For each order of model, errors are calculated for each model constructed using either exact images or one of 16 realisations of noisy ($\mathcal{N}(0, 0.0402)$) images. The average mean square error is calculated in-sample (red plus sign in circle), out-of-sample (blue open circle), and the function fit is determined (green dot). The Akaike and Rissanen information criteria are not shown as they almost completely overlap the in-sample errors. Minima denote the optimal model order using the various error measures.

sample and out-of-sample prediction errors monotonically decrease to order 55 and then retain that value out to order 64. (The plateau is due to the tolerance level for the SVD calculation, which effectively limits the number of centres when the corresponding singular values are much smaller than the first singular value.) When the tolerance level is set to be more lenient, the out-of-sample error for the model constructed from exact data increases beyond order 55 (similar to the increase seen in changing from 11 to 12 centres above). The function fit error for the exact model has a local minima at 5, but the global minima is as for the prediction errors, *i.e.* at order 55. The information criteria errors are not shown as they deviate little from the in-sample error, for both exact and noisy models. The results for the noisy models vary with different realisations: the in-sample error realisations are on average minimal at order 34, while the out-of-sample prediction errors and function fit errors are minimal at order 29, on average. The optimality of models of orders $\simeq 30$, as defined by the error measures, is coherent with the system representations shown in figure 2.10 above. Note that, in this scenario, over-estimating the optimal order is far preferable to under-estimating it; *e.g.* the (out-of-sample) error of an order 50 model is significantly less than that of an order 10 model.

If the system is unknown, multiple noisy learning sets may be obtained either by over sampling the data and then re-sampling it, *e.g.* at 1 in 8, from different initial values, or by splitting long time series into distinct segments. Finding an alternative measure of model optimality, necessary since both the out-of-sample prediction error and function fit error require either knowledge of the system or noise free data, is less trivial and is the subject of the following chapter.

2.3 Other systems: lasers, annuli and the atmosphere

2.3.1 Reconstructing a state space

The governing equations of a system are not usually known; the information available is a time series of noisy observations. In order to reconstruct a state space of the system, we need the measured variable of the time series to be intrinsic of the system, *i.e.* the variable must be such that it contains sufficient information to fully describe the system. The classical way to proceed is then to consider derivatives of the variable, giving an *embedding* of the data in some reconstructed state space [56]. However, differentiation may amplify errors and thus make this approach unsuitable for experimental time series [52]; alternative methods of embedding are described below.

Embedding Consider a (scalar) time series of observations

$$\{y_k\} = \{y(t_k)\}, \quad k = 0, 1, \dots, N; \quad (2.12)$$

recorded with uniform sampling time τ_s from a stationary dynamical process². The y therefore correspond to the application of a measurement function to the true state of the process. If the dynamics are confined to an attractor of dimension d , then in order to prevent self-intersection of the embedded trajectory, and hence preserve determinism, it is necessary, but not sufficient, to embed the data in a state space of dimension m , with $m \geq d$. If the attractor is contained in a smooth n -dimensional manifold, work by *Takens* [91] shows that, for a generic measurement function, it is sufficient that $m \geq 2n + 1$ for an embedded reconstruction of the dynamics to be diffeomorphic to the dynamics of the system in the original state space. In practice the dimension of the attractor is unknown³, and the embedding dimension is often determined by trial and error, although there are some alternative methods, *e.g.* *Fraser and Swinney* [28] and references thereof. Further details about embedding can be found in *Sauer et al.* [75].

Method of delays The method of delays, or time delay embedding, introduced by *Packard et al.* [58], reconstructs a state vector, \mathbf{Y}_k , in m dimensions from a time series $\{y_k\}, k = 1, \dots, N$, of a single observed value, defined as

$$\mathbf{Y}_k = [y_k, y_{k-\tau_d}, y_{k-2\tau_d}, \dots, y_{k-(m-1)\tau_d}] \quad (2.13)$$

for $k = 1, \dots, K$ where $K = N - m + 1$.

The delay time τ_d , where $\tau_d = j\tau_s$ for some $j \in \mathcal{Z}_0^+$, is fairly arbitrary for sufficiently large N , but should be characteristic of the time over which y varies; the choice of m is more important.

Singular Systems Analysis Singular systems analysis involves forming a trajectory matrix A , whose rows are normalised delay vectors as above, *i.e.* $\frac{1}{\sqrt{K}}\mathbf{Y}_k$. Singular value decomposition, described in [10], may then be applied either to A or to the covariance matrix $C = A^T A$, to give the decomposition $U\Sigma V^T$; the data is projected onto a subset of the orthonormal basis given by the right singular vectors, the columns of V . If the decomposition is of the covariance matrix then this projection optimises

²If the system evolves on an attractor then the trajectory will automatically be stationary [26].

³Indeed, it is not known *a priori* if the system is deterministic.

the variance of the data, and minimises any (white) noise contained in the projection by virtue of the least squares nature of determining the singular values. Further details may be found in Appendix C.

Once a state space of the system has been reconstructed, a model of the system may be built, as described in section 2.2.1.

2.3.2 Analytical systems

2.3.2.1 Logistic map

The logistic map is a 1d map used by biologists as a simple model for seasonally breeding populations in which generations do not overlap, as found in many natural populations including temporal zone aphropods which are crop and orchard pests [49]. The population in a given year, x_{n+1} is related to that of the previous year by

$$x_{n+1} = \lambda x_n(1 - x_n), \quad (2.14)$$

as illustrated in figure 2.12. The model displays chaotic behaviour for many values of

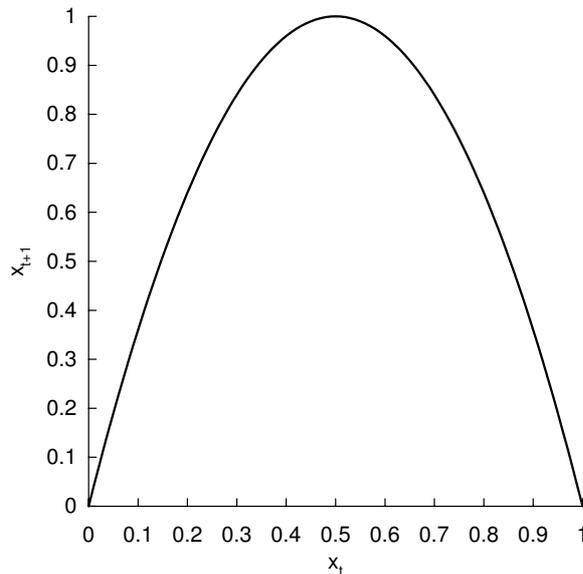


Figure 2.12: Logistic map for $\lambda = 4.0$.

$\lambda \in [3.57, 4]$. For $\lambda = 4.0$ it is possible to transform the logistic map, using a change of variables $x = \sin^2\left(\frac{\pi y}{2}\right)$ where y is defined to be in the interval $[0, 1]$, to give the tent map, defined as

$$y_n + 1 = \begin{cases} 2y_n & 0 \leq y \leq \frac{1}{2} \\ 2 - 2y_n & \frac{1}{2} \leq y \leq 1 \end{cases} \quad (2.15)$$

which can be proven to be chaotic [56].

It is of use to note that, for $\lambda = 4.0$, the invariant measure is given by

$$\rho(x) = \frac{1}{\pi\sqrt{x(1-x)}} \quad (2.16)$$

(see figure 2.13), and so the interval $[0, 1]$ can be divided into K equal probability

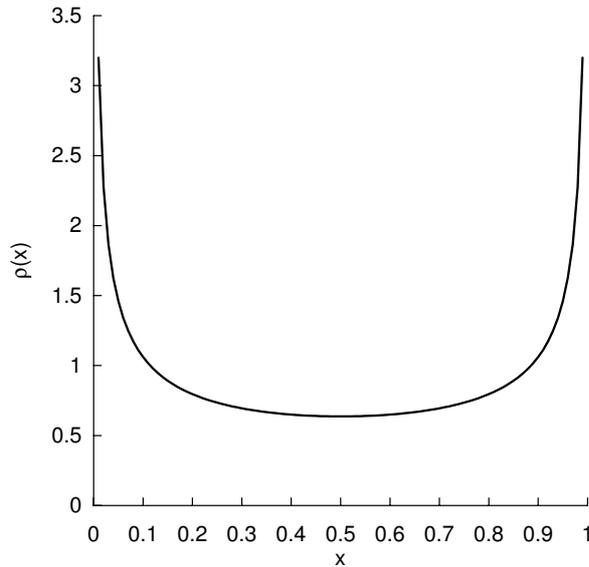


Figure 2.13: Invariant measure of the logistic map for $\lambda = 4.0$.

cells with boundaries $[y_i, y_{i+1}]$ where

$$y_i = \sin^2\left(\frac{i\pi}{2K}\right). \quad (2.17)$$

2.3.2.2 Ikeda map

Consideration of light pulses from a laser travelling round a rectangular system, the corners of which are the location of partially transmitting mirrors, led to the creation of the Ikeda map [39, 34]. If the complex number $z_n = x_n + iy_n$ is taken to represent the amplitude and phase angle of the n th pulse at a given point on the circuit, then the map is given by

$$x_{n+1} = 1 + \mu(x_n \cos(t) - y_n \sin(t)) \quad (2.18)$$

$$y_{n+1} = \mu(x_n \sin(t) + y_n \cos(t)) \quad (2.19)$$

where $t = a - \frac{b}{1+x_i^2+y_i^2}$, $a = 0.4$, $b = 0.6$ and $\mu = 0.83$ are believed to give a chaotic system, whose attractor is shown in figure 2.14. As indicated by the structure of the attractor, a line segment between two points will be rapidly rotated when the points are iterated under the model.

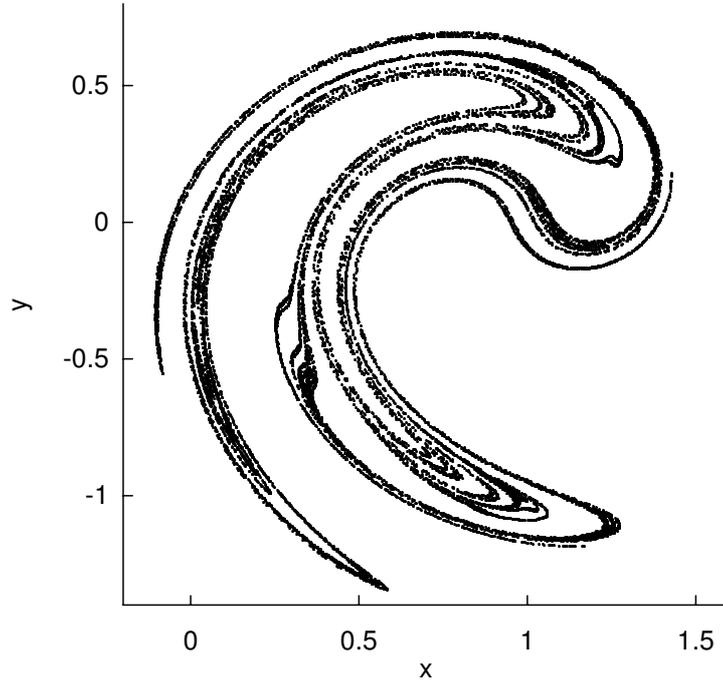


Figure 2.14: Attractor of the Ikeda map for $a = 0.4, b = 6.0$ and $\mu = 0.83$.

2.3.2.3 Sinai map

The Sinai map is another 2d map given by [76, 25]

$$x_{n+1} = (x_n + y_n + \Delta \cos 2\pi y_n) \bmod 1$$

$$y_{n+1} = (x_n + 2y_n) \bmod 1$$

with $\Delta = 0.1$ to give a chaotic system, the attractor for which is illustrated in figure 2.15. It is adopted for investigating the effect of model error in section 4.2 since the trigonometric function may be expanded in a Taylor's series about the origin after a change of variables to translate the map to the region $[-0.5, 0.5] \times [-0.5, 0.5]$. For clarity map is expressed as

$$x'_n = x_n + y_n + 0.5 - \Delta \cos 2\pi y_n \quad (2.20)$$

$$y'_n = x_n + 2y_n + 1 \quad (2.21)$$

$$x_{n+1} = (x'_n + 0.5) \bmod 1 - 0.5 \quad (2.22)$$

$$y_{n+1} = (y'_n + 0.5) \bmod 1 - 0.5 \quad (2.23)$$

Approximate models are then given by

$$x'_n = x_n + y_n + 0.5 - \Delta \text{appcos}_k(2\pi y_n)$$

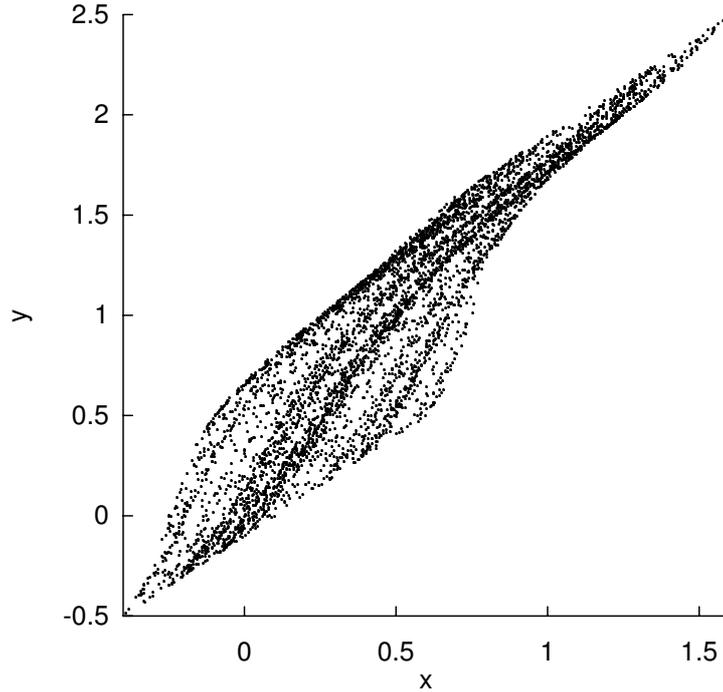


Figure 2.15: Attractor of the Sinai map for $\Delta = 0.1$.

$$y'_n = x_n + 2y_n + 1$$

where $\text{appcos}_k(z) = 1 - z^2/2! + z^4/4! - \dots + (-1)^{2k} z^{2k}/(2k)!$ for a given k , the order of the expansion, and

$$x_{n+1} = (x'_n + 0.5) \bmod 1 - 0.5$$

$$y_{n+1} = (y'_n + 0.5) \bmod 1 - 0.5$$

as above. Note that the approximating function is discontinuous. Two different imperfect models are used, one with a fourth order expansion, the other with a fifth order expansion.

2.3.2.4 Lorenz equations

Having described a number of maps, we now consider a 3d flow described by *Lorenz* in his 1963 paper [45]. The flow of fluid between two rigid horizontal plates subject to gravity and to the two plates being held at different temperatures, with the top plate typically cooler, is known as Rayleigh-Bernard convection. The fluid near the lower plate expands, and buoyancy causes the fluid to rise, while the cooler more dense fluid near the top plate falls. For some ranges of temperature difference between the plates, ΔT , steady convective cellular flow occurs. As ΔT increases the flow becomes chaotic.

If the variations in the fluid are assumed to have two spatial dimensions, then it is possible to substitute double spatial Fourier series (with time dependent coefficients) into the fluid equations governing the flow. This substitution leads to an infinite set of coupled first order o.d.e.s, which may be truncated to give a Galerkin approximation. Saltzman derived a 7 variable system of o.d.e.s describing finite amplitude convection [73], further truncated by *Lorenz* [45] to give the 3d system now known as the *Lorenz equations*:

$$\begin{aligned}\dot{x} &= -\sigma x + \sigma y \\ \dot{y} &= xz + rx - y \\ \dot{z} &= xy - bz\end{aligned}\tag{2.24}$$

where the classical values of parameters are $\sigma = 10, b = 8/3, r = 28$, giving low dimensional chaotic behaviour, and for which a trajectory is illustrated in figure 2.1 on page 8. x is proportional to the circulatory fluid flow velocity, y characterises the temperature difference between rising and falling fluid regions, and z characterises vertical temperature variation; σ and r are proportional to the Prandtl number and the Rayleigh number respectively.

The Lorenz equations are integrated numerically using a fourth order Runge-Kutta scheme, for details see *Press et al.* [65]. The perfect model is taken to be that with a time step of 0.01, while a time step of 0.02 is used to provide an imperfect model. While the Lorenz equations provide a simple flow which exhibits chaotic dynamics and which has much rich behaviour (see *Sparrow* [86] for a full discussion), the symmetry of the system is sometimes undesirable. The Lorenz equations are one of a family of Marzec-Spiegel systems, another asymmetric member of which is described below.

2.3.2.5 Asymmetric Marzec-Spiegel system

The Marzec-Spiegel systems are a family of third order nonlinear ordinary differential equations which may be interpreted as describing the motion of a particle in a time dependent potential described by a polynomial [48]. Here we consider a member of this family which is asymmetric, and which exhibits chaotic behaviour for the given parameter values, namely

$$\begin{aligned}\dot{x} &= y \\ \dot{y} &= -x^3 + \lambda + \delta x \\ \dot{\lambda} &= -\epsilon(\lambda + x - x^3)\end{aligned}\tag{2.25}$$

with $\delta = 0.625$ and $\epsilon = 10^{-3/2}$. The strange attractor is a complex structure, a section of which is given in figure 2.16. Again, the numerical integration method used is a

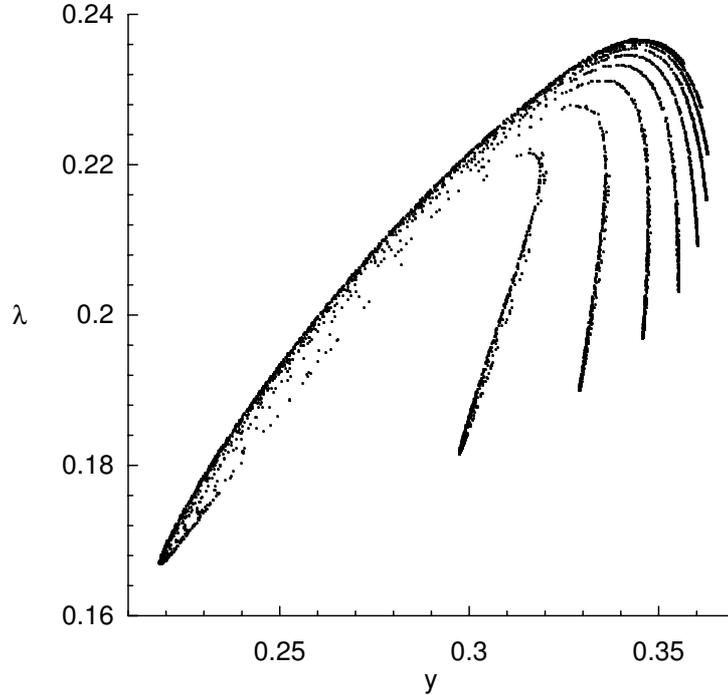


Figure 2.16: The section $x = 0, \dot{x} > 0$ of the strange attractor of a Marzec-Spiegel system in (x, y, λ) space for $\delta = 0.625, \epsilon = 10^{-3/2}$.

Runge-Kutta fourth order method, with a time step of 0.01 taken to define a perfect model.

2.3.3 Laboratory systems: the thermally driven rotating annulus

Analytical systems are useful since knowledge of the system equations enables thorough investigations. We are most interested, however, in the application of methods to observed physical systems for which no perfect model exists. Before addressing models of the Earth's atmosphere, let us consider a laboratory analogy to the planet's mid-latitude large-scale circulation. The circulation may be considered to be driven by two basic processes, as illustrated in the upper panel of figure 2.17. Due to the position of the Earth in relation to the Sun, the atmosphere near the equator is heated more strongly than that near the poles; the rotation of the Earth must also be considered.

These processes are imitated in a laboratory apparatus [36, 29] which consists of two differentially heated, concentric cylinders between which fluid is contained, with insulating boundaries at the top and bottom, all set on a rotating turntable, as

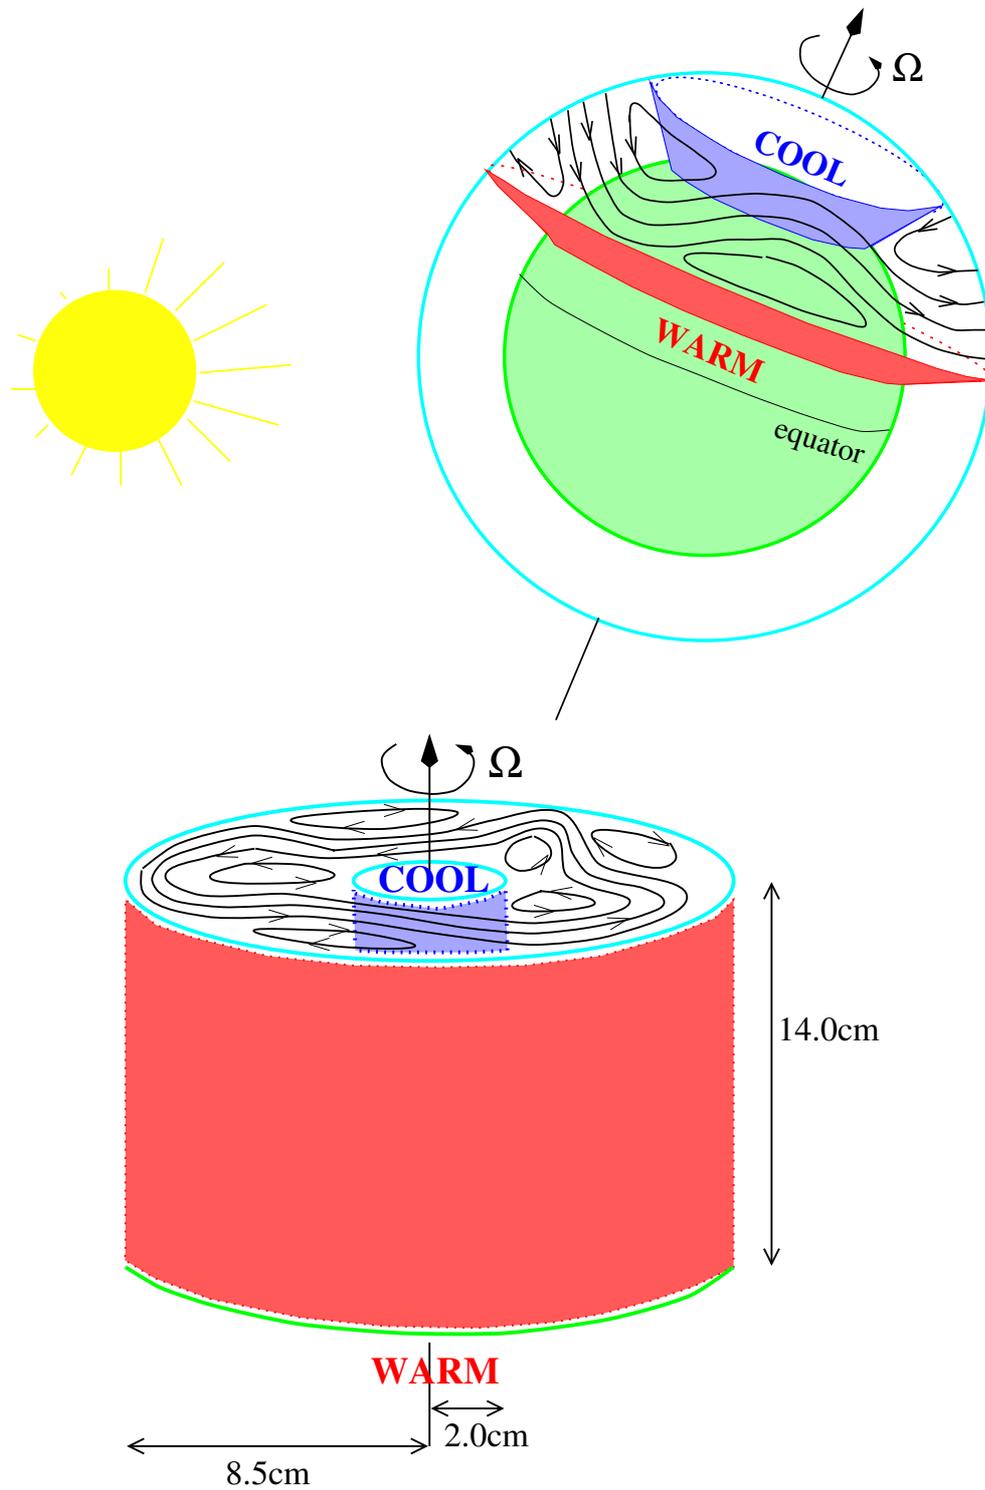


Figure 2.17: Upper panel: The mid-latitude circulation of the Earth's atmosphere may be considered to be driven by the differential heating of the equator and the pole, and by the rotation of the Earth. Lower panel: A laboratory analogy, which comprises of fluid contained between two differentially heated concentric cylinders mounted on a rotating turntable.

illustrated in the lower panel of figure 2.17. The outer cylinder represents a ‘wall’ from the Earth up into the atmosphere near the equator, which is warmer than the equivalent ‘wall’ near the pole. The rotation of the turntable represents the Earth’s rotation. The cylinders are 14.0 cm high, have radii of 2.0 and 8.5 cm respectively, and are held at given temperatures, the inner cylinder being cooler. The temperature of the fluid, a mixture of glycerol and water, is measured by a set of 32 copper-constantan thermocouples uniformly distributed in azimuth at mid-height and mid-radius.

The advantage of the laboratory apparatus is that the parameters, the temperature difference between the cylinders, ΔT , and the rotation rate, Ω , may be varied. For low Ω steady axisymmetric overturning occurs; increasing Ω introduces a second instability causing a wave like disturbance in the azimuthal plane, known as a *baroclinic* instability, which may settle to a spatially regular wave state which is steady (bar rotational drift), periodically varying in amplitude or shape, or varying chaotically in amplitude. If Ω is sufficiently large the flow is spatially irregular and highly chaotic [68]. This information enables construction of a transition diagram akin to the bifurcation diagrams discussed earlier in section 2.1.2. Observational data used in this thesis is from the chaotic regime.

2.3.4 Real world systems: the Earth’s atmosphere

Models of the Earth’s atmosphere used for Numerical Weather Prediction (NWP) are in essence models built from physical principles. It is not the purpose of this thesis to discuss the construction of weather models; a brief description of the NWP models needed to describe later investigations is given below, further details may be found in *e.g.* James [40].

The aim of NWP is to take observed values of atmospheric variables at some initial time and use physical laws to predict meteorological variables at some later time. The physical laws may be separated into the equations of motion, the equation of continuity and the thermodynamic equation; together they form a set of partial differential equations. The dominant terms in the momentum equation are due to the vertical pressure gradient and acceleration due to gravity, and are much larger than the other terms enabling approximation to a hydrostatically balanced atmosphere, which only breaks down at small scales (less than about 10km). The vertical velocity component is significantly less than the horizontal component and vertical motion is further inhibited by stable stratification of the atmosphere and rotation of the Earth. Establishing which are the dominant processes leads to a *primitive set of equations*, which describe a frictionless, adiabatic atmosphere. In order to solve

this set of equations, numerical methods are needed which require discretisation of the system. Finite difference methods using a grid which is regular in longitude and latitude have singularities at the poles; these can be avoided with Fourier filtering, as employed by the UK Meteorological Office. Alternatively *spectral methods* may be used, with quantities represented as Fourier series which are easily truncated as required. Given that the Earth is 3d, and almost spherical, spherical co-ordinates are an obvious choice. Fourier series are replaced with *spherical harmonics* which are products of sinusoidal functions (representing variations in the zonal directions) and Legendre functions (representing variations in the meridional directions); an associated ‘Gaussian’ grid is used (for further details see [40] and [33]).

If the model is to be more realistic then three elements not considered in the primitive equations must be considered, namely radiation, heat exchange between the Earth’s surface and the atmosphere, and the effects of subgridscale processes. The first requires consideration of the effects of clouds, which exist on scales much smaller than that of the grid. Modelling of all three processes is practically achieved by parameterisations of their effects, again described in more detail in *James* [40]. The resulting NWP models are typically of the order of one million dimensions, and are therefore computationally intensive; the resolution at which they are run is largely governed by the time available to produce forecasts and is further constrained by the operational use of ensembles. Initiation of the models from observations is non-trivial due to the sparseness of data (there are typically many more grid points in a region than observations) and the differences between the system state space and the model state space; implications of the latter are discussed in section 5.1.3. In chapter 6 it is shown that current operational ensemble prediction schemes may not be considered to be internally consistent.

2.4 Summary

In Chapter 2 dynamical systems have been introduced, the question of how best to model such a system from data has been discussed and an overview of the systems and models used in this thesis presented.

Systems, defined as continuous time flows or discrete time maps, were shown to exhibit different behaviours as the system parameter(s) are varied. For an analytically tractable system, namely the Moran-Ricker map, the changes in behaviour with variation in parameter were summarised in a new bifurcation diagram. For observed

physical systems this is not usually possible as the parameters can not often be varied. Attractors were introduced to facilitate the definition of Lyapunov exponents, with a positive Lyapunov exponent shown to indicate sensitive dependence on initial conditions, and hence a chaotic system.

After the introduction of a variety of model building techniques, the question of how to determine the optimal order of a model was discussed, and new results were presented comparing information criteria with prediction errors using RBF model of the Moran-Ricker map. Models of orders 1 to 12 were constructed from exact data. For each order of model the in-sample and out-of-sample one step prediction errors, the function fit and the Akaike and Rissanen information criteria (designed to indicate an optimal balance between model order and model accuracy) were calculated. Minima for in-sample errors (in-sample prediction error and the information criteria) occurred at the maximum order of 12. Minima of out-of-sample errors (out-of-sample prediction error and function fit) occurred at order 11. The difference was explained by the in-sample model over-fitting the model to sparse data. The reflection of this over-fitting in the out-of-sample errors, but not the in-sample errors, illustrates the importance of testing a model out-of-sample. The information criteria did not differ significantly from the in-sample prediction error since the penalty of a higher order was more than compensated by the improvement in accuracy of the model. It was concluded that if the data is exact, then the model should be of as high an order as possible providing that the model did not become overly constrained to limited data; the optimal order is indicated by both the out-of-sample prediction error and the function fit.

Models using up to 64 centres were then determined using either exact data or one of 16 noisy realisations. When exact data was used, all error measures considered decreased to minima at order 55; for orders of 56 and above the limitation due to tolerance level used in the SVD calculation was found to prevent further improvement. For each of the 16 noisy models, all errors have minima at \simeq order 30. Above this order, the additional freedom of the model was found to fit the noise rather than the data; the errors showed however, that it was far better to over estimate the best model order than to under estimate it. Again the orders suggested by the information criteria were similar to from those suggested by the in-sample error.

It was noted that calculation of these error measures requires knowledge of exact data which is unavailable for observed physical systems. Chapter 3 will consider alternative measures which may be calculated from observational data. The other systems used in Chapter 3, and throughout this thesis, were also introduced.

Chapter 3

Evaluating erroneous models from uncertain data

For observed physical systems the governing equations are unknown, thus we wish to verify the model as a representation of the system from observations. Further, the lack of knowledge of the governing equations leads to a choice of models; we wish to quantify the optimality of a model as a representation of the system by contrasting the coherence of each model with the observations. As discussed in the previous chapter, the root mean square (r.m.s.) error is often used to quantify optimality as many model construction techniques aim to minimise one step prediction error in the r.m.s. sense. If the model state space is equivalent to that for the system and the data with which the errors are calculated is exact, then the r.m.s. error will show the system equations, *i.e.* the perfect model, to be optimal. Models of physical systems are however often built in a state space inequivalent to the system, using a model structure whose span does not contain the system and with parameters approximated from noisy observations. Evaluating optimality is further hampered by the inexact nature of data available for testing the models.

Before formulating an alternative measure of model optimality, we discuss the complications of chaos and the implications of uncertainty in the observed data, and of the uncertainty in the models. Even when considering the evolution of an analytical system, for which the governing equations are known, approximate numerical integration techniques are often required. Due to the sensitivity to initial conditions demonstrated by chaotic systems, the finite precision of digital computers calls into question the validity of numerical solutions; establishing the existence of a system trajectory close to a numerical solution is addressed by the Anosov-Bowen ϵ -shadowing lemma, introduced below. After illustrating the possible discrepancies of the traditional r.m.s. measure we introduce a new method, constructed in light of the

uncertainties present and data available, to give the ι -shadowing time as a measure of the consistency of a model with the data. Details of the algorithm to determine the ι -shadowing times are contained in Appendix B; an example of how the algorithm searches for the trajectory with the longest ι -shadowing time is given for a simple analytical system. Applications of ι -shadowing, including its use in evaluating model optimality, are discussed in the following chapter.

3.1 Complications of chaos

Implementing numerical integration techniques for chaotic systems is non-trivial due to inherent rounding errors caused by the floating point nature of digital computation. These cause small ‘errors’ in the least significant memory which may rapidly propagate to the more significant bits as a result of chaos, creating large differences depending upon the machine precision used; a simulation may exhibit behaviour not seen in the system.

Since any digital computer is a finite state machine, it must eventually return to a previously visited state, even if the real system does not. Consider a single precision simulation of the Hénon map [35]

$$\begin{aligned}x_{i+1} &= 1 - ax_i^2 + y_i \\y_{i+1} &= bx_i\end{aligned}\tag{3.1}$$

where $a = 1.4$ and $b = 0.3$ so the map is chaotic, and suppose that there exists a numerically computed periodic point, \mathbf{x}_0 say, with period n (n large) [77]. Then if we take this exact value \mathbf{x}_0 and iterate it n times exactly analytically, we find that $\mathbf{f}^n(\mathbf{x}_0) \neq \mathbf{x}_0$ where \mathbf{f}^n represents the iterated Hénon map: the computer simulation has a point \mathbf{x}_0 which exhibits periodic behaviour, with period n , not seen in the actual system. Representing the single precision machine states as a grid, periodicity is equivalent to some n^{th} numerical iterate returning to the grid point of the initial value; figure 3.1 illustrates numerical periodicity without system periodicity.

Realisation of the possible implications of chaos calls into question the validity of simulations; verification of simulations should establish whether the behaviour of the simulation is comparable to that of the system while allowing for small errors. The construction of such a test usually appeals to the *shadowing* work of *Anosov* [4] and *Bowen* [7].

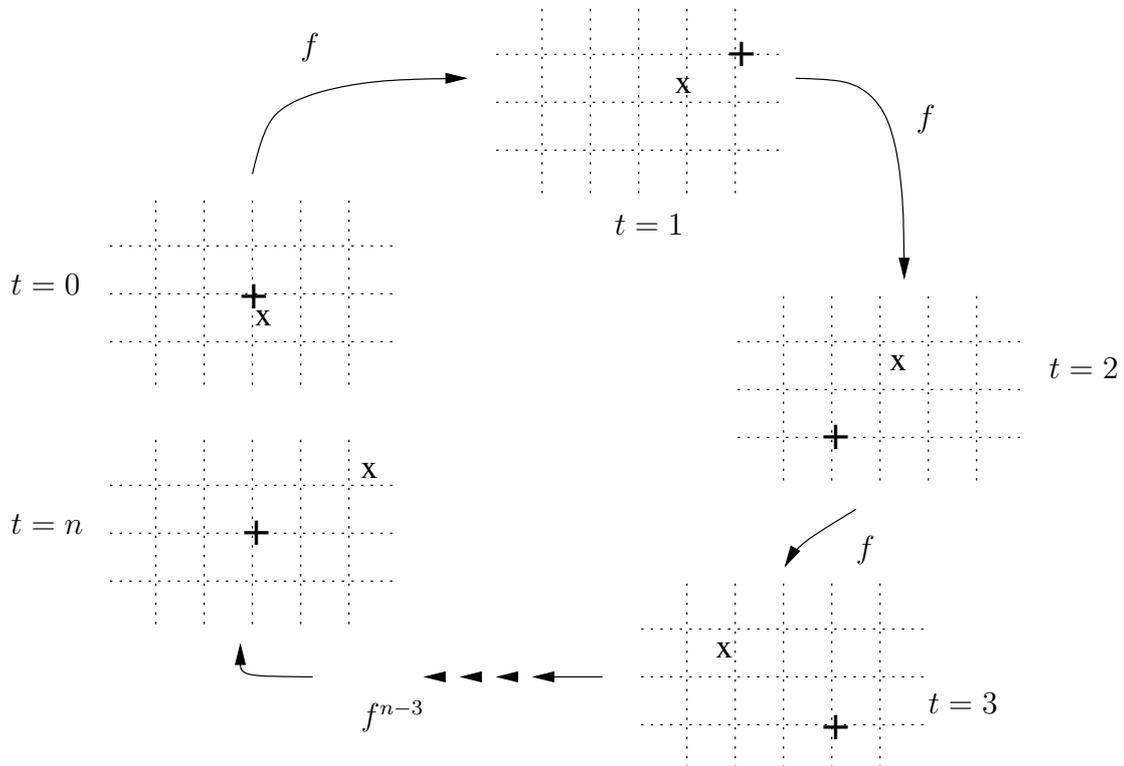


Figure 3.1: A grid representing single precision values, with numerical iterates (+) and exact iterates (x) of an initial value shown for $t = 1, 2, 3, n$. While the initial value shows periodicity, period n , when iterated numerically, exact iteration of the same initial value shows that it is not periodic, period n , under the system.

3.1.1 Anosov-Bowen ϵ -shadowing

Consider a system in which discrete time exact observations are made of a system variable, and suppose that there exists some model approximating the system, *e.g.* a numerical algorithm. Then we have a true trajectory given by the time series of the system variable, and a pseudo-trajectory given by the model. The *Anosov-Bowen shadowing lemma* [4, 7] says that, providing the system satisfies certain given conditions¹, there exists some true trajectory with a slightly different initial condition which ϵ -shadows, *i.e.* stays within a distance ϵ of, the pseudo-trajectory² over some interval $a < t < b$ (see figure 3.2). Note that the true trajectory of the system which shadows the pseudo-trajectory need *not* be generic of the system. For further details and a formal statement of the lemma see Appendix A.

¹Specifically that the system is uniformly hyperbolic and invertible.

²Defining the true trajectory to be ' ϵ -shadowed' by the pseudo-trajectory appears more natural. While neither Anosov [4] or Bowen [7] mention computers, the definition of the Anosov-Bowen shadowing lemma, *e.g.* [16], arises in the verification of computational pseudo-orbits.

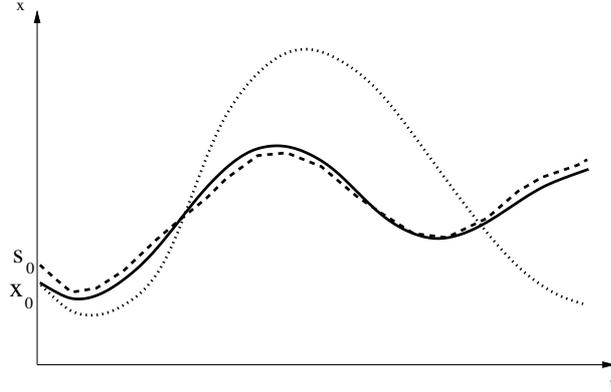


Figure 3.2: Given a pseudo-trajectory (solid line) with initial condition x_0 , there may be some nearby initial condition, s_0 , for which the true trajectory (dashed) stays nearby, or ϵ -shadows, the pseudo-trajectory. Note that the true trajectory with the same initial condition (dotted) will deviate from the pseudo-trajectory.

Shadowing has been employed to overcome the problem of rounding errors in digital computation, characterising global errors in a numerical solution while allowing small perturbations in initial conditions and iterated values. Since it requires knowledge of the system orbit over the time interval in which we are interested it is *not* validation of a model as a predictor; shadowing is a necessary, but not sufficient, condition for the model to succeed as a predictor, *e.g.* for verifying N -body problems in gravitation theory [66].

While few dissipative physical processes satisfy the conditions of the Anosov-Bowen ϵ -shadowing lemma [56], the variation allowed in the initial condition has enabled much work to be done to generalise the implications of the lemma to such systems, *e.g.* van Vleck’s application of ϵ -shadowing to the van der Pol and Lorenz equations [97]; the ‘shadowing time’ and the distance, ϵ , within which the pseudo-trajectory is shadowed are used to quantify the ability of numerically generated solutions [32, 101], with typical shadowing times of the order of 10^4 . If ϵ is of the order of the radius of the system’s attractor, then while the generalisation may be true, it is useless. Even when the generalisations yield a useful value of ϵ , Anosov-Bowen shadowing verifies and contrasts approximations when the theoretical governing equations are known exactly, and the initial conditions are given; the comparison is between the exact solutions and their numerically generated counterparts. For an observed system verifying and contrasting approximations requires comparing the consistency of model solutions with the observational data and uncertainty distribution. Such a comparison is described in section 3.3.2.

3.2 Uncertainty of observation

Theoretical systems are fully described by a set of governing equations and necessary boundary conditions, initial conditions and parameter values; observation does not directly yield any of these details. The available information is measurements of one or more variables, inherent in which is observational uncertainty, at intervals over some period of time. In addition, system uncertainty will arise from approximation of the system by the model. Problems associated with *uncertainty of the model* and *uncertainty in the data* are discussed below. Practically we also encounter a restriction on the amount of observational data available, which may exacerbate difficulties in uniquely determining the system trajectory from the data.

3.2.1 Uncertainty in the data

Ideally the uncertainty in each observation is independently identically distributed. In order to use the observational data effectively we need to establish the uncertainty distribution, which is dependent on the method and accuracy of measurement, and on the level of noise present in the system. Experimental data is often accurate, but *always* to a finite precision. Providing the level of noise in the system is small as compared to the measurement precision, then the error can be classified as quantisation uncertainty³. Otherwise, if the level of noise is larger then the distribution is often assumed to be well approximated as Gaussian.

3.2.1.1 Quantisation uncertainty

If the measurement is a truncated system value and the level of noise in the system is below the level of quantisation, then the uncertainty may be considered to *quantise* the state-space: given an observation \mathbf{y} and a bound on the truncation error ζ , each component of the system variable, \mathbf{s} , has uniformly distributed probability of lying in the hyper-cube $[(\mathbf{y})_l - (\zeta)_l, (\mathbf{y})_l + (\zeta)_l]$, $l = 1, \dots, m$, and negligible probability of lying outside.

The observations and associated uncertainties can therefore be represented as a series of m -dimensional hyper-cubes, or *neighbourhoods of observation*, through which the system trajectory is known to pass, with details of the intermediate sections

³Consider a value of the system of 59.25137..., given an observation of 59.251, if the noise has standard deviation $\sigma = 0.00037$ then noise almost always occurs below the level of quantisation. Conversely, if $\sigma = 0.25$, then the noise is reflected in the observation and the observational uncertainty is no longer solely quantisation error.

unknown. Quantisation uncertainty will arise whenever analogue to digital converters are employed, and is used here for simplicity.

3.2.1.2 Gaussian uncertainty

If an observation is the combination of many measurements then the associated uncertainty is often assumed to have Gaussian distribution, and the mean, μ , and standard deviation, σ , may be approximated from the measurements [72]. A Gaussian distribution is also used to approximate the uncertainty if the level of noise exceeds the measurement precision. Using a set confidence limit we may define neighbourhoods of observation through which the system trajectory has *e.g.* 99% probability to pass, but the distribution within the hypercube defined in this way will not be uniform, in contrast to above.

The techniques presented here for quantisational uncertainty may be generalised to any well specified distribution.

3.2.1.3 Non-uniqueness of observed trajectories

Once the state space is quantised each of the finite number of possible observations, *including any initial condition*, corresponds to an infinite number of possible states of the system: an uncountable number of distinct trajectories of the system may yield identical observational data. Thus a finite time series is unlikely to be sufficient to determine the system trajectory uniquely, even if the governing equations are known. However, for hyperbolic, invertible systems *Farmer and Sidorowich* [27] show that given the governing equations, a unique initial condition, and hence the system trajectory, may be determined to arbitrary precision given a sufficiently long observational time series. Physical systems are not often both hyperbolic and invertible; the Hénon map, logistic map and Lorenz equations each fail to satisfy these conditions. An example of non-uniqueness for the logistic map is given in section 4.1.1.

3.2.2 Uncertainty of the model

The observed system may be a physical realisation of a theoretical system, *e.g.* the double rotor experiment [71], in which case ‘the’ governing equations will be known and used as the model⁴. Usually this is not the case and we need to formulate a model from the observations, the simplest approach being to assume persistence (*i.e.*

⁴Even here, the governing equations will fail to describe the system exactly, since any set of equations summarise a large number of observations; they do not predict external influences, *e.g.* variation in the rigidity of the rotor or frictional heating of a bearing, for single experiments [8].

the current state persisting), followed by a ‘collapse to the climatology’, *e.g.* the mean value of the data, or by some simple periodical variation. Alternatively, a wide range of model construction techniques are available as described in section 2.2.1. For an observed system, the only information available is a set of observations and the system trajectory is unknown; in the following we will therefore refer to system observations and model trajectories.

Regardless of how the model is formulated, it will contain *model error* which is cumulative, *i.e.* given a system \mathbf{f} and an approximating model $\tilde{\mathbf{f}}$ with errors $\boldsymbol{\epsilon}$ then

$$\begin{aligned}\tilde{\mathbf{f}}(\mathbf{x}_0) &= \mathbf{f}(\mathbf{x}_0) + \boldsymbol{\epsilon}_1 \\ \tilde{\mathbf{f}}^2(\mathbf{x}_0) &= \tilde{\mathbf{f}}[\mathbf{f}(\mathbf{x}_0) + \boldsymbol{\epsilon}_1] \\ &= \mathbf{f}^2(\mathbf{x}_0) + \boldsymbol{\epsilon}_2\end{aligned}$$

where $\boldsymbol{\epsilon}_2$ depends on $\tilde{\mathbf{f}}$, $\boldsymbol{\epsilon}_1$ and \mathbf{x}_0 . Model error arises from the approximating nature of the model and should not be confused with *system sensitivity*, *i.e.* sensitive dependence on initial conditions, which is an inherent property of the system (and hopefully also of the model). While large deviations of the model from the observational data are easily discernible by an r.m.s. measure, small but significant model error may go undetected. An alternative method for quantifying model optimality, able to discriminate between model error and system sensitivity, is discussed in the next section.

3.3 Quantifying model optimality

Recall the Ikeda map (section 2.3.2.2), with $\mu = 0.83$ to give a chaotic map, and suppose that we have an observational time series and know the uncertainty distribution to be quantisation uncertainty bounded by 0.05, *i.e.* the observations are accurate to 1 decimal place. The time series is divided into two distinct sets: a learning set and a test set. A radial basis function (RBF) model (see section 2.2.1.2), model α , is constructed from the learning set. The average of the learning set is calculated; model β is given by the RBF model for the first 8 time steps, and thereafter collapses to the average value. The time at which the switch occurs, $t_c = 8$, is the average time at which the r.m.s. error between the observed value iterated under the system equations and the RBF value exceeds the corresponding r.m.s. error for the average value for the system. Trajectories from a chosen point in the test set are iterated under both models, and are illustrated along with the system trajectory in figure 3.3.

Model α gives an oscillating trajectory, while that for model β soon ‘collapses to the climatology of the system’.

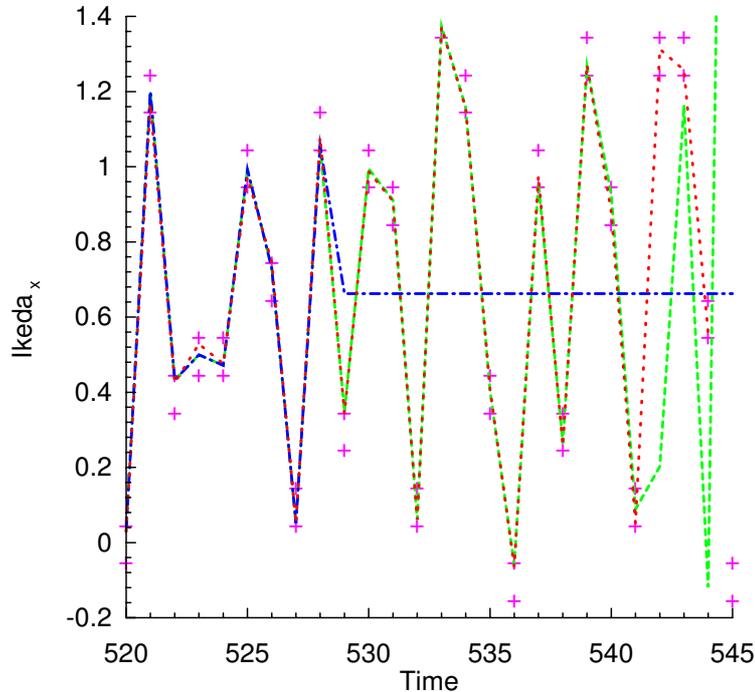


Figure 3.3: Projections onto the x -axis of the system trajectory (red dotted line) and trajectories for model α (green dashed) and model β (blue dot-dashed) for the Ikeda map, with the bounds of the neighbourhoods of observation shown (magenta plus signs). Model α passes through the neighbourhoods of observation for 21 time steps, model β passes through 8.

Given these model trajectories and the observational data and uncertainty distributions, which model is *preferable* as an approximation to the system?

3.3.1 Discrepancies of r.m.s.

A traditional approach is to accept the model which minimises r.m.s. error between observed and predicted values over the time period in which we are interested. In the particular situation above minimising r.m.s. error over time intervals greater than the ‘switch time’, t_c , would choose model β over model α ; over times less than t_c the models would be considered equal.

As discussed above, we *know* the system trajectory to pass through the neighbourhoods of state space at each time step. As soon as a trajectory fails to pass through one of these neighbourhoods it ceases to be consistent with the known information, and can not be the system trajectory. Hence if a model is unable to generate a tra-

jectory consistent with the system observations the model can not be the system, *i.e.* the equations describing the model are not the governing equations of the system.

Constructing neighbourhoods of observation for the Ikeda map, we see that model α gives a trajectory consistent with the information available for 21 time steps, while the trajectory for model β is consistent for 8 time steps; yet model β is ‘preferable’ in the r.m.s. sense. Further, if we had a model γ given by

$$\gamma = \begin{cases} \text{system equations} & 1 \leq t \leq 8 \\ \text{mean value} & t \geq 8 \end{cases} \quad (3.2)$$

then, due to the possible discrepancy between observed and system values arising from observational uncertainty, the r.m.s. measure would reject the system equations in preference for model γ . While this approach may be defended as a good method of prediction⁵, it prefers models inconsistent with the known information, and tends to reject the exact model; definitely not properties desired of the measure of optimality.

3.3.2 ι -shadowing

If a model trajectory exists which stays near the system observations (given the observational uncertainty distribution) for a given time τ_ι , let us say that the model ι -*shadows* the observed system for time τ_ι , for that trajectory; ι -shadowing thus verifies a model trajectory from the observations. Note that ι -shadowing of the observations by a model trajectory is a necessary, but not sufficient, condition for the model to be perfect: the residual errors between the model trajectory and the observations must also have the same distribution as the observational uncertainty. Adopting the notation used to define the Anosov-Bowen shadowing lemma (see Appendix A), we denote the observations of the system $\mathbf{Y} = \{\mathbf{y}_t\}_{t=a}^b$. A trajectory $\mathbf{P} = \{\mathbf{p}_t\}_{t=a}^b$, generated by a model $\tilde{\mathbf{f}}$ approximating the system \mathbf{f} , is said to be a δ trajectory for $a \leq t \leq b$ if

$$|\mathbf{f}(\mathbf{p}_t) - \mathbf{p}_{t+1}| \leq \delta \quad (3.3)$$

for every t , $a \leq t \leq b$. Further, \mathbf{x} , a point in the system, is said to ϵ -shadow this δ trajectory for $a \leq t \leq b$ if

$$|\mathbf{f}^t(\mathbf{x}) - \mathbf{p}_t| \leq \epsilon \quad (3.4)$$

for every $a \leq t \leq b$.

Formally we define ι -*shadowing*:

⁵Indeed, the ‘best’ given an r.m.s. cost function and requiring that the initial observation be used as the initial condition.

Definition 1 (ι -shadowing) Given observations $\mathbf{Y} = \{\mathbf{y}_t\}_{t=0}^{\tau_\iota}$ of a true trajectory of the system, $\mathbf{S} = \{\mathbf{s}_t\}_{t=0}^{\tau_\iota}$, and a model which simulates the system, suppose there exists a δ trajectory $\mathbf{P}_\delta = \{\mathbf{p}_t\}_{t=0}^{\tau_\iota}$ given by the model. If the model is ϵ -shadowed by the observational time series (using an appropriate norm given the observational uncertainty), then the model is said to ι -shadow the observed system within a distance ϵ for time τ_ι .

For quantised observational uncertainty considered here, the appropriate norm is the supremum norm, *i.e.* $\|\cdot\|_{L_\infty}$.

A model may be considered preferable if it is consistent with data available; a *consistent model* of the system is one which ι -shadows within the bounds of the quantised observational uncertainty distribution for all time and for all trajectories emanating from the domain of the system. If the only information known about the system comes from observation, then it will only be possible to test whether or not a model is consistent for observable trajectories; unless the system is in a transient state these will lie on the attractor of the system.

Lemma 1 Given observations $\mathbf{Y} = \{\mathbf{y}_t\}_{t=0}^{\tau_\iota}$ of a true trajectory of a system, a consistent model $\tilde{\mathbf{f}}$, which ι -shadows the system for time τ_ι within the magnitude of the quantisation ζ , is indistinguishable from an exact model for time τ_ι ; it is consistent with all available information for $t \in [0, \tau_\iota]$.

Proof A consistent model ι -shadows the observational system within ζ for time τ_ι ; the model has a δ trajectory $\mathbf{P}_\delta = \{\mathbf{p}_t\}_{t=0}^{\tau_\iota}$, where $\mathbf{p}_t = \tilde{\mathbf{f}}^t(\mathbf{x}_0)$, for some \mathbf{x}_0 in the domain of the model and $\delta > 0$, which is ϵ -shadowed by \mathbf{Y} with $\epsilon = \zeta$. From the definition of ϵ -shadowing (see Appendix A)

$$\begin{aligned} \|\mathbf{y}_t - \mathbf{p}_t\|_{L_\infty} &\leq \zeta & \forall t \in [0, \tau_\iota] \\ \Rightarrow (\mathbf{p}_t)_l &\in [(\mathbf{y}_t)_l - \zeta, (\mathbf{y}_t)_l + \zeta] & \forall l \in [1, m], t \in [0, \tau_\iota] \end{aligned} \quad (3.5)$$

the model trajectory is consistent with all known information, as required.

For finite trajectories there may be more than one consistent model indistinguishable from the observations; if a stable manifold falls within a neighbourhood of observation then multiple model trajectories will exist which converge within the neighbourhoods of observation in evolution and are therefore indistinguishable for some finite time. Non-uniqueness of trajectories is discussed in sections 3.2.1 and 4.1.1, with an example in the latter.

Applying the definition of ι -shadowing to models of the Ikeda map for the trajectory illustrated above in figure 3.3, model α ι -shadows for 21 time steps, model

β for 8 time steps. While neither model α or β are optimal, the ι -shadowing time *gives a measure of the ability of a model to characterise the system* and quantifies the optimality of a model; in the ι -shadowing sense, model α is preferable to model β . We need some measure against which to compare the ι -shadowing time however, since in the case of an exact model the ι -shadowing time is unlimited.

All of the models used here are implemented using floating point digital computations. The accuracy of the initial condition is limited by the machine precision and so we may consider an average time taken for error due to finite machine precision to grow to the scale of the observational uncertainty to give a suitable measure against which to compare the ι -shadowing time. It is not suggested that this time has any implications for the possible duration of shadowing trajectories, but rather that it provides a measure of rate of error growth for a system. The largest global Lyapunov exponent, and its limitations, are discussed in section 2.1.3; it provides a measure of the average error growth rate for a system for an infinitesimal perturbation. We wish to calculate an average time for error due to finite machine precision to grow to the scale of the observational uncertainty. Note that the machine precision, ϵ_c , is finite, so the Lyapunov exponent, Λ , will not give the exact growth rate. Given that the observational uncertainty is bounded by a magnitude ζ , we should compare ι -shadowing times against a time τ_e given by

$$\begin{aligned}\zeta &= \epsilon_c \exp\{\Lambda\tau_e\} \\ \Rightarrow \tau_e &= \frac{1}{\Lambda} \log_e \left(\frac{\zeta}{\epsilon_c} \right).\end{aligned}\tag{3.6}$$

As stated previously, observational uncertainty is present in *all* observations, hence even in the case of a perfect model the observed initial condition will, with $P(1)$, *not* give the optimal ι -shadowing time; the set of optimal initial conditions (those with optimal ι -shadowing times) in the neighbourhood of observation has zero measure. Finding an optimal initial condition within the initial neighbourhood of observation for an imperfect model is a non-trivial problem; full details of the algorithm are given in Appendix B and a demonstration of its implementation for a simple analytical system is given below in section 3.4.

Averaging over the system The rate of uncertainty growth varies with the state, therefore any measure of model optimality must be averaged over the attractor in some way. The simplest approach is to take a random selection of observations as observed initial conditions, calculate the ι -shadowing times, and take the average of

the optimality measure over the set of trajectories. By taking the random selection from observations, they will tend to be distributed according to the natural measure as the duration over which the observations are distributed becomes infinite.

3.4 Demonstration of the ι -shadowing algorithm

Determining the ι -shadowing time distribution of a model for a given observed trajectory requires finding the initial condition $\mathbf{x}_0 \in \Xi_0$, where Ξ_n denotes the neighbourhood containing all states consistent with the observation at time $t = n\tau_s$, with associated ι -shadowing time, τ_ι , which maximises

$$\min_{\tau_\iota} \left\{ \tau_\iota \geq 0 \text{ s.t. } \tilde{\mathbf{f}}^{\tau_\iota}(\mathbf{x}_0) \notin \Xi_{\tau_\iota} \right\}. \quad (3.7)$$

In other words, the ι -shadowing time is the maximum time for which a model trajectory passes through each neighbourhood of observation. If the observational uncertainty has Gaussian distribution then the ι -shadowing time is the maximum time for which a trajectory both passes through each neighbourhood of observation *and* for which the errors are not biased. Below is an illustration of the refinement procedure, explained in detail in Appendix B, for a simple analytical system.

Given an observed initial condition, it is iterated under the model until the trajectory it yields is inconsistent with the observations and their associated uncertainty, at which point the trajectory is said to fail to ι -shadow. At the time of ι -shadowing failure, the difference between the evolved trajectory and the observation at that time may be considered as an approximation to the error in the model trajectory at that time. The algorithm to find the optimal initial condition consists of a method of linear feedback (which is essentially Newton's method), which utilises the linear propagator to estimate the correction needed in the initial condition to achieve the desired perturbation at the time of ι -shadowing failure. The algorithm incorporates a random search method used when the linear dynamics are not a good approximation for the magnitude of perturbations considered.

The method of linear feedback is illustrated by considering a simple 2d saddle point system for which the governing equations are known, and the Jacobian is known exactly. The model is taken to be the system equations, and a time step of 0.01 is used with a fourth order Runge-Kutta numerical integration scheme. The system is described by the pair of ordinary differential equations:

$$\begin{aligned} \dot{x} &= ax + by \\ \dot{y} &= bx + ay \end{aligned} \quad a, b > 0$$

with eigenvalues and corresponding eigenvectors

$$\lambda_1 = a + b \quad , \quad v_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

$$\lambda_2 = a - b \quad , \quad v_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

which will be unstable and stable directions respectively, providing $b > a$, as shown in figure 3.4.

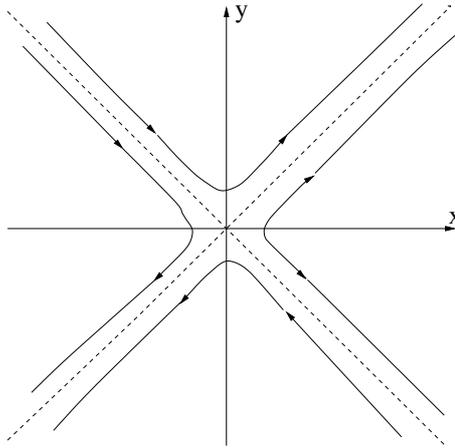


Figure 3.4: Typical trajectories for the saddle system, showing stable and unstable directions.

Given

$$\mathbf{f}(\mathbf{x}) = \begin{pmatrix} a & b \\ b & a \end{pmatrix} \mathbf{x}$$

then the Jacobian is

$$\mathcal{J}(\mathbf{x}_i) = \begin{pmatrix} a & b \\ b & a \end{pmatrix},$$

i.e. the Jacobian is independent of position. Since the system describes a flow, the tangent linear propagator over ν time steps is given by

$$\begin{aligned} \mathcal{M}(\mathbf{x}_0, \nu) &\simeq \left\{ \prod_{i=0}^{\nu-1} (1 + \mathcal{J}\delta t) \right\} \\ &= \begin{pmatrix} 1 + a\delta t & b\delta t \\ b\delta t & 1 + a\delta t \end{pmatrix}^{\nu} \end{aligned}$$

where $\delta t = 0.01$, assumed to be sufficiently small for the approximation to hold.

A system initial condition of $(-0.1896, 0.2004)$ is chosen with $a = 2, b = 4$, and iterated to give a system trajectory. ‘Observational’ data is obtained by rounding the system data to 3 decimal places, and therefore has a quantisation uncertainty

distribution with a maximum amplitude of 5×10^{-4} in each component. The initial observation of $(-0.190, 0.200)$ is taken as a trial initial condition to give a trajectory, which fails to ι -shadow at the first time step, as illustrated in figure 3.5. The mea-

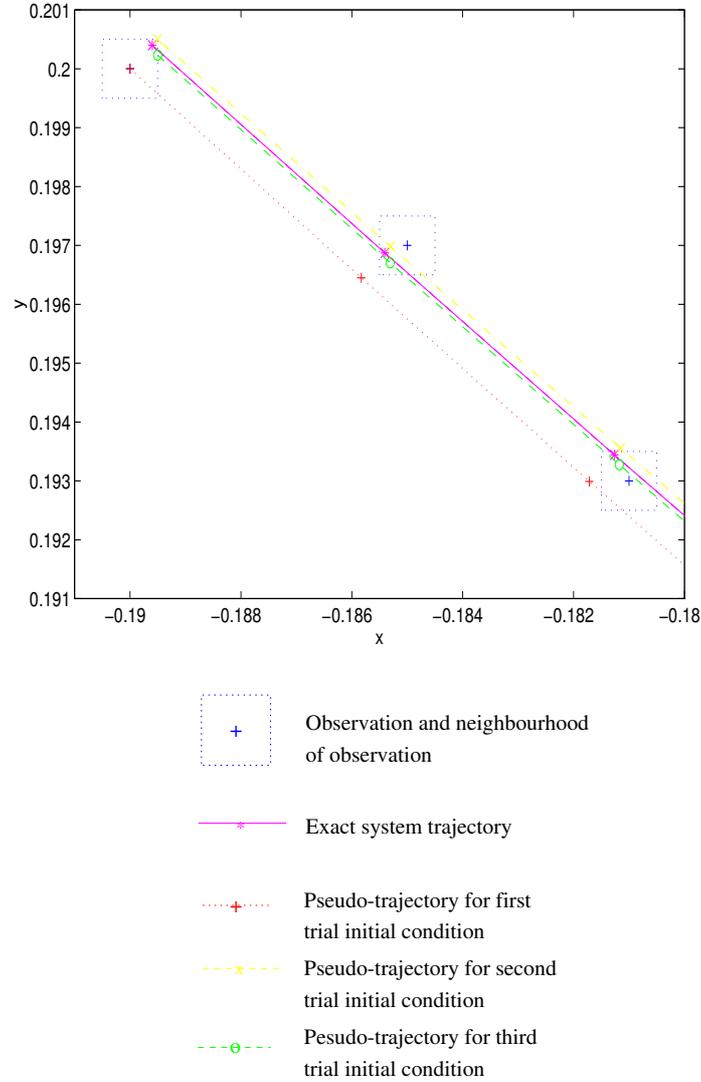


Figure 3.5: Illustration of localised linear feedback implemented for the saddle system. The initial observation has a model trajectory which fails to ι -shadow at the first time step. The error at time $t = 1$ is fed back to give a revised initial condition, limited to be within the neighbourhood of observation, which passes closer to the system value at time $t = 1$, as intended, but fails to ι -shadow at time $t = 2$. Again the error is fed back, and gives a trial initial condition whose pseudo-trajectory passes close to the system value at times $t = 1, 2$.

asurable error between the evolved model trajectory and the observation at this time, ϵ_1 , is calculated to be

$$\epsilon_1 = \begin{pmatrix} 8.296 \\ 5.521 \end{pmatrix} \cdot 10^{-4}.$$

In order to suggest a perturbation to the initial condition to correct this error we invert the tangent linear propagator, by means of singular value decomposition (SVD) and application of the Moore-Penrose pseudo-inverse, to find the perturbation Δ_0 such that $\mathcal{M}_\nu \Delta_0 \simeq \epsilon_\nu$ (see Appendix C). Firstly, we apply SVD to \mathcal{M} to obtain

$$\begin{aligned} \mathcal{M} &= U\Sigma V^T \\ &= \begin{pmatrix} 0.707 & -0.707 \\ 0.707 & 0.707 \end{pmatrix} \begin{pmatrix} 1.06 & 0 \\ 0 & 0.98 \end{pmatrix} \begin{pmatrix} 0.707 & 0.707 \\ -0.707 & 0.707 \end{pmatrix}. \end{aligned}$$

Defining $E = U^T \epsilon_1$, it is calculated to be

$$E = \begin{pmatrix} 9.770 \\ -1.962 \end{pmatrix} \cdot 10^{-4}.$$

In order to apply the Moore-Penrose pseudo-inverse we calculate $\mathcal{K} = \frac{E}{\sigma}$ where σ is taken to be the singular value providing that it exceeds some minimum threshold, here

$$\mathcal{K} = \begin{pmatrix} \mathcal{K}_1 \\ \mathcal{K}_2 \end{pmatrix} = \begin{pmatrix} 9.217 \\ -2.002 \end{pmatrix} \cdot 10^{-4}.$$

Then we take $\Delta_0 = V\mathcal{K}$ so that

$$\begin{aligned} \mathcal{M}_\nu \Delta_0 &= U_\nu \Sigma_\nu V_\nu^T V_\nu \mathcal{K}_\nu \\ &\simeq U\Sigma V^T V \Sigma^{-1} E_\nu \\ &\simeq U\Sigma V^T V \Sigma^{-1} U^T \epsilon_\nu \\ &\simeq \epsilon_\nu. \end{aligned}$$

Here

$$\Delta_0 = \begin{pmatrix} 7.933 \\ 5.101 \end{pmatrix} \cdot 10^{-4}$$

giving a new trial initial condition,

$$\mathbf{x}'_0 = y_0 + \Delta_0 = \begin{pmatrix} -0.18921 \\ 0.20051 \end{pmatrix},$$

limited to $\begin{pmatrix} -0.18950 \\ 0.20050 \end{pmatrix}$ to fall within the neighbourhood of observation. This initial condition is then iterated and gives a value at the first time step near to the observed value, as intended. The revised initial condition is iterated under the model and found to fail after 2 time steps (see figure 3.5). This process is repeated, with the next refinement ι -shadowing for 19 time steps, until ι -shadowing is achieved for the entire trajectory; successive approximations to the initial condition and their corresponding trajectories are illustrated below in figure 3.6.

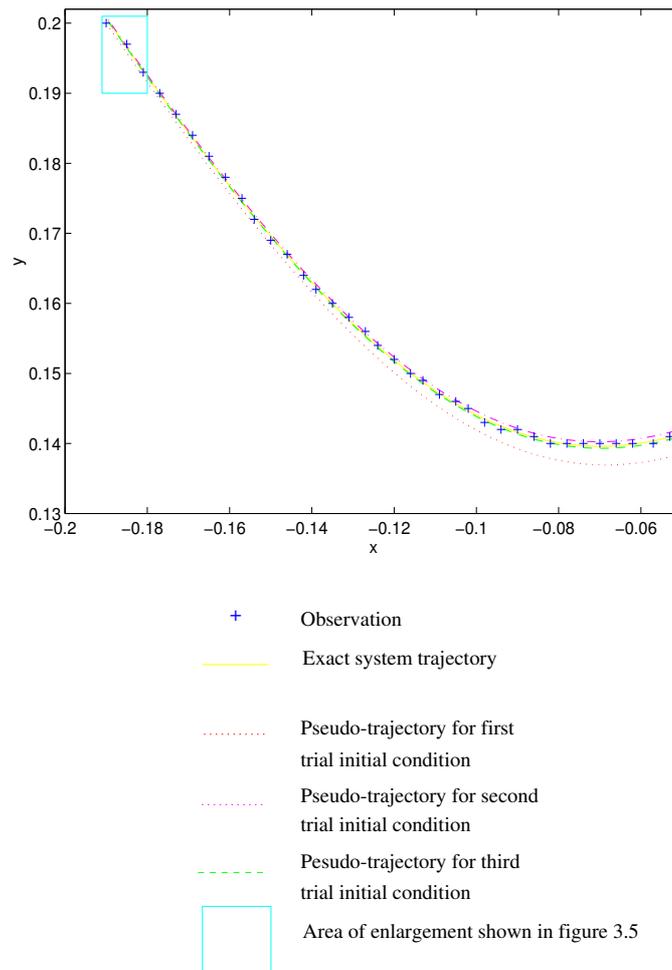


Figure 3.6: An exact trajectory, system values and observed values for the saddle system, with a succession of trial initial conditions and their model trajectories which ι -shadow for increasing durations. Figure 3.5 is an enlargement of the region indicated.

3.5 Summary

In Chapter 3 we have considered how to compare imperfect models of a system from observations. The complications of implementing numerical integrations of a chaotic system on finite precision computers, and the implications of uncertainty in both the data and the models were discussed. The traditional root mean square measure was considered in light of these discussions and shown to reject the exact system in favour of an imperfect model; an alternative measure to quantify model optimality, the ι -shadowing time, was introduced. The two measures were compared using imperfect models of the Ikeda map, highlighting the suitability of ι -shadowing times as a model evaluation measure, and the algorithm used to determine the ι -shadowing time was illustrated for a simple saddle system.

It was noted that, while models may be compared using prediction error measures when the exact values are known (as in Chapter 2), physical systems are usually observed, necessitating the use of an alternative measure of model evaluation. For observational physical systems, the concern is whether an imperfect model is able to display behaviour seen in the system. Quantisational observational uncertainty was introduced and taken as standard (since it will arise whenever analogue to digital converters are employed); neighbourhoods of observation, through which the system is known to pass at each time step, were then defined. Model error, arising from imperfect model structure and estimation of parameters from inexact observations, was contrasted with system sensitivity and it was stated that a measure of model optimality should distinguish between the two.

Two imperfect models of the Ikeda map were then defined and evaluated in terms of r.m.s. error. Given that the system is known to pass through the neighbourhoods of observation of the system, it was argued that if a model is able to characterise the behaviour of the system, then the model will have a trajectory which passes through all the neighbourhoods of observation (and will hence be consistent with the known information). This motivated the definition of an ι -shadowing time of a model for a trajectory as the maximum time for which there exists a trajectory of the model which is contained within the neighbourhoods of observation. Illustration of ι -shadowing times as a model evaluation measure for the imperfect models of the Ikeda map demonstrated that ι -shadowing exploits (model) system sensitivity while penalising model error and that, consequently, the ι -shadowing time distribution will always show a perfect model to be optimal. This was in contrast to the r.m.s. measure which was shown to reject the exact system in favour of an imperfect model. The algorithm

used to determine the optimal initial condition (with the maximum ι -shadowing time) was demonstrated to work using a simple saddle system. Applications of ι -shadowing are the topic of the next chapter.

Chapter 4

Evaluating models: enlightening shadows

ι -shadowing provides an alternative approach to evaluating models which better distinguishes between model sensitivity and model error. One motivation for defining an alternative to the traditional root mean square error was to quantify the ability of imperfect models to represent a given system from observations. In this chapter the distribution of ι -shadowing times are compared for various low-dimensional models of maps and discussed in light of the information available, which is considerable since the perfect models are known. When an imperfect model is used to try and shadow observations of the Lorenz system, the regions in which shadowing fails are found to be correlated to regions of rapid error growth. A systematic study of the relation between model error and regions of shadowing failure is presented for the Sinai map, and the results applied to improve a radial basis function (RBF) model of the rotating fluid annulus. Other applications of ι -shadowing are discussed briefly. The chapter concludes with a discussion of the limitations to improvements possible by employing ι -shadowing; specifically uncertainty will still remain in the initial condition, advocating the use of ensembles.

4.1 Quantifying model optimality

Three low-dimensional systems are considered and both the choice between direct or iterative models and the distribution of base points about which the model is defined (see section 2.2 for details) are discussed. Polynomial models of the logistic map are considered; quadratic maps can fit the logistic map exactly, enabling the comparison of ι -shadowing time distributions of perfect and imperfect model scenarios. No (finite order) polynomial model fits the Moran-Ricker map perfectly; the performance of

imperfect models is discussed with respect to the estimated maximum shadowing time. For 1d maps, finding the ι -shadowing trajectories is usually straightforward; for the Lorenz equations it is shown that the method used does affect the results.

In contrast to observational noise, dynamical noise affects the state of the system and hence the future trajectory of the system. We investigate how well the Lorenz equations shadow observations generated by adding stochastic dynamical noise at regular intervals to a numerical integration of the Lorenz equations; this is equivalent to trying to shadow observations of a high (arguably infinite) dimensional system with a much lower dimensional model. The whole of this chapter represents new research.

4.1.1 Logistic map

Observational data is generated by iterating the logistic map of section 2.3.2.1

$$x_{t+1} = \lambda x_t(1 - x_t) \quad (2.14)$$

and observing the trajectory to three decimal places. Discontinuous segments of 100 observations are then considered as separate sets of system observations. If the model is taken to be the exact logistic equation, *i.e.* equation 2.14 with $\lambda = 4.0$, then, for most sets, an initial condition is found whose trajectory ι -shadows the observations for the duration of the set. (The method used to find the initial condition is that of interval convergence, as described in section B.1.) For a small number of sets, however, the ι -shadowing times remain limited and far from optimal.

Often, the shadowing times are limited because the algorithm stops due to splitting of the region, *i.e.* given a region which ι -shadows for time τ_1 we refine the region, set up a new ensemble of points and find that the ensemble members which have model trajectories which ι -shadow for time $\tau_2 > \tau_1$ form disconnected subintervals of the region. On refining these subintervals further, it is possible to find two trajectories which only differ significantly at one particular time and both of which ι -shadow for a long time. This raises the possibility of distinct trajectories being indistinguishable from the system trajectory even in the limit of infinitely long observational time series. Such non-uniqueness of initial conditions giving trajectories consistent with observational data for infinite time is in contrast with the result for hyperbolic, invertible systems noted in section 3.2.1.3.

Turning to local polynomial interpolation models introduces a choice of the number, n , of base points about which the model is defined, a choice of polynomial and a choice of direct or iterative model (*i.e.* whether a model is built to directly approximate the, say, p^{th} iteration of the map or whether a model is built to approximate the

map and then applied iteratively p times). There is a further choice of whether the base points are equally spaced, or spaced according to the invariant measure. Note that quadratic iterative models can fit the logistic map system exactly.

Perhaps the simplest models are direct linear interpolation models with equally spaced base points; increasing the number of base points increases the ι -shadowing time distributions as shown in figure 4.1. This is to be expected since increasing the

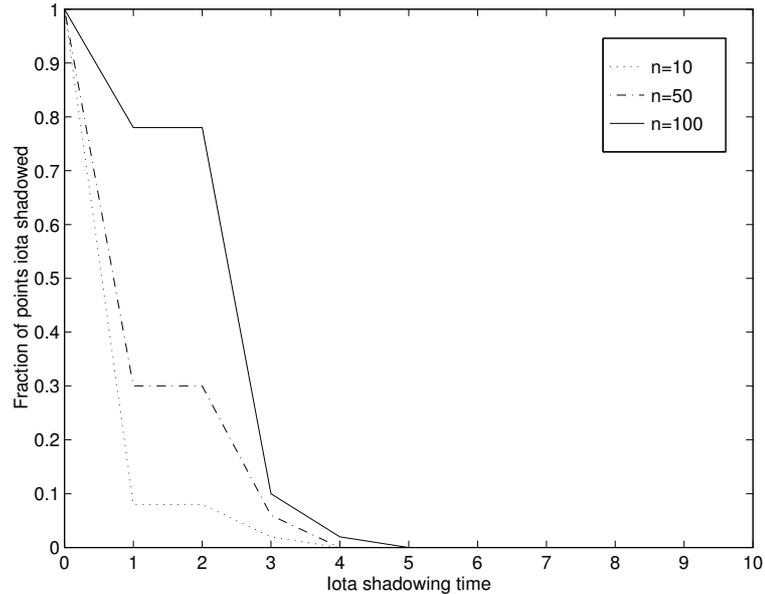


Figure 4.1: ι -shadowing results for the logistic map using direct linear interpolation models with equally spaced base points and observations correct to three decimal places. The fraction of initial conditions ι -shadowed is plotted as a function of time for varying numbers, n , of base points: $n = 10$ (dotted line), $n = 50$ (dash-dotted) and $n = 100$ (solid).

number of base points increases the accuracy with which the direct map is estimated.

It is expected that iterative models will out perform their direct equivalents since linear and quadratic approximations are able to fit the (one step) quadratic system much more accurately than the complex structure of multiple step iterations of the system [26]. This is reflected in the ι -shadowing times, with iterative models ι -shadowing for longer, see figure 4.2 for a comparison of linear direct and iterative models with equally spaced base points.

Only when a model's shadowing time distribution lies clearly to the right of and above another is it an unambiguously better model. From figure 4.2 we can deduce that 100 base point linear iterative models are superior to direct models with the same number of base points, but no such clear distinction can be made between a direct linear model with 100 base points and an iterative linear model with 10 base

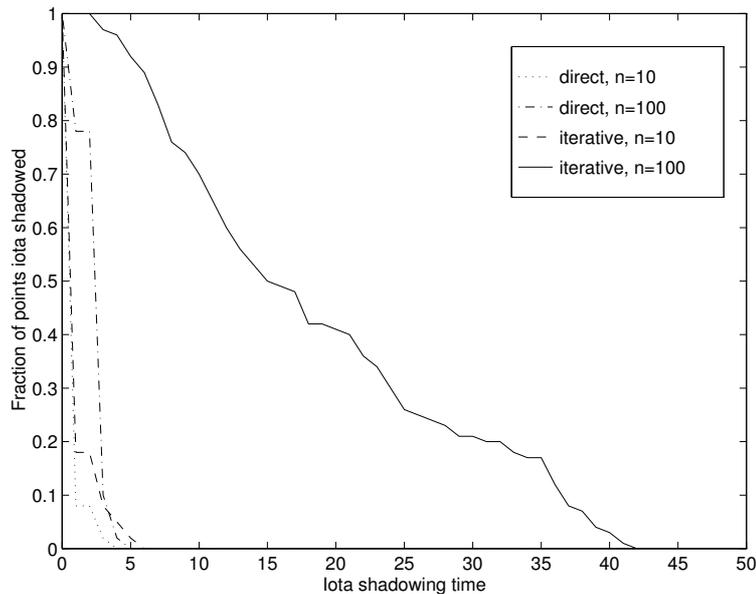


Figure 4.2: ι -shadowing results for the logistic map using linear interpolation models with varying numbers of equally spaced base points, comparing direct and iterative models.

points. Similar results (not shown) are found for quadratic models and using base points distributed according to the invariant measure.

The best distribution of base points depends on the ability of the model as a predictor. For small n , the models are not very accurate, ι -shadowing times are short and so model trajectories will not explore much of the model domain; it is sensible to distribute base points in the most frequently visited regions. As n increases, the models become more accurate, ι -shadow for longer and trajectories begin to visit less frequented regions; the density of base points distributed according as the invariant measure is sparse in less frequented regions, making equally distributed base points preferable. The ideal distribution of base points will be a combination of points which are equally spaced and those which are spaced according as the invariant measure; determining the ideal distribution of base points using ι -shadowing times is illustrated for a model of the annulus in section 4.3. These results are not affected by whether the model is direct or iterative, as illustrated for linear models in figure 4.3.

As shown in figure 4.4, quadratic interpolation models ι -shadow longer than their linear counterparts. Quadratic iterative models have ι -shadowing times far longer than any other model, expected since the logistic map is quadratic, and the quadratic iterative models are limited only by ‘cumulative computational error’. As discussed in section 3.3.2, one estimate of the ‘average’ time for errors due to the finite machine

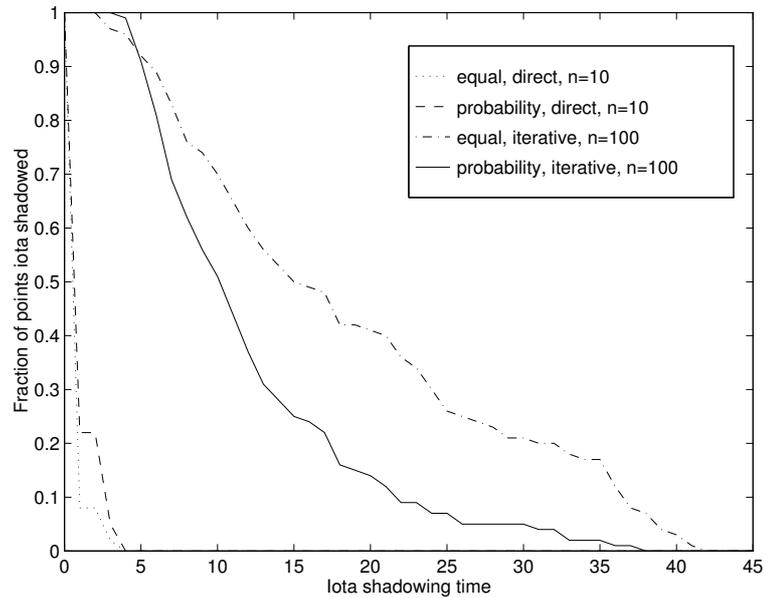


Figure 4.3: ι -shadowing results for the logistic map using linear interpolation models, comparing models with equally spaced base points with those with base points distributed according as the invariant measure (probability).

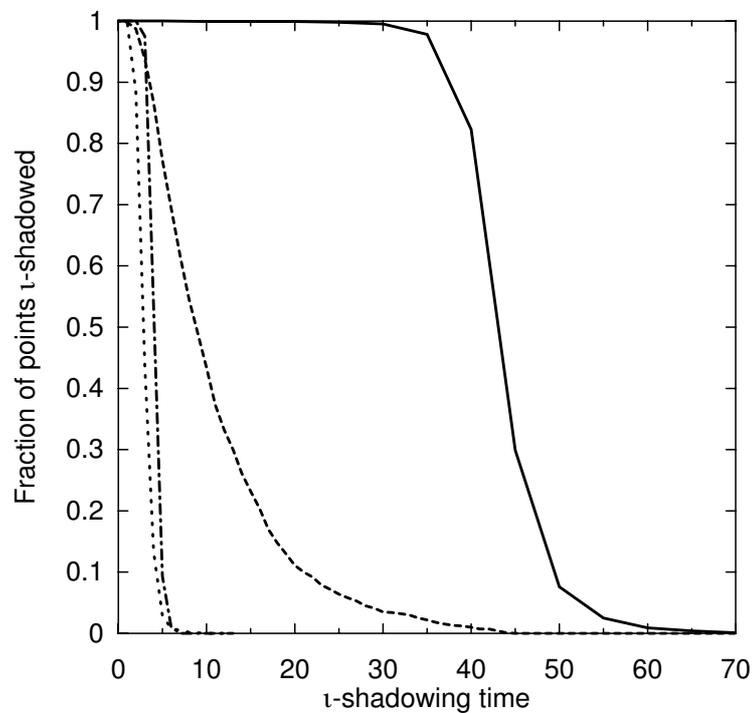


Figure 4.4: ι -shadowing results for the logistic map using models with 100 base points spaced according to the invariant measure, comparing direct linear (dotted line), iterative linear (dash-dotted), direct quadratic (dashed) and iterative quadratic (solid) models.

precision to grow to the magnitude of the observational uncertainty is given by¹

$$\tau_e = \frac{1}{\Lambda} \log_e \left(\frac{\zeta}{\epsilon_c} \right) \quad (3.6).$$

For the logistic map the Lyapunov exponent is $\log_e 2$ [3], the machine precision is found to be 1.110×10^{-16} and the observational uncertainty is 5×10^{-4} ; τ_e is calculated to be about 42. Since the finite time Lyapunov exponent for 42 time steps is similar to the global Lyapunov exponent, we may expect τ_e to be a relevant estimate in this case. From figure 4.4 we can see that $\tau_e = 42$ coincides with the sharp drop, between 40 and 45 time steps, of the number of initial conditions which can be shadowed using a quadratic iterative model: over 80% of initial conditions can be shadowed for 40 time steps while less than 30% can be shadowed for 45 time steps.

4.1.2 Moran-Ricker map

The Moran-Ricker map (see figure 2.2) is given by

$$x_{t+1} = x_t \exp\{r(1 - x_t)\} \quad (2.3),$$

and has chaotic behaviour for $r = 3.7$, used here. While it is similar to the logistic

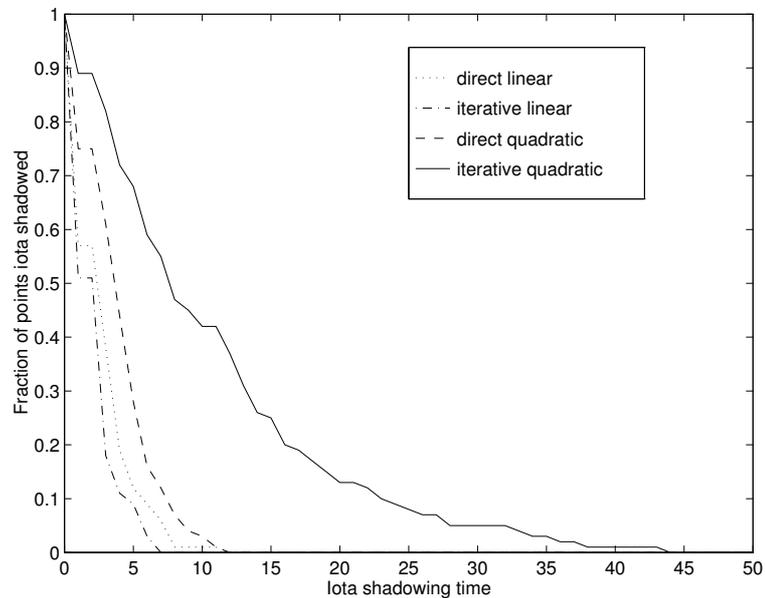


Figure 4.5: ι -shadowing results for the Moran-Ricker map, comparing direct and iterative, linear and quadratic models using 400 base points.

map, being exponential it does not possess the symmetry of the logistic map and

¹The limitations of τ_e as such an estimate are also discussed in section 3.3.2.

most methods of prediction which fit the logistic map well due to its quadratic nature, fail with the Moran-Ricker map. The invariant measure is not known analytically, hence base points for the models are equally spaced. It is found that there is much less occurrence of splitting of optimal regions when the model is taken as the exact system equation and the method of interval convergence used. Relative performance of local linear and quadratic, direct and iterative, models reflects that for the logistic map.

In contrast to the logistic map, a direct piecewise linear model out performs its iterative counterpart, as shown in figure 4.5. For the logistic map the local linear fit to the map is fairly accurate, and iterating the one step map is more accurate than a direct fit to the iterated map which rapidly becomes complicated. The more complex nature of the Moran-Ricker map leads to a higher degree of error in the one step map, which propagates on repeated application. In comparison to the logistic map, the iterated Moran-Ricker map is not significantly more complex than the one step map, and hence a local linear fit directly to the iterated map is more accurate than iterating a fit to the one step map. ι -shadowing thus highlights a result not immediately evident from the system equations, but not unexpected. There is no dramatic increase in ι -shadowing times for the quadratic iterative model, since we are no longer dealing with a quadratic map, but it is still the best model out of those considered.

The Lyapunov exponent is estimated to be 0.67, giving a value of τ_e of about 43. Only the quadratic iterative model ι -shadows for that length of time, and then only for a few trajectories. This suggests that the averaged rate of error growth for the logistic map and the Moran-Ricker map are similar. This suggests that the quadratic iterative model of the Moran-Ricker map really is worse representation of the system than the equivalent model is of the logistic map (rather than disparate ι -shadowing times being an artifact of different average error growth rates), as expected.

4.1.3 Lorenz equations

Consider first a ‘perfect model scenario’ where both the system and the model are the Lorenz equations with classical parameter values, integrated using fourth order Runge-Kutta with a time step of 0.01. Observational data is observed to three decimal places in disjoint segments of 250 or 500 observations. Perfect model ι -shadowing times are very long and all the sets of observations are shadowed for their duration. Typically, perfect model trajectories are found which ι -shadow system observations

for at least 3000 time steps of 0.01, or 30 time units, where one cycle is about 0.85 time units.

If the model is made imperfect by increasing the time step of the integration to 0.02 time units then the ι -shadowing times are shortened, but the model still ι -shadows for considerable times (over 50% of initial conditions are shadowed for at least 2.5 time units). Applying both the random search method and the combined random search and local linearised feedback methods of finding the optimal initial condition (see Appendix B for details) to the same set of trajectories, the combined method performs better, as illustrated in figure 4.6. The combined method is also applied to

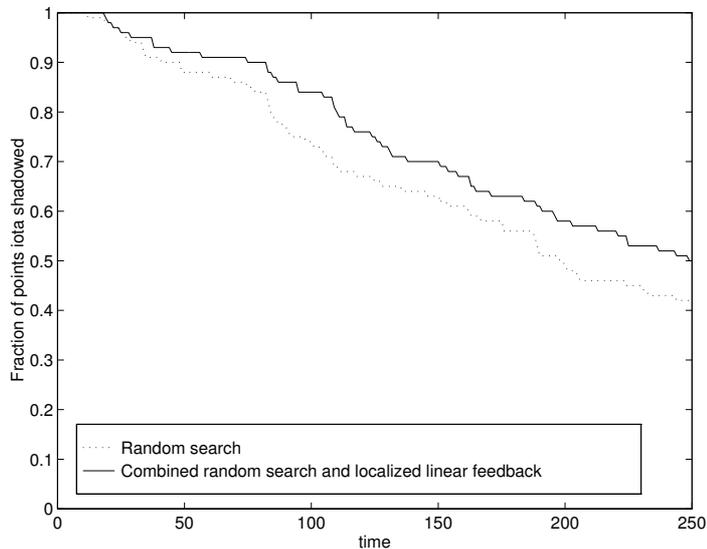


Figure 4.6: ι -shadowing results for the Lorenz equations comparing the random search and combined methods for finding the optimal initial condition; time is shown in 0.01 time steps. The model is the system equations integrated with a time step of 0.02.

the longer sets of 500 observations, see figure 4.7, enabling the distribution of the last points of the ι -shadowing model trajectories, or ‘failure’ points, to be determined. These failure points are shown in figure 4.9 and are compared, in section 4.2, with regions of rapid error growth of the system to see if ι -shadowing may be used to identify regions of model error. Firstly, we consider the effects of dynamical noise.

ι -shadowing in the presence of dynamic noise In order to investigate the effect of dynamic noise on the ability of a model to ι -shadow, a system trajectory was generated by iterating the Lorenz equations (using a time step of 0.01) and adding stochastic noise at regular intervals; the trajectory was observed to three decimal places to give a set of system observations. The stochastic noise was ‘trimmed’

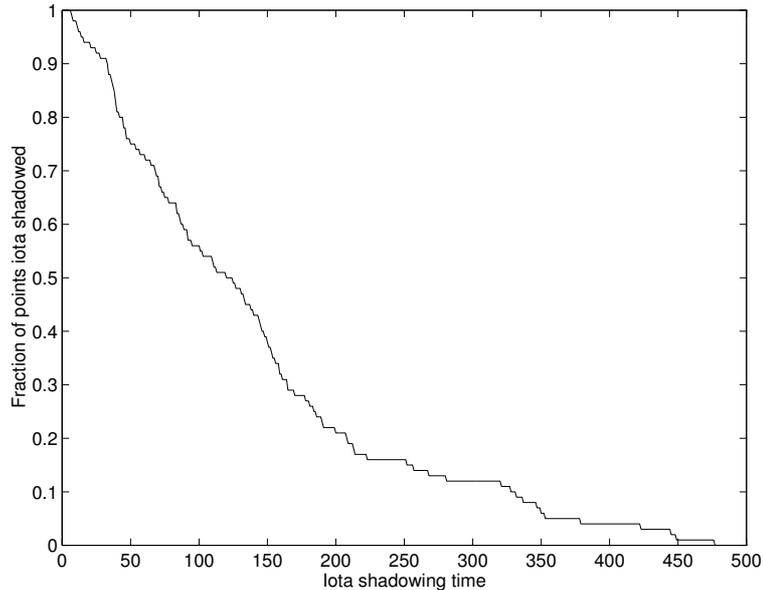


Figure 4.7: ι -shadowing results for the Lorenz equations using the combined method for finding the optimal initial condition, the model being the system equations integrated with a time step of 0.02; time is in 0.01 time steps.

Gaussian noise added to each component, constrained to lie within a certain interval; noise for each of n components chosen from a Normal distribution of mean μ and standard deviation σ and confined to lie within s standard deviations of the mean will be denoted $\mathcal{N}_n(\mu, \sigma|s)$.

Four different sets of observations were generated with stochastic dynamical noise; (1) $\mathcal{N}_3(0, 0.01|5)$ and (2) $\mathcal{N}_3(0, 0.05|2)$ noise added every 5 steps and (3) $\mathcal{N}_3(0, 0.05|5)$ and (4) $\mathcal{N}_3(0, 0.10|2.5)$ noise added every 10 steps. These values were chosen to see the effect of adding larger error at the same frequency by comparing (1) and (2), and in order to to compare the effect of adding noise at the same frequency and with the same bound, but with different standard deviations ((3) and (4)).

The Lorenz equations, integrated using the same 0.01 time step, but with no dynamical noise was used as a model and the resulting ι -shadowing time distributions are shown in figure 4.8.

Note that we are comparing the performance of the same model for various systems, rather than comparing the performance of different models. Comparing the distributions of ι -shadowing times in figures 4.8 and 4.7, we can see that the presence of dynamic noise in the generation of observations does not prevent an imperfect model, effectively of much lower dimension than the system, from shadowing for considerable times. As expected, lower levels of dynamic noise allow the imperfect model

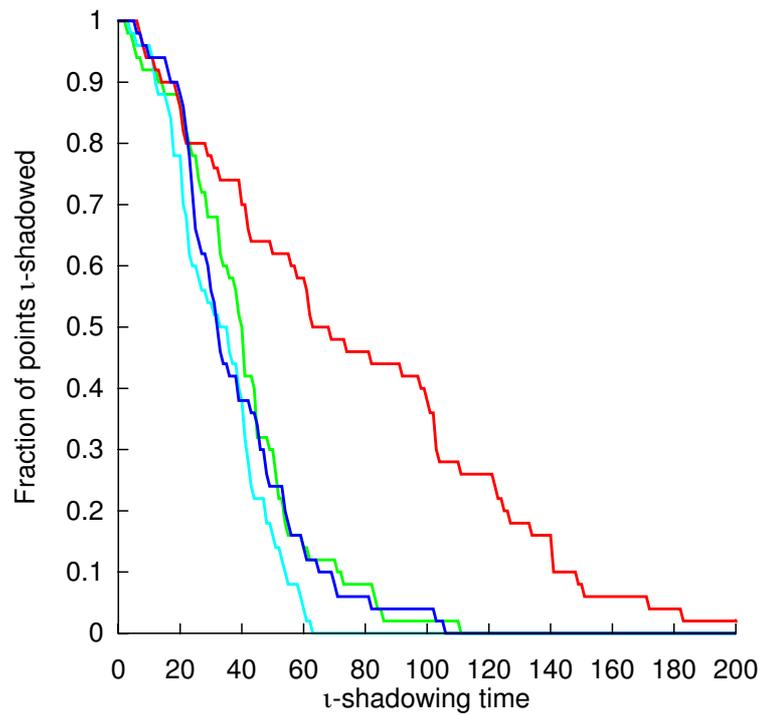


Figure 4.8: ι -shadowing results using the Lorenz equations to model observations generated with stochastic dynamic noise added at regular intervals. Various levels of stochastic noise were used in generating the ‘system’ observations: $\mathcal{N}_3(0, 0.01|5)$ noise added every 5 time steps (red), $\mathcal{N}_3(0, 0.05|2)$ noise added every 5 steps (blue), $\mathcal{N}_3(0, 0.05|5)$ noise added every 10 time steps (green), and $\mathcal{N}_3(0, 0.10|2.5)$ noise added every 10 steps (cyan).

to ι -shadow for longer. Similarly, a smaller standard deviation in error allows the imperfect model to shadow for longer than error with the same bound but a larger standard deviation.

4.2 ι -shadowing failure, model error and system sensitivity

In this section we demonstrate that ι -shadowing may be used to identify regions of large model error. The correlation between failure points and regions of rapid error growth is illustrated using the Lorenz system; a more thorough investigation using the Sinai map then follows.

Above it was mentioned that, for the Lorenz equations, the points at which an imperfect model fails to ι -shadow were recorded; they are illustrated in figure 4.9 with a trajectory of the Lorenz equations also shown for reference. Comparing the distribution of these failure points with the distribution of points according to the rapidity at which errors double, see figure 4.10 (from *Smith* [80], figure 1a), it appears that the regions visited just before ι -shadowability is lost correspond to those in which errors grow rapidly. Comparing the distribution of failure points with the invariant measure (figure not shown), no such correspondence is apparent.

The correlation between regions of ι -shadowing failure and rapid error growth suggests that ι -shadowing may be used to indicate regions of model error; the imperfect model of the Lorenz equations has a larger time step of integration than the perfect model, and hence we expect the greatest model error to be in regions where errors grow most rapidly. In the case of data based models, such as radial basis function models, model error often indicates sparseness of data in that region. If the regions of model error can be found, then a denser concentration of information may be able to be obtained in such regions, and the model thereby improved. This is now done explicitly for the more computationally tractable Sinai map.

The 2d Sinai map, introduced in section 2.3.2.3, may be translated to be centred about the origin and represented as a two step map

$$x'_n = x_n + y_n + 0.5 - \Delta \cos 2\pi y_n \quad (2.20)$$

$$y'_n = x_n + 2y_n + 1 \quad (2.21)$$

$$x_{n+1} = (x'_n + 0.5) \bmod 1 - 0.5 \quad (2.22)$$

$$y_{n+1} = (y'_n + 0.5) \bmod 1 - 0.5 \quad (2.23).$$

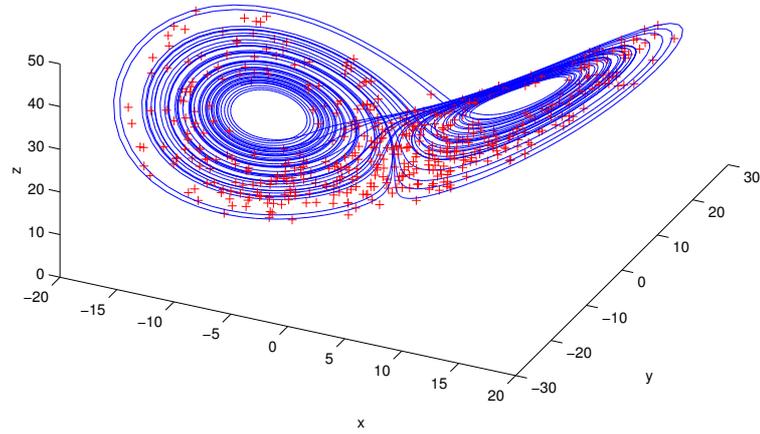


Figure 4.9: Distribution of the final points of the model to ν -shadow observations from the Lorenz equations (red plus signs). A trajectory of the Lorenz equations (blue line) is also shown for reference.

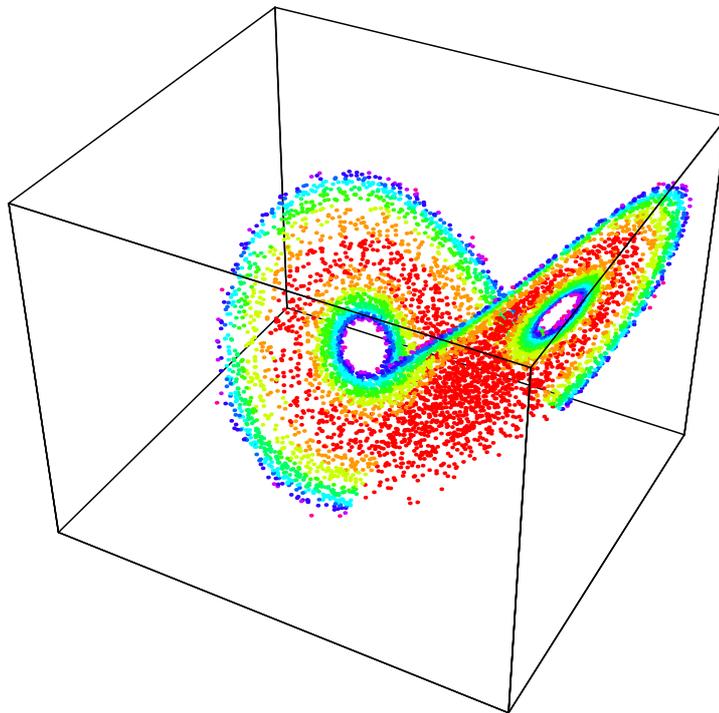


Figure 4.10: Distribution of the error doubling times on the Lorenz attractor, with red representing points with the most rapid, and lavender the slowest, error doubling. Figure 1a from *Smith* [80].

The map is iterated, and the values of x', y' observed to three decimal places. Expanding the cosine term of equation 2.20 in a (convergent) Taylor's series and retaining only the first four terms of the expansion yields an imperfect model whose error may be calculated exactly, as illustrated as a function of y' in figure 4.11 a. The model

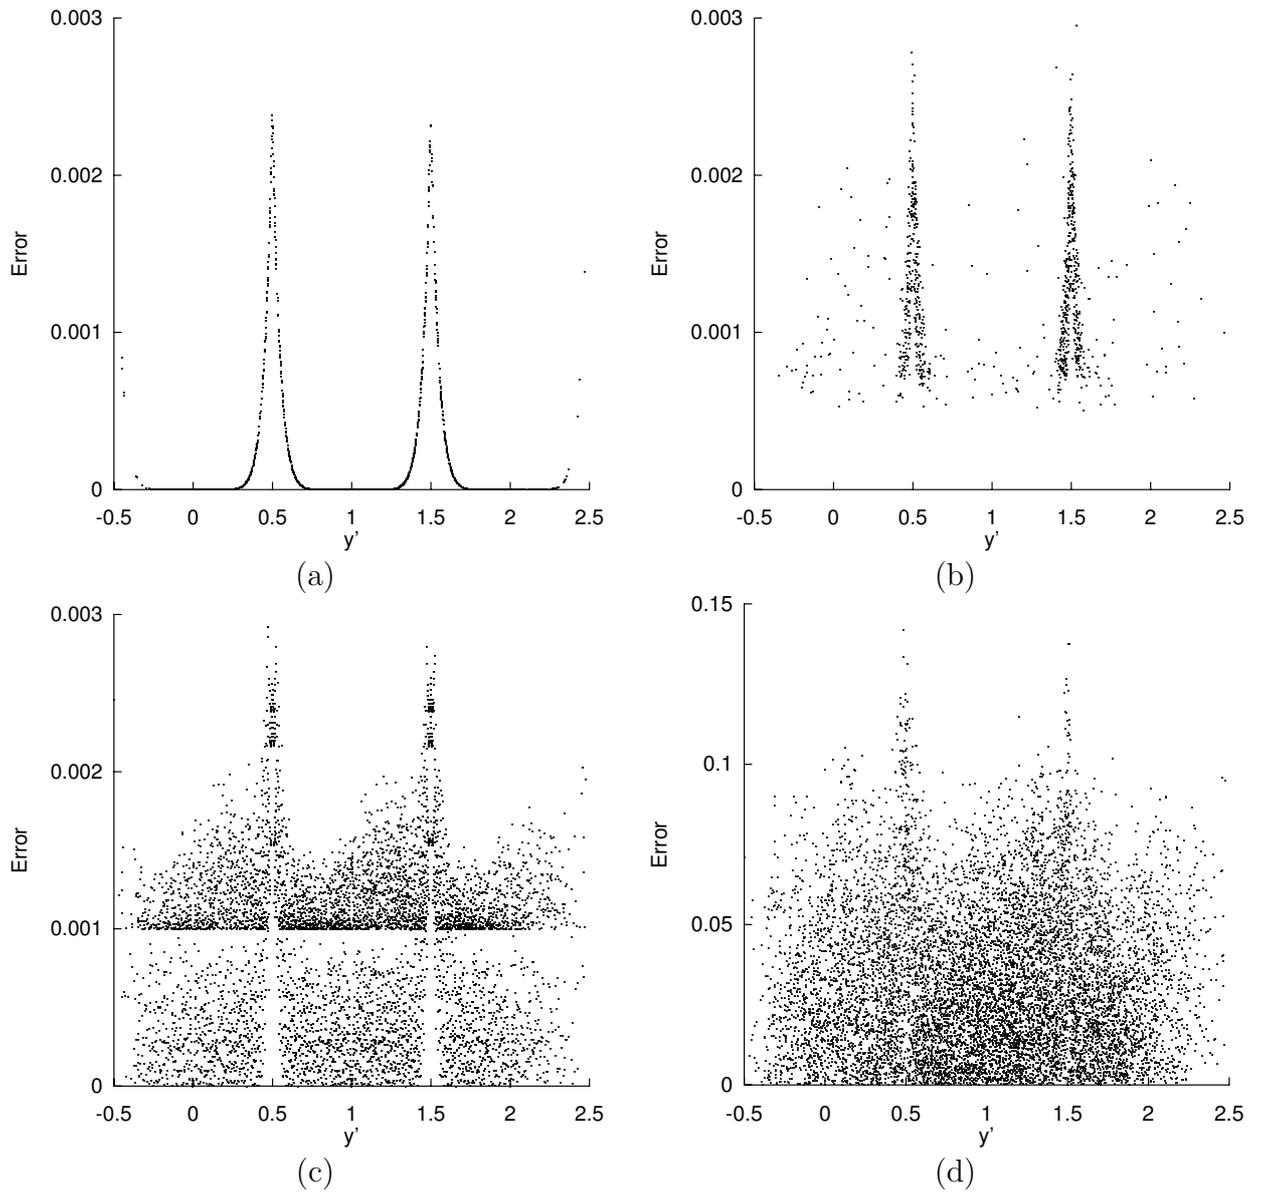


Figure 4.11: (a) Model error for a four term Taylor approximation to the translated Sinai map as a function of y' , (b) error (with respect to the system values) of points which fail to shadow observed data for 5 time steps using the imperfect model, (c) mean one step prediction error with respect to the observations and (d) mean five step prediction error with respect to the observations. Note that the scale of figure d differs from the others.

error is of course periodic in y' , due to the approximation error in the term $\cos 2\pi y$, with the error maximised when $2y$ is an odd integer.

Given only the model and observations, the ‘traditional’ method of estimating the model error is the one step observed prediction error. If x_n is the system value at time n , observed to be y_n and f is the system approximated by the model \tilde{f} , then the model error is $|f(x_n) - \tilde{f}(x_n)|$. The one step observed prediction error is given by $|y_{n+1} - \tilde{f}(y_n)|$, and is illustrated in figure 4.11 c. While there is some evidence of the one step observed prediction error reflecting the model error, it is not clear; note the obvious break at the quantisation level of 10^{-3} .

An alternative approach to localising model error is to recognise that the model is imperfect and ι -shadow the observations. Short disjoint segments of 5 observations are taken, and if no model trajectory can be found which ι -shadows over the 5 observations, the point at which shadowing fails is noted. The associated error for this point is the difference between its iteration under the model and the corresponding observation, which is illustrated as a function of y' in figure 4.11 b. The error of the failure points clearly peaks periodically, as with the model error. While, by definition, the error at a point which fails to ι -shadow is bounded below by the observational uncertainty when quantised, the upper range of error is similar to that for the model error. The distribution of shadowing failures is found to vary little with varying trajectory duration. The five step observed prediction error is also shown (figure 4.11 d) and the model error is found to be less clearly reflected than in the one step observed prediction error.

The figures above show how the model error, mean one step prediction error and error of failure points vary as a function of y' . In order to show the regions in which the different errors are distributed, all points for which any of the three errors exceeds a threshold, chosen as 0.001, are plotted in figure 4.12. (Note that from figure 4.11 it can be seen that the choice of threshold is relatively unimportant.) While the points for which one step observed prediction error are widely distributed over the Sinai attractor, those for model error or error of shadowing failure points are confined to bands for which $y' \simeq 0.5, 1.5$. ι -shadowing allows better localisation of model error.

For the imperfect model used here, it is easy to improve the model by including more terms of Taylor expansion; with just one more term there is a factor 10 decrease in the model error, and the maximum model error is less than the observational uncertainty. The simple approach for using shadowing to locate regions of model error presented above has its limitations: if all the model errors are below the observational uncertainty then there are few points which fail to shadow over short time spans. For physical systems, it is unlikely that the model error will be the same order of magnitude as the observational uncertainty; alternative approaches, using failure points

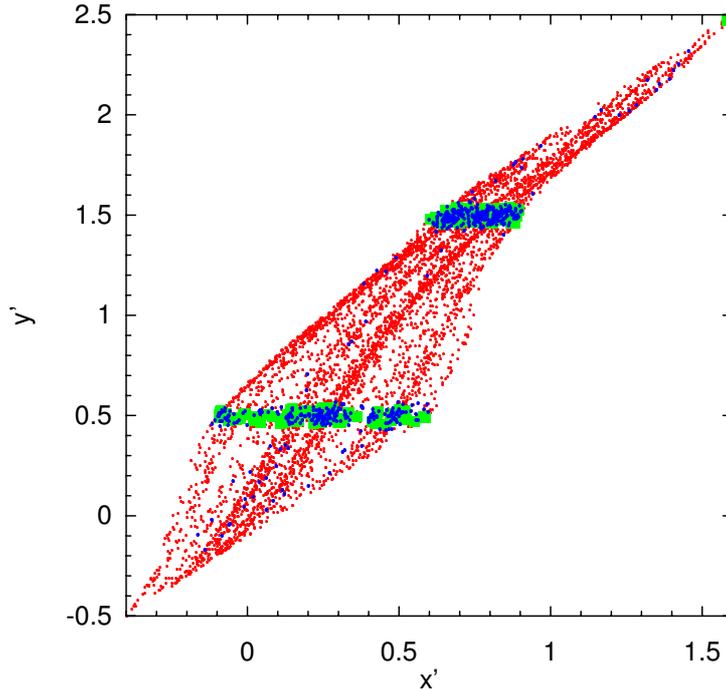


Figure 4.12: Points at which different errors exceed a threshold of 0.001 for the Sinai map: one step observed prediction error (red dots), model error (green filled squares) and error of shadowing failure points (blue dots).

of ι -shadowing trajectories over longer time scales, may be formed to localise model error in such cases. In contrast, the one step observed prediction shows a similar error pattern for the five term expansion; without knowing what the model error is, one could not decide when the one step observed prediction error contains information and when it does not.

Shadowing failure over short time spans has been shown to indicate regions of large model error more effectively than one step observed prediction error. In the following section we will use shadowing failure points to identify regions of model error for an imperfect radial basis function model of the rotating fluid annulus. Using this information the model will be modified and the performance of the two models compared by contrasting shadowing time distributions.

4.3 Improving a RBF model of the annulus

When contrasting distinct models with the aim of identifying (or constructing) a perfect model, we require a model trajectory to be consistent with the observations, given the observational uncertainty, if it is to ι -shadow. In the case of quantisational uncertainty, this implies that each point of the model trajectory *must* lie within a

distance, defined by the observational uncertainty, of the observations; if any point is outside this tolerance, the model trajectory fails to ι -shadow. In contrast, when tuning a particular predictive model the specific tolerance to be set is less obvious: given that we know the model is imperfect, it may be advantageous to choose a tuning tolerance which is greater than the observational tolerance. Ideally, in the limit of infinitely large data sets and potentially perfect model structures, the tuning tolerance will approach the observational tolerance. In the finite data case, a larger value for the tuning tolerance can more clearly identify regions of model error; a tuning tolerance larger than the observational tolerance is used below.

Radial basis function (RBF) models are global interpolation schemes which may be constructed using observational data; details of model formulation are given in section 2.2.1.2 and applied to observations of the Moran-Ricker map in section 2.2.2.2. The complexity of the model is determined by the number of centres used to define the model, and indirectly by the dimension of the data used. We wish to use ι -shadowing failure points to locate regions of model error in which more centres may be chosen with the aim of improving the RBF model.

The system we consider is the thermally driven rotating fluid annulus in a chaotic regime; details of this laboratory experiment are in section 2.3.3. Observations of temperature are taken for a highly chaotic regime with a sampling time τ_s of 2s. Given that the typical period of oscillation seen in the observations is between 320s and 512s, the data is re-sampled at a lower frequency (1 in 8). The data is then embedded (for details see section 2.3.1) in a 5d space with a delay time $\tau_d = 3(8\tau_s)$. For purposes which will become clear below, the data was divided into a 3 part learning set, each part consisting of 400, 300 and 100 points respectively, which is used to build and improve the model, and a test set consisting of about 1000 points, used to test the models out-of-sample.

A basic model is constructed by choosing 64 centres from the 400 (embedded) data points in part I of the learning set and using a Gaussian radial basis function. This model is then tested on part I of the learning set (*i.e.* testing is intentionally in-sample): every fifth data point is taken as an initial condition and a trajectory which ι -shadows for five time steps sought; the (49) points at which shadowing fails after less than five time steps are recorded, along with their associated error, as above. Quantisational uncertainty representing the tuning tolerance, discussed above, is taken to be 0.1°C in each component. The resulting shadowing failure points are hoped to suggest regions in which the model error is large. Due to the structure of RBF models, large model error may be due to sparseness of information in these

regions, and hence the model may be improved by adding centres (and their associated images) in the regions of shadowing failure. From part II of the learning set, the point closest to each of the 49 failure points is found, and these 49 points are then added as centres, and a modified, and hopefully improved, model with 113 centres built. A further model, with 113 centres chosen (in the same way as for the 64 centre model) from parts I and II of the learning set was also constructed.

The three models (the original with 64 centres, the 113 (standardly chosen) centre model and that modified using failure points) were then tested out of sample by establishing their shadowing time distributions, shown in figure 4.13. Both of the

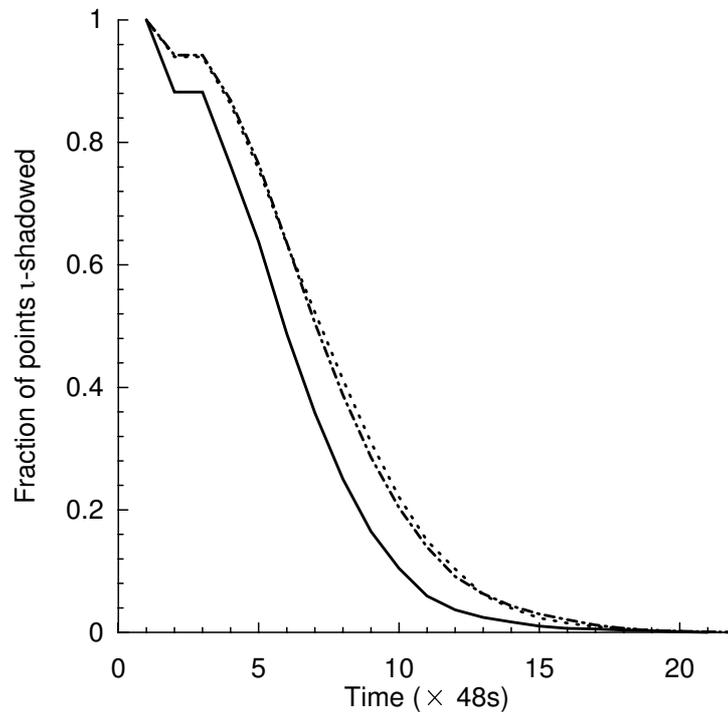


Figure 4.13: Comparison of ι -shadowing time distributions for the original 64 centre RBF model of the annulus (solid line), a modified model with 113 centres chosen from parts I and II of the learning set (dot-dashed line), and the model with 113 centres chosen using the location of shadowing failure points (dotted line).

modified models are an improvement as the ι -shadowing time distributions are greater than that of the original model. The model modified using the location of shadowing failure points is the best of those considered; points of ι -shadowing failure seem to have indicated regions of large model error and thus allowed the model to be improved. Note that shadowing times of 21 time steps ($21 \times 48s$) are achieved: a typical prediction time for the annulus is 18 time steps [78].

The process may be repeated, testing the model modified using failure points on parts I and II of the learning set and finding additional centres close to shadowing

failure points from part III of the learning set. The improvement possible is limited, as the shadowing failures seem to occur in similar regions, so the additional centres are close to or the same as those already included. This problem may be alleviated by the availability of more data, but alternatively, it may indicate a region of state space in which the embedding fails. By considering only those points which fail to shadow and have relatively large error (in comparison to the observational uncertainty), this problem is alleviated somewhat and models can be constructed for which almost all points in the test set shadow for at least 5 time steps; unfortunately most points fail to shadow soon after, so that the model is not an unambiguous improvement. The time for which the shadowing trajectories are sought may also be altered; if a shadowing time of 20 is desired, adding centres close to the failure points results in a modified model which shadows a smaller fraction of points for a short time, but which shadows a larger fraction of points for times above about 15 time steps.

Using regions of ι -shadowing failure to locate regions of model error has been shown to result in improved RBF models for the annulus. Improvements are limited, possibly by the quantity of data available; the type of improvement achieved may be varied according to the purpose of the model.

4.4 From evaluation to prediction

4.4.1 Limitations of ι -shadowing

Used as a tool for comparison ι -shadowing reflects differing abilities of various models. The illustrations above give desired results; perfect models outperform all other models, and the relative performance of other models may be logically deduced for simple systems, *e.g.* the logistic and Moran-Ricker maps. Results from the Lorenz system illustrate the superiority of the combined random search and local linearised feedback method of finding the optimal initial condition in implementation of the ι -shadowing algorithm, and suggest that ι -shadowing may be used to determine areas of phase space in which a model fails to represent the dynamics of the system.

Application of ι -shadowing to a RBF model for fluid flow in a chaotic regime in the annulus demonstrates that ι -shadowing may be implemented to compare models when the governing equations are truly unknown and the method reliable on observational data and a given model.

It should be noted that while ι -shadowing is valid in its use of model evaluation as discussed above, constructing a model which satisfies the ι -shadowing criterion ‘in

sample' is trivial. Given a series of observation σ_i , $i = 0, \dots, N$ correct to Q bits then iterates of the map

$$F(x) = (2^Q x) \bmod 1$$

observed to Q bits will ι -shadow the data, if x_0 is given by

$$x = \sigma_0 + \sigma_1 2^{-Q} + \sigma_2 2^{-2Q} + \dots + \sigma_N 2^{-NQ}. \quad (4.1)$$

For an infinite time series it would be necessary to find an approximation to $\sum_{n=N+1}^{\infty} \sigma_n 2^{-nQ}$.

ι -shadowing does *not* validate a model as a predictor, since it requires knowledge of the observed trajectory; however the ι -shadowing time will give a limit of predictability² for the model, allowing comparison between models to be made from observational data available. ι -shadowing may also be considered as an assimilation technique, which can be applied as a method of noise reduction, or for gap filling of incomplete data sets, as described below.

4.4.2 ι -shadowing as an assimilation technique

Most assimilation techniques look to assimilate the observational data available, often using the dynamics of the model, and provide analysis values which are hopefully closer to the system values than the observations. Methods differ in how the analysis values are chosen through their choice of measure to be optimised; varying the initial condition to obtain varying model trajectories in search of an optimal solution is common.

One approach is to formulate the task into a constrained least squares minimisation problem through the use of an objective function, typically defined to be the sum of squares of differences between observations and the values given by the model trajectory, with the initial condition taken to be the control variable [92]. Assuming the statistics of the observational uncertainty are known, gradient descent methods solve this problem through the use of an adjoint operator, requiring repetitive application of a method in order to converge on the minimising solution due to limited regions in which linearity holds. If the assimilation is only over past observations to give an optimal present analysis value, then such a method is sequential assimilation. Variational assimilation is applied at times in the past using observations which occur both prior to and after the time at which the assimilation is implemented.

As discussed in section 3.3, model trajectories which are optimal in the root mean square sense need not be consistent with the available data. An alternative approach

²Given stationarity of the data.

is given in the form of ι -shadowing, which will gradually adjust the initial condition to give a model trajectory which ι -shadows for longer times. If we seek a trajectory which ι -shadows over an interval which includes the time at which an analysis is required, then the resulting analysis given by the ι -shadowing trajectory will be consistent with the available information in the interval considered. If there is no trajectory which ι -shadows for the given interval, then the trajectory with the longest ι -shadowing time is used. This method therefore not only provides analysis values, but also some measure of the ability of the model to represent the system; it does not always return an analysis value for a fixed assimilation window, but those returned are consistent. (In sections 5.2.2 and 5.2.3 ι -shadowing is used to provide analysis values.)

In addition to providing (hopefully) more relevant analysis values, any assimilation technique may be used as a method of noise reduction or ‘gap filling’. For delay embedded data, normal nonlinear noise reduction involves ‘predicting’ a point from the preceding and following values and taking this ‘prediction’ as the current value, *e.g.* use $(x_{-2}, x_{-1}, x_1, x_2)$ to predict x_0 . Given observations the optimal ι -shadowing trajectory may be determined. Each point in the trajectory, *e.g.* $\mathbf{x}_i = (x_{i-n+1}, x_{i-n+2}, \dots, x_i)$ (assuming that we are using delay co-ordinates in \mathcal{R}^n), then contains estimates of n consecutive system values; the shadowing trajectory provides (upto) n different estimates of each observation. This is superior since any of the suggested improvements to the system value, given by the n estimates, will be consistent with a longer stretch of the dynamics and the distribution of estimates of the system value give a consistency check. Note that noise reduction does not require delay embedded data; the same approach may still be used, with each ι -shadowing trajectory providing a single estimate of each observation.

Such a method may be adapted to provide consistent values to fill gaps in data. Linearly interpolating the data available to give a initial estimate, this may be substituted as an ‘observed’ value, with the corresponding observational uncertainty set to, say, twice the maximum deviation from the mean value of any observed data point. Application of ι -shadowing will yield an ι -shadowing trajectory which will provide a distribution of n estimates if the data is delay embedded (one if it is not) for the unknown data point.

4.5 Finite improvements: the need for ensembles

However much data is available, it will contain observational uncertainty and models of physical systems will be imperfect. Assimilation techniques can provide improved

estimates of the best initial condition, given the model dynamics, but the analysis values will still be inexact. Due to the believed chaotic nature of the observed physical systems which we wish to predict, uncertainty in the initial condition is important. In the following chapters we will consider the use of a group, or ensemble, of initial conditions, each of which is consistent with information available at the time of prediction, to predict the probability of future system values.

4.6 Summary

In Chapter 4 ι -shadowing has been applied as such to compare competing models for a variety of systems. Information about regions of model error was shown to be contained in the distribution of ι -shadowing failure points, and this finding was exploited to build an improved model of the annulus. Applications for noise reduction were presented and the limitations of ι -shadowing discussed.

Polynomial interpolation models of the logistic map were compared, enabling the ability of ι -shadowing to distinguish perfect models from imperfect ones to be evaluated. Comparing the ι -shadowing time distributions for (1) models with varying numbers of base points, (2) both direct and iterative methods, (3) models with base points which are either equally spaced or spaced according as the natural measure, and (4) linear and quadratic interpolation models, all gave results as expected; the perfect quadratic iterative map was by far the best model according to the ι -shadowing time distributions.

Similar comparisons of polynomial models were executed for the Moran-Ricker map since, in contrast to the logistic map, no finite term polynomial model is perfect. Other than for the quadratic iterative models, the results were qualitatively similar to those for the logistic map, with the exception of the comparison of direct and iterative models. Differences in the complexity between the one step map and multi-step iterations between the logistic map and the Moran-Ricker map explained the disparate results.

A perfect model of the Lorenz equations was shown to ι -shadow for very long times. Increasing the integration time step provided an imperfect model which was used to compare the algorithms for finding the optimal initial condition; the combined search method was found to be optimal. Varying levels of dynamic noise were used to form high order systems, observations of which were ι -shadowed using the same deterministic imperfect model; ι -shadowing time distributions varied depending on the magnitude and frequency of the dynamic noise. The imperfect model was found

to perform better when noise of a smaller magnitude, or noise with the same bound but distributed with a smaller standard deviation, was added. The ι -shadowing times were still considerable, demonstrating that ι -shadowing may be used to evaluate the ability of models to exhibit the behaviour of a system with considerably different structure.

Regions of ι -shadowing failure points for the imperfect model of the Lorenz equations were found to correlate with regions of maximum system sensitivity: since regions of systems sensitivity are expected to indicate regions of greatest model error this result motivated the use of ι -shadowing to localise regions of large model error. A systematic investigation using imperfect models of the Sinai map showed ι -shadowing to localise model error better than one step prediction error (calculated with respect to the observations). The application of ι -shadowing to localise model error was applied to a RBF model of the annulus, and the model modified with some success.

It was noted that ι -shadowing does not validate a model as a predictor, but the ι -shadowing time will give a limit of predictability for a model. The use of ι -shadowing as an assimilation technique was then discussed. Noting that while, for a given assimilation window, ι -shadowing would not always return an analysis, those returned would be consistent with the observations over the interval for which the trajectory shadows. It was outlined how ι -shadowing, used as an assimilation technique, may be employed in noise reduction; the use of ensembles was then promoted since the improvements in the accuracy of the initial condition possible using assimilation techniques are finite. The choice of ensembles, and testing the validity of assumptions made in their definition, are the subjects of the following chapters.

Chapter 5

Probabilistic predictions: ensemble construction and performance evaluation

The weather is one of many physical systems whose prediction poses a complex challenge. Uncertainty in observation of physical systems and the limited accuracy of analysis values implies that predicting future system states by evolving the analysis is likely to be inaccurate even if the model is perfect. Coupled with chaotic models, the initial discrepancy between system value and analysis may grow rapidly in evolution; prediction errors due to uncertainty in the initial condition may be large and cannot be reliably estimated using linear growth rates. Probabilistic predictions aim to quantify the potential error in a prediction due to the analysis error or observational uncertainty by evolving *ensembles of initial conditions*. In addition to ensembles over uncertainties in the initial condition considered here, ensembles over other uncertainties, such as ensembles over different models, could be considered in a similar way [18].

In this chapter the various methods of ensemble construction are first introduced. The Marzec-Spiegel system, whose system equations are known, is used to contrast perfect and unconstrained ensembles, while a RBF model of the thermally driven rotating annulus enables the effects of model error on ensemble formation to be investigated. Constrained vector subspaces, used in the definition of ensembles in Numerical Weather Prediction are then introduced. The theoretical justification that a constrained vector ensemble will achieve a given aim is shown to depend upon a number of assumptions. It is judicious to ensure, as far as is possible, that these assumptions are fulfilled in implementation (*i.e.* that the ensemble is consistent with its definition), and to evaluate the performance of the ensemble according to the aim

for which it was designed; the second half of the chapter addresses these issues and applies the resulting measures to evaluate ensembles. Limited computer resources restrict the tests possible on high-dimensional physical systems, while the lack of a perfect model for *any* physical system inhibits a thorough comparison of ensemble formations and performances. Ensembles are therefore constructed and tested using a perfect model of the Marzec-Spiegel system and a RBF model of the annulus. Tests of internal consistency of operational ensembles used with Numerical Weather Prediction (NWP) models, where ensemble size is severely limited by computational resources, are delayed to the next chapter as is a discussion, facilitated by the range of models investigated, of the application of results from low-dimensional systems to high-dimensional ones.

5.1 Ensemble construction

Choosing the members of an ensemble requires care if the resulting predictions are to be informative. Ideally we wish to use a *perfect ensemble* of initial conditions, indistinguishable from an observed state of the system, whose distribution of evolved members at time t defines a probability density function (p.d.f.) to an accuracy limited only by sampling uncertainties arising from the finite nature of the ensemble. A perfect model and an exact understanding of the observational uncertainty are necessary, but not sufficient, requirements to be able to form such an ensemble. *Unconstrained ensembles* of randomly chosen members distributed according to the observational uncertainty distribution may include physically unrealisable states¹, *i.e.* members which do not lie on the invariant manifold of the system, and will not reflect the frequency of occurrence of system states, as described by the natural measure. The interpretation of p.d.f.s defined by unconstrained ensembles is therefore non-trivial; the errors in the p.d.f.s are not accountable to the finite nature of the ensemble and may in fact increase with increasing ensemble size. Constraining the members to the invariant manifold by distributing them according to the convolution of the natural measure of the system with the observational uncertainty distribution forms a perfect *constrained ensemble*. The errors which occur in such a perfect constrained ensemble are *accountable* since they are finite sampling uncertainties; given the size of the ensemble we may calculate the error which will occur in the predicted p.d.f.s. Accountable p.d.f.s

¹Unless the invariant manifold is equivalent to the state space.

are therefore straightforward to interpret and are useful until they are indistinguishable from the invariant measure (projected onto the observed variable) to which they evolve.

Discrepancies between system and model invariant manifolds, usual in imperfect models, render perfect ensembles inaccessible. Constraining ensembles to lie on the model invariant manifold in order to approximate perfect ensembles is impractical for systems with long Poincaré return times, while unconstrained ensembles which adequately sample the region of observational uncertainty rapidly become computationally infeasible as the dimension of the model increases. If reasonable approximations of true p.d.f.s² are unattainable, then the aim of ensemble forecasting must be reconsidered, and alternative constrained subspaces defined to facilitate optimal employment of a limited number of ensemble members. Two motivations used in defining alternative constrained vector (CV) subspaces arise from numerical weather prediction, namely to include the initial condition closest to the system value and to predict the ‘worse case scenario’ by capturing the spread; both require a given norm.

5.1.1 Unconstrained ensembles

Given only an observation and knowledge of the observational uncertainty distribution, an *unconstrained ensemble* may be formed by randomly choosing members distributed as defined by the observational uncertainty. Consider the 3d $((x, y, \lambda))$ Marzec-Spiegel system introduced in section 2.3.2.5, for which the system equations are known. Given a quantised observation, an unconstrained ensemble may be formed by randomly distributing members within the 3d hyper-cube defined by the quantisation uncertainty. Predicted p.d.f.s are obtained by evolving the ensemble under the 3d model, here taken to be a perfect model of the system. Figure 5.1 provides an illustration of the evolution over $t \in [0, 1]$ of a 64 member unconstrained ensemble. Initially the ensemble is distributed according to the observational uncertainty distribution, here a uniform distribution over a 0.01 interval centred about the observation in each of the three dimensions, and denoted as $\mathcal{U}_3(0.01)$. The ensemble is evolved under a perfect model of the Marzec Spiegel system. The illustration is obtained by then projecting the ensemble onto the (x, y) plane. The initially compact ensemble rapidly diverges as it evolves; by $t = 1$ the evolved members are distributed in $x \in (-0.1, 1.1)$ suggesting that no confident prediction can be made. Member concentration is high in $x \in (-0.1, 0.1)$, but sparse near the validating observation at

²Theoretically defined by evolving the observational uncertainty distribution exactly under the system.

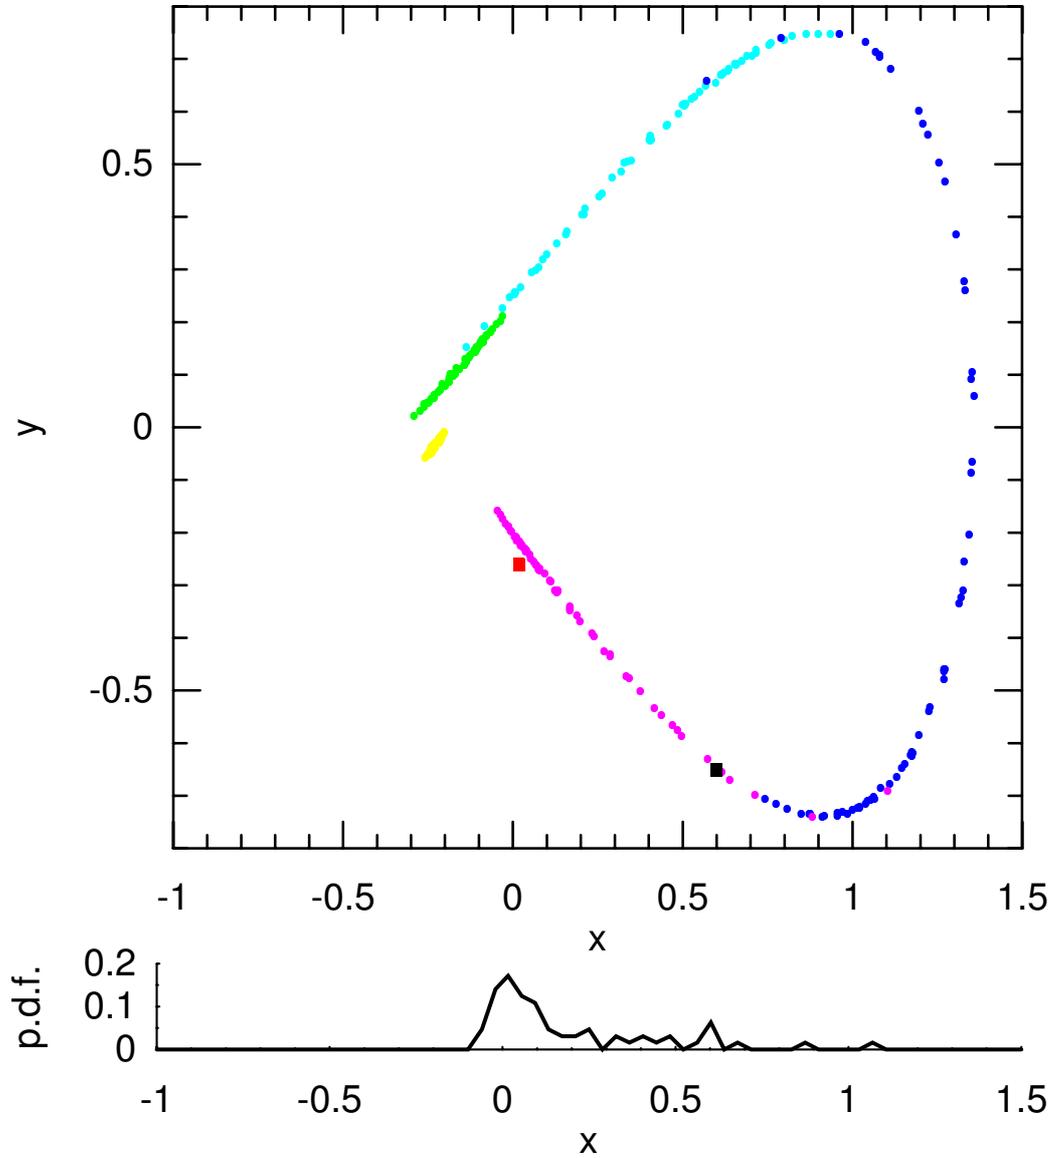


Figure 5.1: Evolution of a 64 member unconstrained ensemble, with $\mathcal{U}_3(0.01)$ distribution, using a perfect model of the Marzec Spiegel system, projected onto the (x, y) plane. Upper panel: An unconstrained ensemble centred on the observation at $t = 0$ is constructed within a hyper-cube (whose projection is red square) and then evolved under the model to give predicted members at $t = 0.2$ (yellow dots), $t = 0.4$ (green), $t = 0.6$ (light blue), $t = 0.8$ (dark blue) and $t = 1$ (magenta). The region containing the unconstrained ensemble constructed about the observation at $t = 1.0$ is also shown (black square). Lower panel: The probability density function given by the unconstrained ensemble initiated at $t = 0$ and evolved to $t = 1$; the p.d.f. reflects the density of ensemble members (magenta dots in the upper panel) as a function of x .

$t = 1$ (*i.e.* near $x = 0.6$), as shown by the p.d.f. in the lower panel of figure 5.1. This

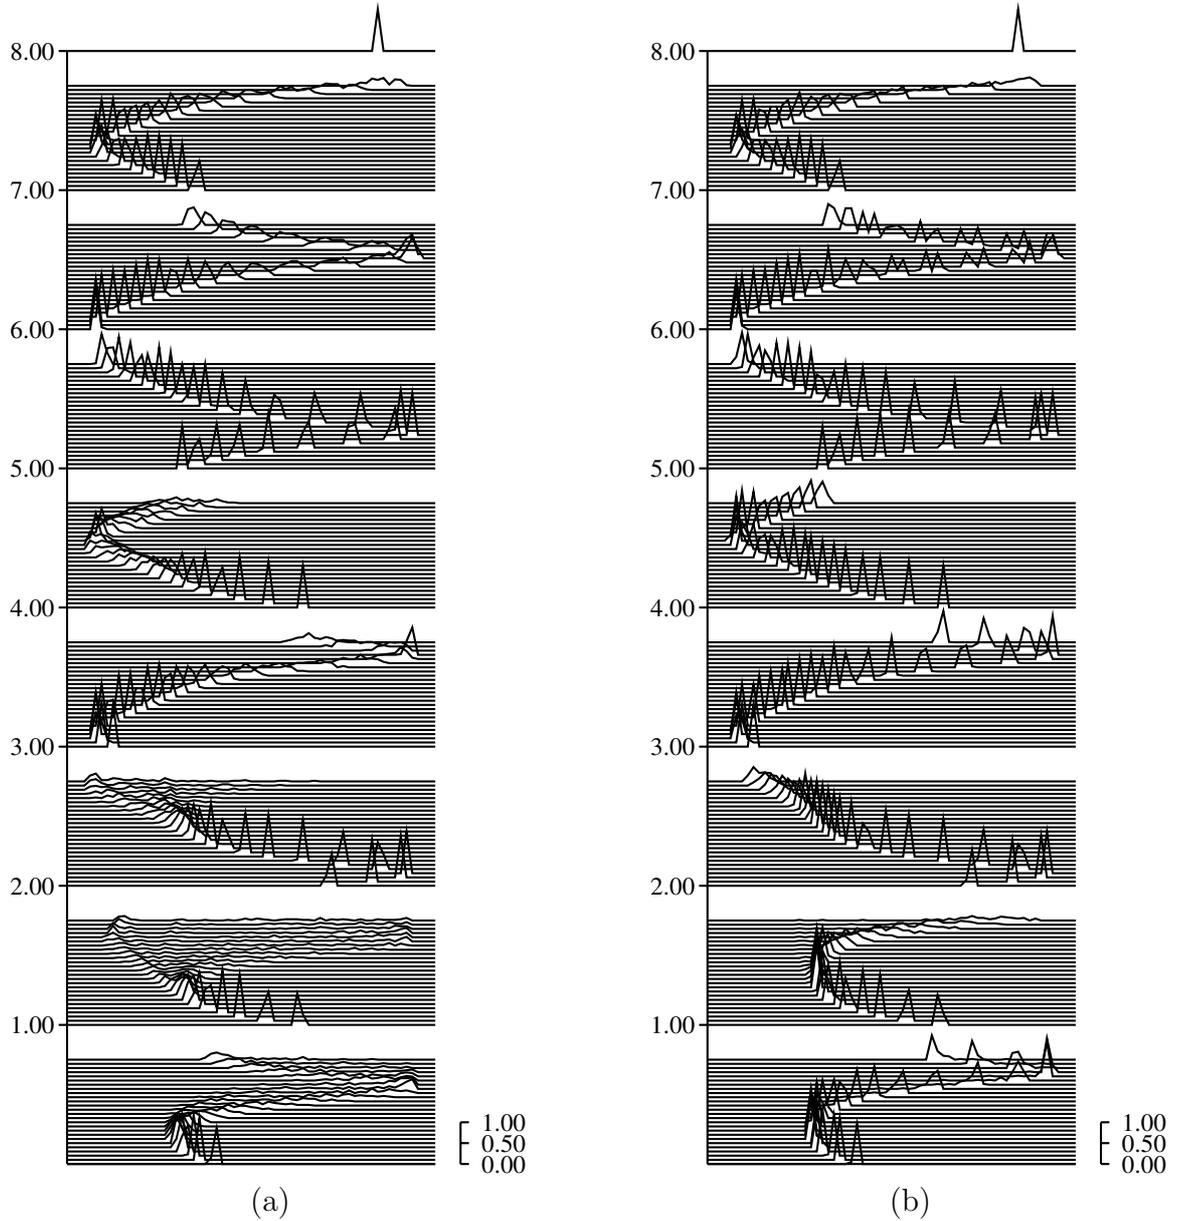


Figure 5.2: Evolution of p.d.f.s constructed from 64 member (a) unconstrained and (b) perfect ensembles (with $\mathcal{U}_3(0.01)$ distribution) evolved under a perfect model of the Marzec Spiegel system, projected onto $x \in [-1, 1.5]$. The gaps in the time axes indicate where new ensembles have been formed about the corresponding observation.

alternative p.d.f. representation is used in figure 5.2 a at time intervals of 0.04, and extended for eight consecutive unconstrained ensembles: the p.d.f.s calculated from the ensembles are displayed clearly by plotting histograms of members with respect to x as a function of time. The gaps in the time axis denote the formation of a new

ensemble about the observation at that time, thus enabling a cursory evaluation of ensemble performance. The divergence of the ensemble initiated at $t = 0$ and noted above is visible, the non-uniformity of the evolved p.d.f. is much clearer and we may observe that the p.d.f. predicted at $t = 1$ from the ensemble initiated at $t = 0$, hereafter denoted $t = 1 \downarrow_0$, has little power at the value of the validating observation, reflecting the sparsity of members at the validating observation in figure 5.1. The differing forms of the evolved p.d.f.s, *e.g.* the compact nature of that at $t = 6 \downarrow_5$ compared to the multi-modal prediction at $t = 3 \downarrow_2$, highlights the varying sensitivity of the model as a function of initial condition. Note also the bifurcation of the ensemble at $t \simeq 1.4$ and return of skill at $t \simeq 3.8, 4.6, 6.7$.

The errors in the unconstrained ensembles are known, a priori, not to be accountable, which complicates the interpretation of the predicted p.d.f.s. The p.d.f.s suggest that there is a higher probability of correctly predicting the observation at $t = 8$ than at $t = 1$ since the p.d.f. at $t = 8 \downarrow_7$ is more concentrated about the validating observation as compared to the diverse p.d.f. at $t = 1 \downarrow_0$, but the error in each p.d.f. is unknown. The model is perfect, thereby banishing the complications of model error, but consider the invariant manifold of this Marzec-Spiegel system (figure 2.16 in section 2.3.2.5 shows a cross section of the invariant manifold for $x = 0$ and $\dot{x} > 0$). Some of the randomly chosen members will not lie on the invariant manifold; they are physically unrealisable states which will never occur in the system. Furthermore, the system may visit different regions of the manifold with varying frequency, as expressed by the natural measure, and we require the members to reflect this distribution if the error in prediction is to be accountable to sampling uncertainties. These issues beget questions such as whether the bifurcation which appears at $t \simeq 1.4$ is a real characteristic of the observed system or the consequence of including members in the unconstrained ensemble which do not initially lie on the manifold (the answer to which is provided using perfect ensembles in section 5.1.2). Without some knowledge of the natural measure, it is impossible to address these questions and the errors in unconstrained ensemble p.d.f. predictions remain variable and unknown.

Model imperfections further complicate unconstrained ensemble p.d.f. interpretation not only by the introduction of model error, but also by raising the question of whether the projected system invariant manifold is contained by that of the model. Figure 5.3 shows 256 member unconstrained ensembles, with Gaussian ($\mathcal{N}(0, 0.05)$) distribution in each component (*i.e.* $\mathcal{N}_5(0, 0.05)$), evolved using an imperfect 5d (short time scale) RBF model of the annulus (see sections 2.2.1.2, 2.3.3 and 4.3 for details). The varying certainty of prediction as a function of initial condition is again seen here,

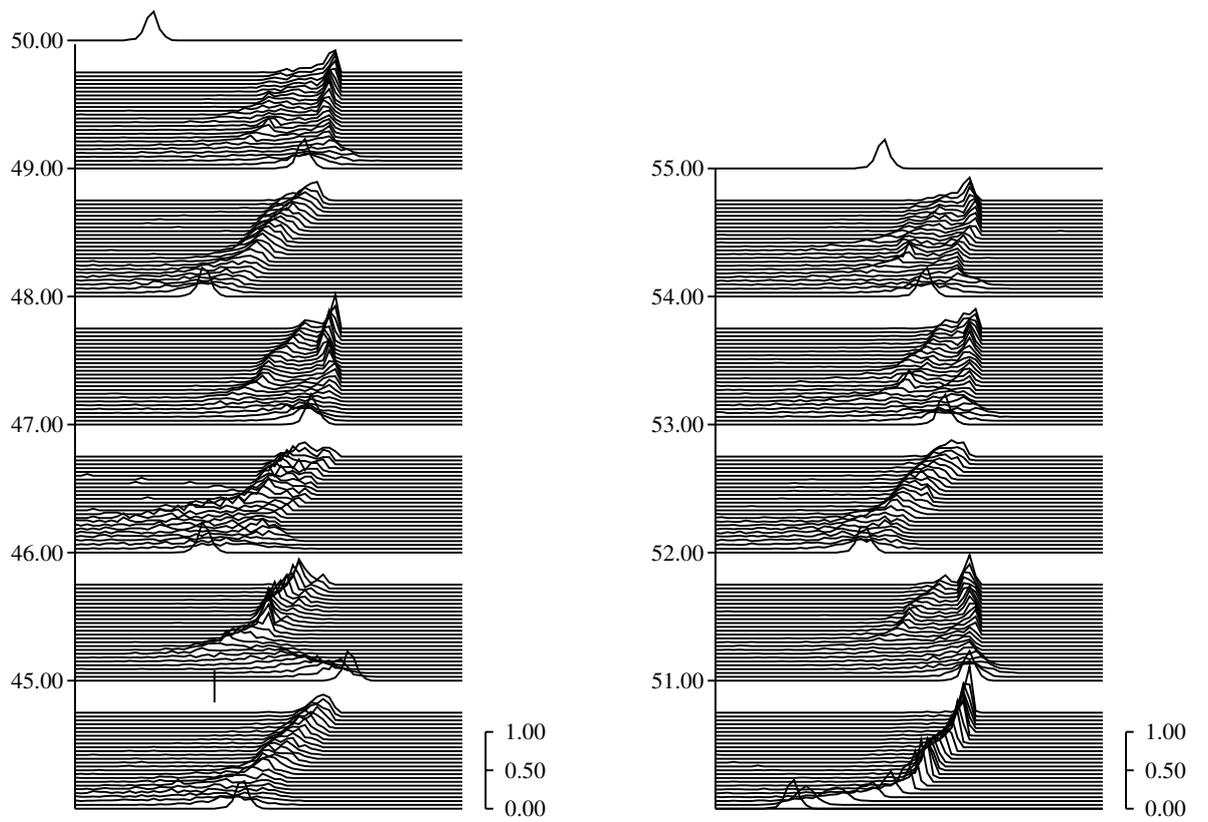


Figure 5.3: P.d.f.s from unconstrained 256 member ensembles with Gaussian ($\mathcal{N}_5(0, 0.05)$) distribution evolved under an (imperfect) 5d RBF model of the annulus, projected onto $x_5 \in [22^\circ\text{C}, 25^\circ\text{C}]$.

but now compact evolved p.d.f.s suggesting confident predictions may be accurate, as at time $t = 51 \downarrow_{50}$, or inaccurate due to model error, as at time $t = 48 \downarrow_{47}$. Many members evolve out of the observed temperature interval of the system $T \in [22^\circ\text{C}, 25^\circ\text{C}]$, *e.g.* in the time interval $t \in [46, 47]$.

There are two manifestations of an imperfect model which we wish to be reflected in the ensemble evolution: model error which causes the ensemble to evolve differently from the system, and the incongruence of the system manifold projection with the model manifold. In contrast, inclusion of initial members which are physically unrealisable representations, a situation exacerbated by manifold incongruence, causes undesirable, unquantifiable errors in the resulting ensembles which obfuscate p.d.f. interpretation as noted above. Unfortunately diagnosing the extent of the different influences is difficult, especially given only observations.

Accepting the limitations of unconstrained ensemble predictions, a large unconstrained ensemble may be used to densely sample the observational uncertainty, ensuring some members lie on the manifold and yielding p.d.f.s which provide some useful information, *e.g.* providing an upper bound on spread. As the complexity of the model increases however, the number of ensemble members which can be evolved in a reasonable time rapidly decreases while the dimensionality of the model space is increasing, *e.g.* a maximum of ~ 50 members may be evolved operationally for NWP models with approximately one million dimensions [59]. The combined effect of model error and finite sampling render unconstrained ensembles useless as the size of the ensemble is restricted while the dimensionality of space to be sampled increases. Consider figure 5.4 (from *Smith & Gilmour* [84]) which shows the evolution of four 128 member unconstrained ensembles (initiated at $t = 11, 17, 23, 29$), with Gaussian ($\mathcal{N}_5(0, 0.075)$) distribution, each evolved using an imperfect (long time scale) RBF model of the annulus. The ensembles are reformed after six time steps since the six step prediction is diverse in each case, despite being initiated within a small region, and if continued many of the ensemble members would evolve out of the vicinity of the validating observation. The ι -shadowing trajectory which shadows out to $t = 26$ shows that the sparseness of members renders the ensemble useless after a relatively short time. (In fact, we may wish to include the initial condition which ι -shadows for a long interval as a member of the ensemble; we return to discuss this issue in section 5.1.3.3.) An unconstrained ensemble may provide more information for longer if it sampled the initial uncertainty more densely, but even for this 5d model larger ensembles are computationally infeasible.

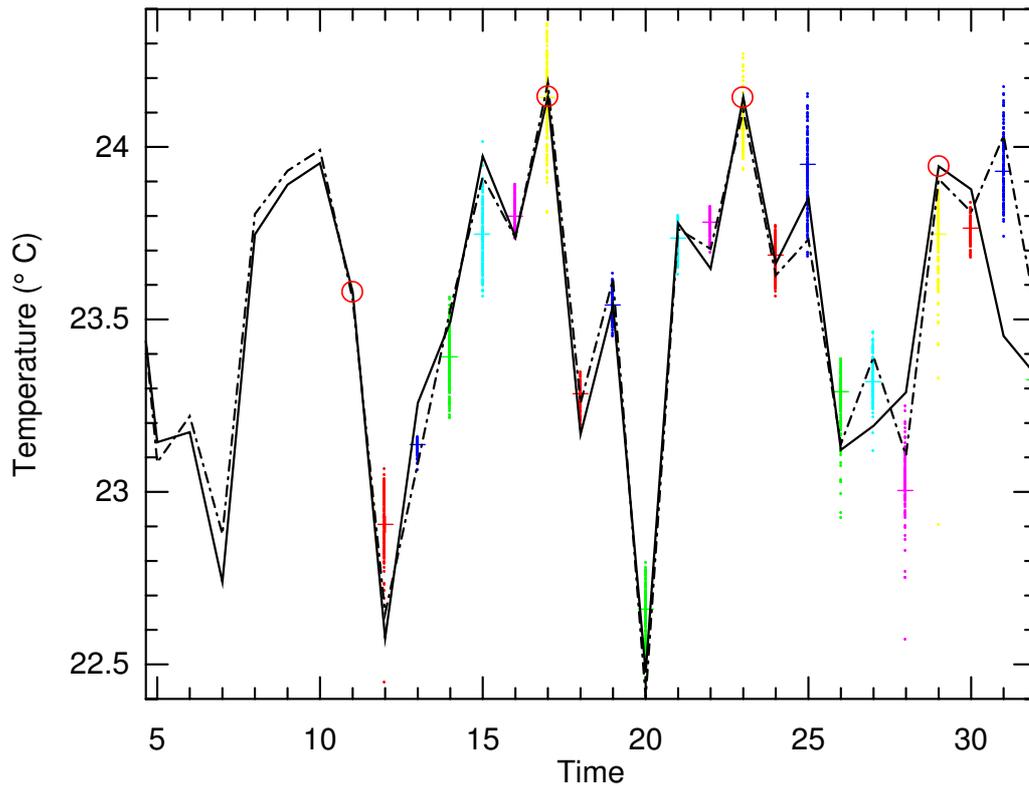


Figure 5.4: An observed temperature time series (solid line) from the annulus and an l -shadowing trajectory (dot-dashed line) from a RBF model which stays ‘close’ to (within 0.2°C of) the observed trajectory for 26 time steps (One time step is 48s). Four 6 step ensemble predictions are shown: an ensemble of 128 points, normally distributed within 0.075°C of (each component of) the initial observation, is initiated at times 11, 17, 23 and 29 (circles) and iterated under the model to give a distribution after 1 (red dots), 2 (dark blue), 3 (green), 4 (light blue), 5 (magenta) and 6 (yellow) steps, the mean of which is denoted (+). (Figure 9 in [84].)

The above examples demonstrate the difficulty of interpreting p.d.f.s from unconstrained ensembles, both due to inclusion of members which are not physically realisable states and to sparse sampling. Overcoming these problems may be achieved, with varying success, by constraining the ensemble members. For analytical low-dimensional systems the motivation to constrain members to the system invariant manifold is the desire for accountable p.d.f. predictions. In predicting physical systems, constraint of members to a lower dimensional model subspace is necessary if sampling is to be sufficiently dense for results to be meaningful.

5.1.2 Perfect ensembles

The discussion above illustrates that a perfect model and an exact understanding of the observational uncertainty distribution are necessary but not sufficient requirements for formation of a perfect ensemble. If the error in the predicted p.d.f.s are to be accountable to finite sampling uncertainties then ensemble members must be physically realisable states lying on the system invariant manifold and distributed according as the convolution of the system natural measure and the observational uncertainty distribution. Such perfect ensembles provide useful information until they are indistinguishable from the projection of the natural measure onto the observed variable, or climatology, to which they evolve asymptotically by virtue of the assumed ergodicity of the system.

If the observational uncertainty is quantised then a perfect ensemble has members distributed according as the natural measure within a hyper-cube $[\mathbf{x}_0 - \boldsymbol{\delta}, \mathbf{x}_0 + \boldsymbol{\delta}]$ where \mathbf{x}_0 is the observation and $2\boldsymbol{\delta}$ the quantisation level. The simplest way to practically form such an ensemble is that introduced by *Smith* [82], following the suggestion of *Lorenz* [45]: evolve the system and collect trajectories emanating from *analogue* points, contained within the hyper-cube and hence observationally indistinguishable; the corresponding trajectories are nevertheless distinct since they are recorded during evolution. This method assumes that analogue points collected sample the manifold without bias; an assumption violated if the numerical integration scheme used in the model iteration yields multiple consecutive values within the hyper-cube. Including all consecutive analogues as a single member, by weighting them accordingly, partially overcomes this problem and is implemented here. If the observational uncertainty has Gaussian distribution then a perfect ensemble may be formed by implementing the above process, using a hyper-cube suitably defined by the ensemble size and variance of the uncertainty, and then weighting the members according to the Normal probability density of the member. Using the system dynamics in this way provides

members of a perfect ensemble for which any errors in predicted p.d.f.s are accountable to uncertainty from finite sampling.

The above procedure was used to form perfect 64 member ensembles for the Marzec-Spiegel system with $\mathcal{U}(0.01)$ uncertainty distribution. Figure 5.2 b illustrates the evolution of ensembles initiated about the eight observations for which unconstrained ensemble evolution is illustrated in figure 5.2 a. The perfect ensemble p.d.f.s demonstrate variation in predictability as a function of initial condition (compare the spread of the p.d.f. at $t = 2 \downarrow_1$ with that at $t = 4 \downarrow_3$) and return of skill (*e.g.* at $t \simeq 0.85$), but are mainly coherent and compare well with the validating observations. Interpretation of the predicted p.d.f.s is now straightforward; given the observation at time t with a quantisation uncertainty of 0.01 in each component, the probability of a system state at time $t + 1$ is given by the predicted p.d.f. to within the sampling uncertainty. If the size of the ensemble is increased, the error in the p.d.f. decreases accountably.

Comparison of corresponding unconstrained and perfect ensemble predictions, figure 5.2 a and b respectively, highlights the variable error in the unconstrained p.d.f.s: unconstrained ensembles initiated at $t = 0, 1, 2, 3, 4$ yield predictions which differ considerably from their perfect counterparts due to the inclusion of members which do not lie on the manifold; those ensembles initiated at $t = 5, 6, 7$ provide more accurate (*i.e.* less in error, but not necessarily more compact) predictions as compared to those initiated at $t = 0, 1, 2, 3, 4$ ³. While the unconstrained predictions suggest that there is a higher probability of correctly predicting the observation at $t = 8$ than at $t = 1$, the perfect predictions demonstrate the reverse. Information provided by the perfect ensemble indicates that the bifurcation in the unconstrained p.d.f. at $t \simeq 1.4$ is caused by the inclusion of members not initially on the manifold and that it is not a real characteristic of the system; interpretation of unconstrained p.d.f.s is non-trivial whereas perfect ensembles provide accountable predictions.

Unfortunately the requirements of a perfect model, namely an exact understanding of the observational uncertainty distribution and (analytical or numerically generated) knowledge of the invariant manifold and the natural measure, are too demanding to be met for the prediction of physical systems. *Paparella et al.* [62] describe a method of random analogue prediction (RAP) utilising previously observed near neighbours as a step towards perfecting ensembles which requires little knowledge about the

³Increased accuracy may arise either from the manifold spanning more of the state space near the observation, in which case more members which lie on the manifold are included, or from the system visiting a region of state space in which perturbations off the manifold are rapidly damped.

system. Practical implementation of RAP is restricted to systems with short return times relative to the period of observation, thereby excluding many high-dimensional systems; alternative methods of constraint are discussed below.

5.1.3 Constrained vector subspaces

Aiming to predict true p.d.f.s is not viable for ensembles in high-dimensional models. Numerical weather prediction centres have progressed the development of low-dimensional vector subspaces to which small ensembles, limited in size by computational resources, may be constrained so as to provide useful information. The two vector fields dominating operational NWP ensemble formation are *singular vectors* (SV) developed and employed by the European Centre for Medium-range Weather Forecasting (ECMWF) [53], and *breeding vectors* (BV) developed and employed by the American National Centre for Environmental Prediction (NCEP) [94]; the respective motivations are to capture the spread of possible outcomes, and to include the best possible prediction in the ensemble. These motivations emanate from analytical results conditional on many assumptions: the definitions are given before discussing the implicit assumptions. When employed operationally, the CV perturbations are initialised: ECMWF rotates their SV perturbations in order to avoid localisation effects [14]; NCEP applies a mask to highlight regions of interest [95]. While initialisation procedures may (or may not) constrain the ensemble members to lie on the model attractor, it is *not* known if they lie on the (projection of the) system attractor; operational initialisation procedures do not yield ensembles with accountable error.

Evaluation of the performance of these ensembles is discussed below, while testing the internal consistency of the ensembles by quantifying the the validity of the assumptions made in construction is the subject of the next chapter. *Dream perturbations* and *Lyapunov vectors* (LV), more relevant to low-dimensional systems, are also defined to aid validation and to clarify motivation of BV formation respectively. The use of twin (equal and opposite) pair perturbations in the formation of NWP ensembles is standard [38]; it is introduced here as some definitions are dependent upon them and they are exploited in the following chapter.

5.1.3.1 Singular vectors

Singular vectors of a point \mathbf{x} and time τ_{opt} are the right singular vectors of the forward linear propagator of a system, \mathcal{M} , for a given optimisation time τ_{opt} . On application of SVD to yield

$$\mathcal{M}(\mathbf{x}, \tau_{\text{opt}}) = U\Sigma V^T, \quad (5.1)$$

the SV are the columns of V (for details see Appendix C). They represent the directions of the linearised system which grow most rapidly over $[0, \tau_{\text{opt}}]$. It is usual to rank the vectors according to the corresponding singular values, σ_i . In 2d a circle evolves into an ellipse under the linearised system; the first and second singular vectors are the perpendicular radial vectors of the circle which evolve to the major and minor axis of the ellipse respectively, as depicted in figure 5.5.

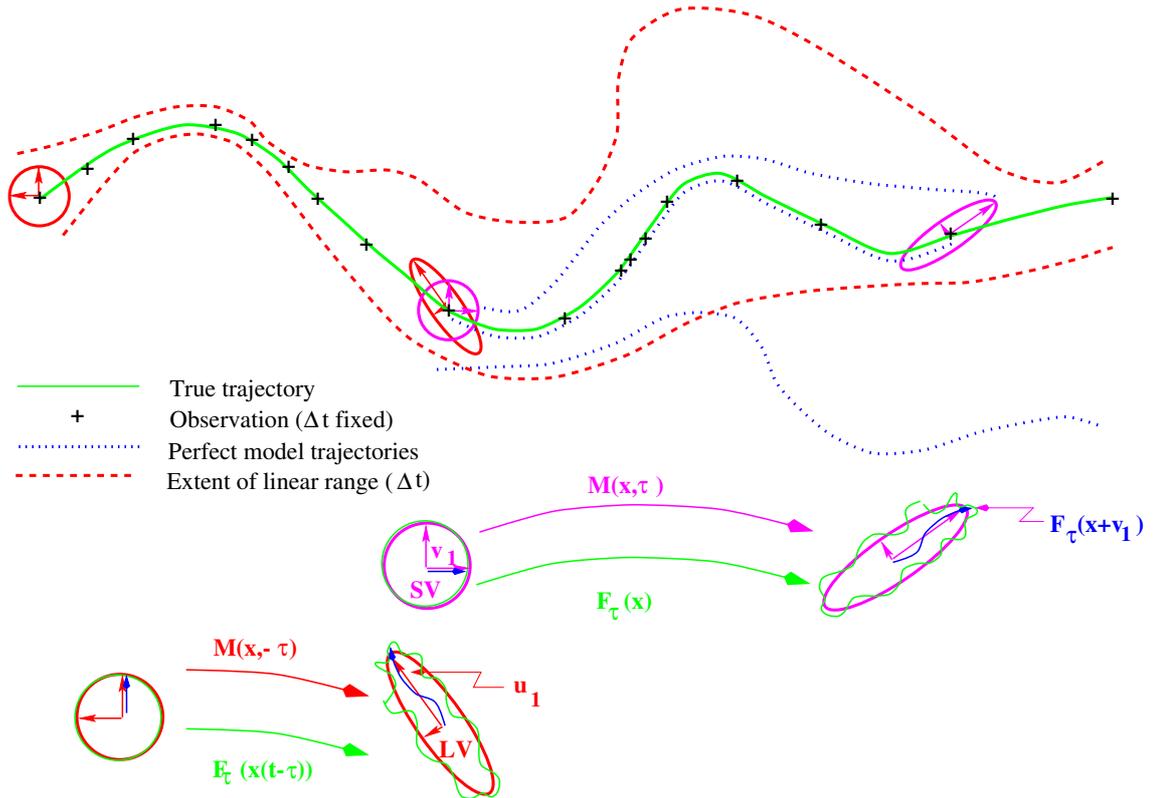


Figure 5.5: A schematic illustration representing a trajectory of the system (solid line) in its state space. Observations (plus signs) are at equally spaced intervals (Δt) in time. Evolution under the linear propagator reflects the growth of infinitesimal perturbations (the circle evolves into the ellipse); forward evolution defines the singular vectors (SV) for an optimisation time τ . The linear propagator is only approximate for finite uncertainties which evolve nonlinearly under the model, F , (the circle evolves into a ‘wavy’ closed curve near the ellipse); an initial magnitude can be defined so that the approximation error is bounded for any given fixed time (in this case Δt) defining a linear range (whose extent is indicated schematically by the dashed line) which will vary with location. The linear evolution over historical times yields the finite time Lyapunov vectors (LV). The three dotted lines are trajectories from “nearby” initial conditions.

Singular vector (SV) ensembles are defined as those which are restricted to a subspace spanned by the leading singular vectors of $\mathcal{M}(\mathbf{x}, \tau_{\text{opt}})$, where the optimisation time, τ_{opt} , is fixed. By constructing SV perturbations about the observation

(or analysis) of a magnitude comparable to the observational uncertainty (or analysis error), SV ensembles aim to capture the ‘worse case scenario’. Various assumptions are implicit in this motivation:

Denote by \mathbf{M}_0 the model representation of the system state \mathbf{S}_0 at time $t = 0$ (equivalent in a perfect model scenario), observed as (or with an analysis value) \mathbf{A}_0 to within the observational uncertainty (analysis error) $\boldsymbol{\delta}$.

1. Given \mathbf{M}_0 in a perfect model scenario, suppose we wish to find maximum spread of perturbations of magnitude $|\boldsymbol{\delta}|$ after time τ_{opt} using the SV of $\mathcal{M}(\mathbf{M}_0, \tau_{\text{opt}})$. The image of finite perturbations evolved under the fully nonlinear flow will in general, as indicated in figure 5.5, differ from their image under the linear propagator. The issue is whether or not this difference is significant since employment of SV ensembles assumes that the linearisation is a good approximation for the magnitude of perturbation used over the interval $[0, \tau_{\text{opt}}]$; specifically does

$$\tilde{f}^{\tau_{\text{opt}}}(\mathbf{M}_0) + \mathcal{M}(\mathbf{M}_0, \tau_{\text{opt}})(\boldsymbol{\delta}) \simeq \tilde{f}^{\tau_{\text{opt}}}(\mathbf{M}_0 + \boldsymbol{\delta})?$$

An acceptable approximation error bound for a given fixed time defines an initial magnitude which is the extent of the linear range (illustrated as both the dashed lines in figure 5.5 and as the circle in figure 5.6). Variation of error growth rate with location causes the extent of the linear range to vary with location also. Use of a fixed optimisation time requires satisfying the linearity assumption for the most rapidly growing regions, *i.e.* the regions of least linear range. An alternative is nightmare vectors, defined by *Smith* [82] and *Smith et al.* [85], formed by adjusting the optimisation time, up to a given maximum, for each initial condition by monitoring the growth rate of the perturbations. Further details of the formation and performance of nightmare vectors can be found in *Smith et al.* [85].

2. Still supposing a perfect model scenario, we actually observe (or calculate an analysis from the observations) \mathbf{A}_0 and construct the SV about this point, assuming that they approximate the SV at \mathbf{M}_0 . If this assumption is to hold then we require that

$$\mathcal{M}(\mathbf{A}_0, \tau_{\text{opt}}) \simeq \mathcal{M}(\mathbf{M}_0, \tau_{\text{opt}});$$

a necessary but not sufficient requirement is that \mathbf{M}_0 is contained within the linear range of \mathbf{A}_0 , as shown in figure 5.6.

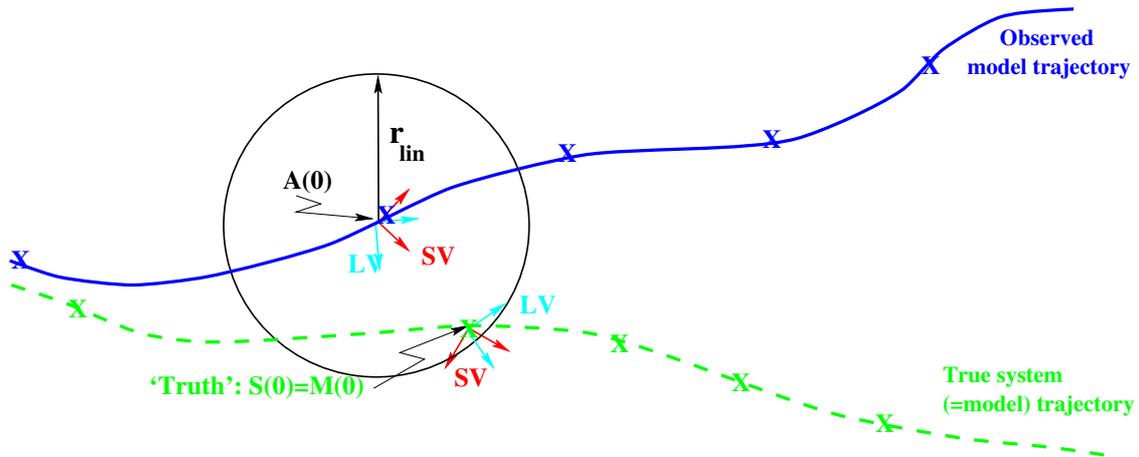


Figure 5.6: Schematic illustration, in perfect model scenario, of true model trajectory (green line) with corresponding singular vectors (SV, red) and Lyapunov vectors (LV, light blue) at $\mathbf{S}(0)$. SV are also shown about the observed value $\mathbf{A}(0)$, as is the model trajectory (blue line) emanating from this point and assumed to extend to $t = -\Delta t$, enabling the definition of LV at $\mathbf{A}(0)$. The extent of linear range, r_{lin} , at $\mathbf{A}(0)$ for τ_{opt} is also represented (black circle) and assumed here to contain $\mathbf{S}(0)$.

3. In general the model is imperfect (unlike figure 5.6), necessitating consideration of whether a projection of the system SV at \mathbf{S}_0 into the model state space exist. If one does, is it meaningful?
4. Similar to the model state to observation approximation above, if the projected system SV at \mathbf{S}_0 do exist and are meaningful, are they well approximated by the model SV at \mathbf{A}_0 ?

Operationally, only the linearity assumption, (1), may be tested, using the twin pair nature of the perturbations, detailed below in chapter 6. Assumption (2) is addressed below for the Marzec-Spiegel system and the rotating fluid annulus.

5.1.3.2 Lyapunov vectors

The global Lyapunov vectors, LV_∞ , of a point \mathbf{x} are the left singular vectors of the linear propagator $\mathcal{M}(\mathbf{x}, -\Delta t)$ in the limit $\Delta t \rightarrow \infty$, *i.e.* the columns of U on application of SVD (see equation 5.1). For the LV_∞ to be well defined, \mathbf{x} must lie on the invariant manifold so that there exists a trajectory along which to take the limit. If well defined, the LV_∞ represent the directions of sustainable growth, knowledge of which tells us little about local behaviour; in fact perturbations in the direction of the leading Lyapunov vector (that with largest corresponding singular value) may decrease with time for arbitrarily long finite time intervals. Finite time, or so called

‘local’, Lyapunov vectors (LV) are defined as for LV_∞ except for finite $\Delta t < 0$, in order to provide information about directions of locally (past) rapid growth (see figure 5.5) as described in *Smith et al.* [85] following the suggestion by *Lorenz* [46]. The requirement for \mathbf{x} to lie on the invariant manifold for the LV to be well defined still holds.

Lyapunov vector (LV) ensembles are defined as for SV: as those which are restricted to a subspace spanned by the leading (finite time) Lyapunov vectors of $\mathcal{M}(\mathbf{x}, -\tau_{\text{opt}})$. LV perturbations about the observation (or analysis) of a magnitude comparable to the observational uncertainty (or analysis error) then aim to capture the directions which have grown the most over $[-\tau_{\text{opt}}, 0]$. Again, various assumptions are implicit in this motivation. In a perfect model scenario \mathbf{A}_0 must lie on the manifold so that the LV are well defined, as well as providing a good approximation to those at \mathbf{S}_0 , as for the SV, illustrated in figure 5.6. Relevance of LV in the presence of model imperfections demand the same additional assumptions as for SV, namely that the projection of the system LV at \mathbf{S}_0 into the model state space must exist, be meaningful and be well approximated by the model LV at \mathbf{A}_0 .

5.1.3.3 Dream perturbations

What perturbation(s) would ideally be included in an ensemble? Usually one wishes to include as a member of the ensemble, an initial condition which reflects the future behaviour of the system as accurately as possible, given the imperfections in the model. ι -shadowing defines such an initial condition, if it exists, once future observations are available. A trajectory which ι -shadows over $[-\tau_1, \tau_2]$ is indistinguishable from the system given the corresponding observations. As τ_1 and τ_2 increase, the value of the trajectory at $t = 0$ will become a closer approximation of the (projected) system value given the model imperfections. With $P(1)$ such an initial condition will not be the observation (and is unlikely to be the corresponding analysis); the perturbation from the observation (or analysis) to such a ι -shadowing initial condition is defined as the *dream perturbation*, since such an initial condition will yield a best prediction under the given model providing that τ_2 is sufficiently large, see figure 5.7. If the dream perturbation is taken relative to the analysis value then it approximates the analysis error, and we may wish to include such a perturbation in an ensemble. Finding the dream perturbation requires knowledge of future observations so it does not provide a direct means of ensemble formulation for prediction. Knowledge of the dream perturbation does however facilitate the investigation of whether systematic

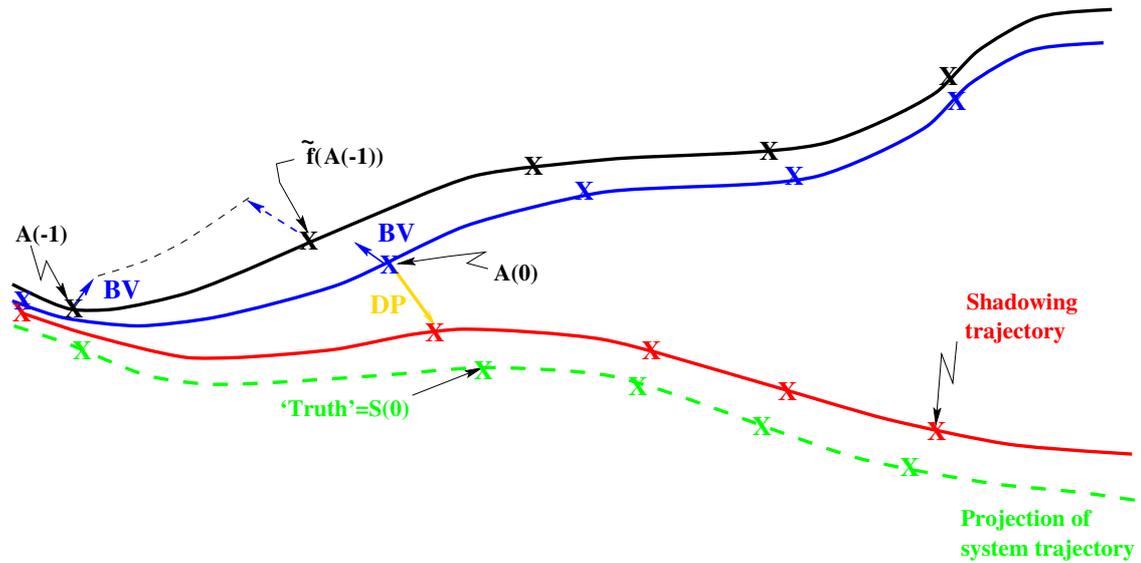


Figure 5.7: The projection of a system trajectory (dashed green line) into the imperfect model domain is best approximated by an ι -shadowing trajectory (red solid line). The dream perturbation (DP, orange arrow) is the perturbation from the observation (analysis) $\mathbf{A}(0)$ to the shadowing trajectory value at that time. A breeding vector (BV, blue arrow) is bred by iterating the BV at $\mathbf{A}(-1)$ forwards under the model \tilde{f} (to dotted blue arrow), rescaling the vector and moving it to the new observation (analysis) $\mathbf{A}(0)$ etc.

information about such a perturbation is contained in any operational constrained vectors and provides an estimate of the analysis error.

5.1.3.4 Breeding vectors

Unlike the preceding constrained vectors, breeding vectors (BV) are not based upon the application of theoretical definitions, but were developed with the aim of capturing the directions in which the analysis error may grow, using only past observations and model dynamics [94]. Data assimilation to update analysis values is performed over a cycle time τ_c . In order to ‘breed’ the growing (analysis) error ‘modes’, a perturbation of arbitrary direction and fixed magnitude, ϵ_{BV} , comparable to that of the believed analysis error, is added to the analysis at some time τ . Both the analysis and perturbed point are iterated forwards under the model, and the difference between the iterated points provides the orientation of the new BV perturbation, see figures 5.7 and 5.8 *i*. The procedure is repeated, adding a perturbation in this orientation and with magnitude ϵ_{BV} as before, to the analysis at time $\tau + \tau_c$ and both iterated forwards, their difference taken as the new orientation, etc. It is intended that after several cycles of breeding the BV field is some composite of the local directions in which there is recent fast error growth [94]; if the model is perfect and the perturbation

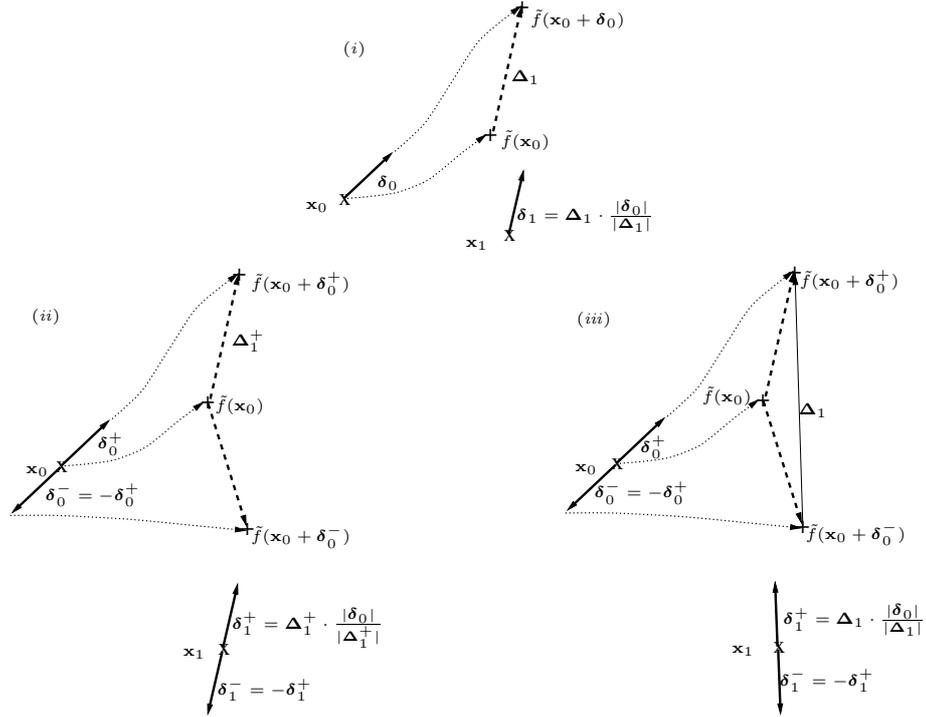


Figure 5.8: Schematics of breeding. (i) The ideas of breeding vectors: A perturbation δ_0 (solid line at $t = 0$) is added to the analysis \mathbf{x}_0 at initial time and both are iterated forward under the model (dotted lines). The evolved perturbation Δ_1 (dashed line) is then rescaled to the magnitude $|\delta_0|$ of the initial perturbation to give the new BV δ_1 (solid line at $t = 1$) which is used as the perturbation from the new analysis \mathbf{x}_1 . If the dynamics are linear for magnitude $|\delta_0|$ of perturbation, then twins of BV may be generated by considering $\pm\delta_0$ which will evolve to $\pm\Delta_1$ and be rescaled to $\pm\delta_1 \dots$. If the dynamics are not linear there is a choice of how to generate BV twins. (ii) Method A: Chose one of the evolved twins (here the positive) over the other, re-scale and introduce its symmetric image. (iii) Method B: Alternatively the difference between the evolved twins may be rescaled and its symmetric image introduced. Method B is employed operationally.

infinitesimal the BV evolve toward the first (finite time) Lyapunov vector. Note however, that the definition of BV does not depend on the analysis values lying on the invariant manifold. BV are bred from finite perturbations for an imperfect model using the model dynamics to give information about the analysis error, hence they also aim to approximate the dream perturbation. The set of BV obtained will depend on the magnitude ϵ_{BV} , and to some extent on the initial direction of the perturbations. It is claimed that the subspaces spanned by BV initiated by different perturbations have ‘significant similarities’ after a few breeding cycles [37, 94]. Houtekamer and Derome [37] suggest that the BV magnitude, operationally held constant, should be varied according to the estimated magnitude of the analysis error, which varies with location, if the BV are to optimally capture the growing ‘modes’ of the analysis error.

Operational implementation of BV ensembles, *e.g.* by NCEP [94], differs slightly: initially a perturbation is both added to and subtracted from the analysis and the two perturbed values iterated to give the BV. There is then a choice of how to define the new BV orientation: either as the difference between either the positive or negative perturbation (method A), or as the difference between the two iterated perturbations (method B), illustrated in figure 5.8 *ii* and *iii* respectively. Whichever method is used, the new BV is then added to and subtracted from the next analysis value and the process repeated. These two methods are equivalent if evolution is linear over time τ_c for the magnitude ϵ_{BV} . The importance of nonlinear growth effects are investigated in the next chapter. Method B is employed operationally in NWP by NCEP.

5.2 Performance evaluation

The definitions of constrained vector subspaces above are motivated by differing aims of the resulting ensembles. While perfect ensembles will always yield p.d.f.s accurate to within sampling uncertainties, errors in unconstrained ensembles will be system dependent, and will vary with location within the phase space of any system. Given the definitions of constrained vector subspaces there are two questions which need to be considered in evaluating the performance of the CV ensembles: how successful are the ensembles at achieving the aims for which the subspaces were defined, and are the assumptions made in definition of the ensembles valid? The former we address for systems where data is relatively plentiful and where we either know the exact system value or can approximate it by appropriate use of ι -shadowing. Testing the relevance of some of the approximations also requires some knowledge of the true state and these tests are discussed here. Investigation of the validity of the linearity

assumption, which pertains to both singular vector and breeding vector subspace definitions, is facilitated by the use of twin pair perturbations in operational NWP ensembles, and is discussed in the next chapter.

After introducing the practical details used in CV ensemble construction, we describe the measures which may be used, given the information available, to evaluate CV ensemble performance and to establish the validity of certain assumptions. Results are presented for CV ensembles formed using a perfect model of the low-dimensional Marzec-Spiegel system, and using an imperfect RBF model of the annulus.

5.2.1 CV ensemble evaluation measures

Operationally, ensembles are formed about analysis values \mathbf{A}_0 , which are constructed using past observations and model dynamics, as described in section 4.4.2. The use of analysis values is implemented here since the BV are defined with the aim of capturing analysis error rather than the observational error. Given the observations, observational uncertainty and model, we define the analysis to be the value at $t = 0$ of the shadowing trajectory which shadows over at least $t \in [-\tau_a, 0]$ for a prescribed τ_a , usually giving an analysis value ‘closer’ to the exact system value (given the model dynamics) than the observation, as illustrated in figure 5.9. This definition ensures that LV (of the analysis) are well defined for $\tau_{LV_{opt}} \leq \tau_a$. If the model representation, \mathbf{M}_0 , of the exact system state of the system, is known (or able to be approximated) we may form *system singular vectors* (SSV) of the exact system value as well as singular vectors of the analysis (SV), both for a prescribed optimisation time, $\tau_{sv_{opt}}$. Breeding vectors are constructed using method A (BVA) and/or method B (BVB) for a prescribed cycle time, τ_{BV_c} . Disjoint sections of continuous observed system trajectories are used, thus the initiation of BV is only required once for each trajectory: BV perturbations (if multiple BV are required) are orthogonal at inception and are then ‘pre-bred’ over $\tau_{BV_{init}}$ steps before their first implementation to allow ‘alignment’ of the BV with the fastest growing ‘modes’ of the analysis error. CV ensembles are then formed by adding and subtracting perturbations $\mathbf{CV}_n, n \in [1, CV_{max}]$, where CV_{max} is the number of CV required, of magnitude ϵ_{CV} , to the analysis. (SSV perturbations are added to (subtracted from) the system state.)

The *local manifold* is defined by applying local singular value decomposition (see Appendix C and [9]) to the local points on the system manifold; points on the system manifold are assumed to be approximated by the observations when system values are not known. For purposes of ensemble evaluation, ‘truth’, \mathbf{S}_0 , is taken as the

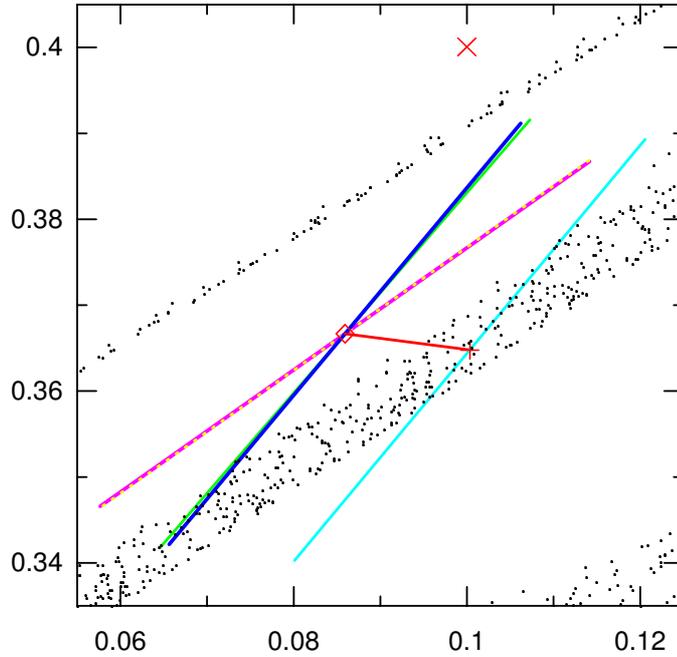


Figure 5.9: Constrained vectors (CV) of magnitude 0.02 (in each component) for an initial condition of the Marzec-Spiegel system, projected onto the (x, y) plane. The analysis value (red diamond) lies on a shadowing trajectory extending over at least $t \in [-2, 0]$ and provides a better estimate of the exact system value (red plus sign) than the observation (red cross). The dream perturbation (red line) represents the analysis error. Also shown are twin (equal and opposite) pairs of a finite time Lyapunov vector (LV) (green line, with $\tau_{LV_{opt}} = 2$), singular vector defined about the analysis (dark blue) and system value (light blue) (both with $\tau_{SV_{opt}} = 2$), and breeding vector defined using method A (magenta) and method B (yellow dotted) (both with $\tau_{BV_{init}} = 10, \tau_{BV_c} = 1$), defined using a perfect model of the system. The relation of the CV to the organised structure of the local system manifold (black dots) can be seen. While the BV reflect the manifold much better than the SV, note that an initialisation method which chooses as ensemble members the nearest point on the manifold may yield an SV ensemble preferable to the BV ensemble (since the SV ensemble might contain members on two sheets of the manifold).

system state when known, and otherwise as the value of the shadowing trajectory which shadows over at least $t \in [-\tau_1, \tau_2]$, for prescribed τ_1, τ_2 , when such a trajectory can be found. The dream perturbation is then defined as the perturbation from the analysis to the ‘truth’. Figure 5.9 illustrates the observation, analysis, ‘truth’, constrained vectors, dream perturbation and local manifold for an initial condition given a perfect model of the Marzec-Spiegel system.

Some of the evaluation measures require comparison of subspaces; similarity is quantified by summing appropriate inner products of vectors describing orthonormal bases of the subspaces of interest to give a projection index. For an md model, form two $m \times m$ matrices, A and B , the first d_A and d_B (respectively) columns of which comprise the orthonormal bases of the subspaces of interest, and which are otherwise filled by zeros. A projection matrix P is then calculated, with entries $P_{ij} = (A_{ik}, B_{kj})$ (using summation convention), where (\cdot, \cdot) is the inner product. The projection index for the projection of A onto B is then the sum of squares of the rows of P divided by d_A , while that for the projection of B onto A is the sum of squares of the columns of P divided by d_B . Such a projection index yields results expected intuitively: the projection of a subspace with basis of rank l onto one with basis of rank n has a maximum projection of 1 if $l > n$ and of l/n if $l \leq n$. The similarity between two subspaces is then taken as the projection of one subspace onto the other, if their bases are of the same rank; if not, the average of the projection index of subspace of constrained vector A (CVA) onto that of constrained vector B (CVB) and the projection of subspace CVB onto that of CVA is used. Orthonormal bases of SV, SSV and LV subspaces are available by virtue of the use of singular value decomposition in formation. When required, orthonormal bases of BV and local manifold subspaces may be similarly obtained by use of singular value decomposition. Form a $m \times BV_{\max}$ matrix B (where m is the dimension of the model considered), whose rows are transposed BV, thus

$$B = \frac{1}{\sqrt{m}} \begin{bmatrix} \mathbf{BV}_1^T \\ \mathbf{BV}_2^T \\ \vdots \\ \mathbf{BV}_{BV_{\max}}^T \end{bmatrix}, \quad (5.2)$$

and similarly for local manifold point vectors. Then application of SVD to the covariance matrix, $C = B^T B = V \Sigma^2 V^T$, yields an orthonormal basis, the columns of V , optimal in the least squares sense of variance. Further, the effective rank of the basis may be determined from the associated singular values [10] (for more details see section C.2).

5.2.1.1 Evaluation of SV ensembles

As explained in section 5.1.3.1, singular vector ensembles aim to capture the ‘worse case scenario’. Evaluating whether they succeed in this aim is achieved by comparing the *spread* of the evolved SV perturbation with the real spread, given by the magnitude of the evolved dream perturbation, at optimisation time, $\tau_{\text{sv,opt}}$; specifically, comparing

$$|\tilde{f}^{\tau_{\text{sv,opt}}\text{SV}}(\mathbf{A}_0 + \mathbf{SV}_n) - \tilde{f}^{\tau_{\text{sv,opt}}\text{SV}}(\mathbf{A}_0)| \text{ with } |\tilde{f}^{\tau_{\text{sv,opt}}\text{SV}}(\mathbf{M}_0) - \tilde{f}^{\tau_{\text{sv,opt}}\text{SV}}(\mathbf{A}_0)|,$$

$n = 1, \dots, SV_{\text{max}}$. If multiple SV are used, then the spread for each SV is calculated and the average taken⁴. SV spread equal to or exceeding real spread indicates possible success of the associated ensemble in capturing the ‘worse case scenario’; other CV spread are calculated for comparison.

Of the assumptions implicit in formulation of SV ensembles, described in section 5.1.3.1, the latter two depend on the model and system of interest, and require exact knowledge of both for validity to be ascertained. The validity of the first assumption, that the linear approximation is valid for perturbations of magnitude ϵ_{sv} for the optimisation time $\tau_{\text{sv,opt}}$, can be addressed without knowledge of the system, providing that twin pair perturbations are used; this is discussed in detail in the following chapter for models of the Marzec-Spiegel system, the annulus and NWP. The second assumption is that the singular vectors of the analysis (SV) approximate those of the ‘truth’ (SSV). The validity, or otherwise, of this assumption is quantified by the similarity between the two vector subspaces, with a similarity index (as defined above) near 1 indicating validity of the assumption.

5.2.1.2 Evaluation of BV ensembles

Breeding vectors were developed with the aim of capturing the directions in which the analysis error may grow, thus including in the ensemble the initial condition closest to the ‘truth’; we expect the BV to have some projection onto the dream perturbation which here is the analysis error. Further, both the BV and dream perturbations use model dynamics and past observations (due to the use of ι -shadowing to define the analysis) in their definition. If the model is perfect and the perturbation infinitesimal, the BV evolve toward the first (finite time) Lyapunov vector. To measure the success of BV in achieving their aim, we measure the similarity between BV and LV

⁴While an average may appear to unfairly penalise a pair of SV perturbations in which one grows sufficiently to capture the verification while the other does not, this scenario is seldom seen in practice.

subspaces to see how well the BV capture the most rapidly past growing (local) directions (a similarity index near 1 denoting success), and the projection of the dream perturbation onto the BV subspace to see how well the BV capture the analysis error (a projection index of 1 indicating that the dream perturbation is contained in the BV subspace). The projection of the dream perturbation onto other CV subspaces is also calculated for comparison.

A further measure of BV achievement is the frequency with which the region between the evolved analysis and evolved BV perturbation (as calculated for the spread measure above) includes the verification, taken as the true value at the evolved time; specifically, how often is

$$(\mathbf{M}_t)_l \in [\tilde{f}^t(\mathbf{A}_0)_l, \tilde{f}^t(\mathbf{A}_0 + \mathbf{B}\mathbf{V}_n)_l] \forall l \in [1, m]? \quad (5.3)$$

Inclusion of the verification in such a region suggests that the BV ensemble may include as a member the initial condition ‘closest’ to the system value which gives the best possible prediction under the model. Again, this is also calculated for the other CV for comparison.

The validity of the linearity assumption, made operationally by the use of method B, is discussed in the following chapter. Here, however, we compare the similarity between BVA and BVB; if they differ then the linearity assumption is not valid. The results of other measures for the two methods are compared when they differ to see if either method is consistently superior.

5.2.1.3 Other measures of evaluation

Comparison of perfect ensemble p.d.f.s with those from unconstrained ensembles (figure 5.2 in section 5.1.1 on page 80) demonstrates the importance of ensemble members lying on the manifold if the ensemble is to accurately reflect possible behaviour of the system, within the constraints of model error. CV ensemble members are given by perturbations from the *analysis* in the directions of the CV. Only if the analysis is defined such that it lies on the manifold (and, further, lies on a section of the manifold ‘connected’ to the system state) will similarity between the CV subspace and the local manifold be a valid indication of whether ensemble members may lie on the manifold. As illustrated by the dream perturbation in figure 5.9, if the analysis does not lie in the manifold (or if lies on a section of the manifold not connected to the system value), then perturbations to points which do lie on the manifold will not lie in the manifold! While the dream perturbation is by no means the only direction for

perturbations from the analysis onto the manifold, good projection of the dream perturbation onto the CV subspace suggests that the CV ensemble may include members which do lie on the manifold *regardless of the position of the analysis in relation to the manifold*. Projection of the CV onto the (leading directions of the) local manifold are included here for completeness. The similarity between SV and BV subspaces are also given for interest.

5.2.2 Marzec-Spiegel system ensemble evaluation

Observations of the Marzec-Spiegel system, with quantisation $\mathcal{U}_3(0.1)$ error, were used with a perfect model to form analysis values using shadowing with $t_a = 2$, 75% of which had analysis error in $\mathcal{U}_3(0.04)$; exact system values were used as ‘truth’. The local manifold is effectively 2d for most of the (995) initial conditions used. Measures were calculated using the following parameters, unless otherwise stated: $CV_{\max} = 1, \epsilon_{CV} = 0.02, \tau_{LV_{\text{opt}}} = \tau_{SV_{\text{opt}}} = 2, \tau_{BV_{\text{init}}} = 10, \tau_{BV_c} = 1$. Measures were also calculated using five alternative parameter values, namely (1) $\epsilon_{CV} = 0.005$; (2) $\epsilon_{CV} = 0.1$; (3) $\tau_{SV_{\text{opt}}} = 4$; (4) $\tau_{BV_{\text{init}}} = 5$; and (5) $CV_{\max} = 2$, in order to examine the parameter value’s effect, or lack thereof, on the measures.

5.2.2.1 Similarity between CV bases

The similarity index between SV and SSV bases was approximately 1.0 (indicating almost exact agreement) for the standard parameter values, and also for alternatives (1,2,3,5); the assumption that the SV approximate the SSV is valid here. The similarity between the two alternative BV subspaces, and of the BV and LV subspaces, is shown in table 5.1. The two BV construction methods produce similar BV for small

$\epsilon_{CV}, CV_{\max}, \tau_{BV_{\text{init}}}$	BVA·BVB	BVA·LV	BVB·LV
0.005 or 0.02, 1, 10	0.99	0.58	0.57
0.1, 1, 10	0.89	0.55	0.57
0.02, 2, 10	0.99	0.63	0.63
0.02, 2, 5	1.00	0.58	0.57

Table 5.1: Similarity indices between BV and LV subspaces for the Marzec-Spiegel system.

perturbations, but when the perturbation magnitude increases to 0.1 differences appear; the assumption of linearity is not valid for $\epsilon_{BV} = 0.1$. The similarity index between BV and LV subspaces is quite high when single CV are considered. While the index increases slightly when 2 CV are used; recall that any two 2d subspaces of a 3d system have a similarity index of at least 1/3.

Similarity between SV and BV bases was about 0.25 for the standard parameter values and alternatives (1,2) in which perturbation magnitude is varied; this increases to about 0.4 when 2 CV are considered, alternative (5).

5.2.2.2 Spread

The ratios of CV spread to the ‘real spread’ (the magnitude of the difference between the evolved analysis and evolved ‘truth’) are shown in figure 5.10. The peak for the BV ratio occurs at a value slightly less than 1, while that for the SV and LV ratios is centred on 1; all the CV ratio histograms have considerable distribution for values greater than 1. The effect of varying the perturbation magnitude on these results is

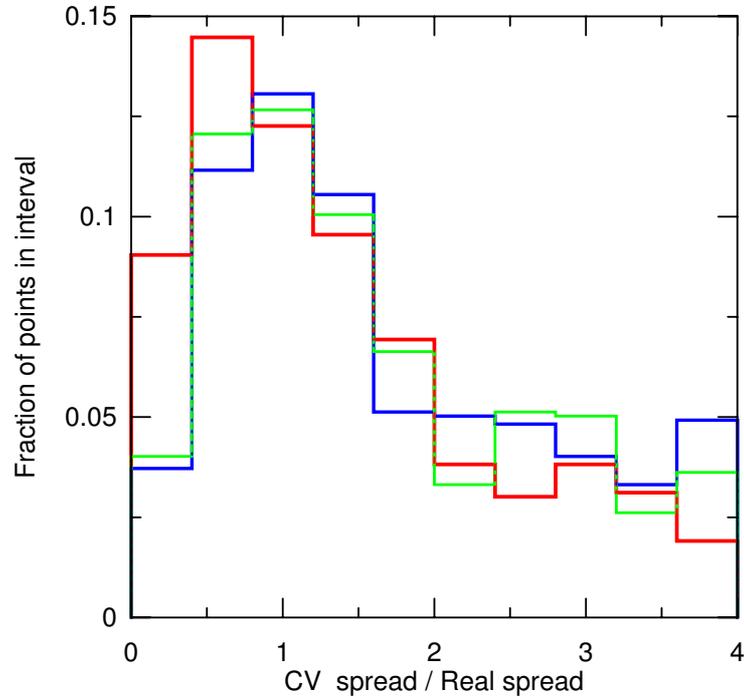


Figure 5.10: A histogram of the ratios of constrained vector spread to real spread for SV (blue), BVA (red) and LV (green), calculated using a perfect model of the Marzec-Spiegel system.

shown by similar histograms for SV and BV in figure 5.11 a and b respectively. The SV results show that SV ensembles are able to capture the spread if the perturbation amplitude is of similar magnitude to that of the analysis error. If the perturbations are too small, real spread is not captured, while if too large the SV spread drastically overestimates the real spread. LV spread to real spread ratios behave similarly to those of SV, while those for BV peak at less than 1 unless the perturbation magnitude

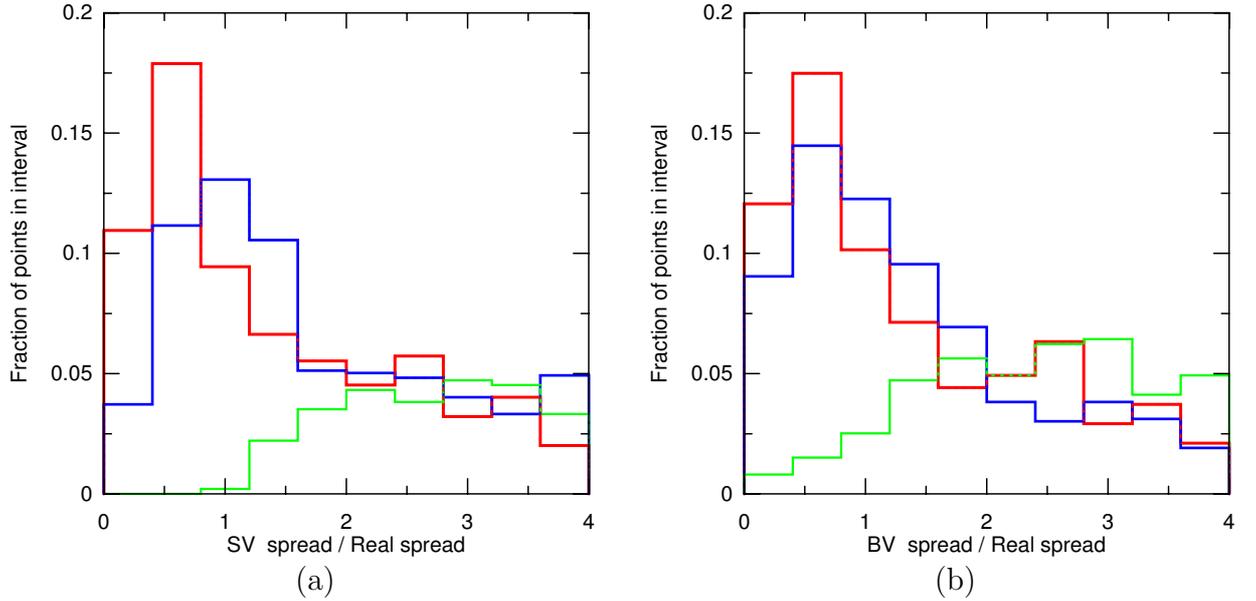


Figure 5.11: A histogram of the ratios of (a) SV spread to real spread and (b) BV spread to real spread, for perturbation magnitudes 0.005 (red), 0.02 (blue), 0.1 (green), calculated using a perfect model of the Marzec-Spiegel system.

is larger than the analysis error ($\epsilon_{BV}=0.1$), when the real spread is again drastically over estimated.

The percentage of cases for which the region between the evolved analysis and the evolved CV perturbation encompass the verification are shown in table 5.2. Evolved

ϵ_{CV}	SV	SSV	LV	BVA	BVB
0.005, 0.02	96	87	76	64	54
0.1	93	85	71	61	52

Table 5.2: Percentage of cases for which evolved CV perturbations encompass the corresponding verification for the Marzec-Spiegel system.

SV perturbations usually encompass the verification. It is interesting to note that while increasing the perturbation size to 0.1 increases the spread ratio, the resulting evolved perturbations encompass the corresponding verification less frequently for all the CV. SSV encompassing percentages are slightly less than those of SV; this result does not conflict with similarity between SV and SSV subspaces since the SV perturbations are from the analysis, while those of the SSV are from the system state. Method A BV perturbations (formed by taking one evolved twin) consistently encompass the verification more frequently than those constructed by method B.

Figure 5.12 is a scatterplot of CV spread vs. real spread, with symbols denoting whether the verification value is included in CV spread or not. Failure to encompass

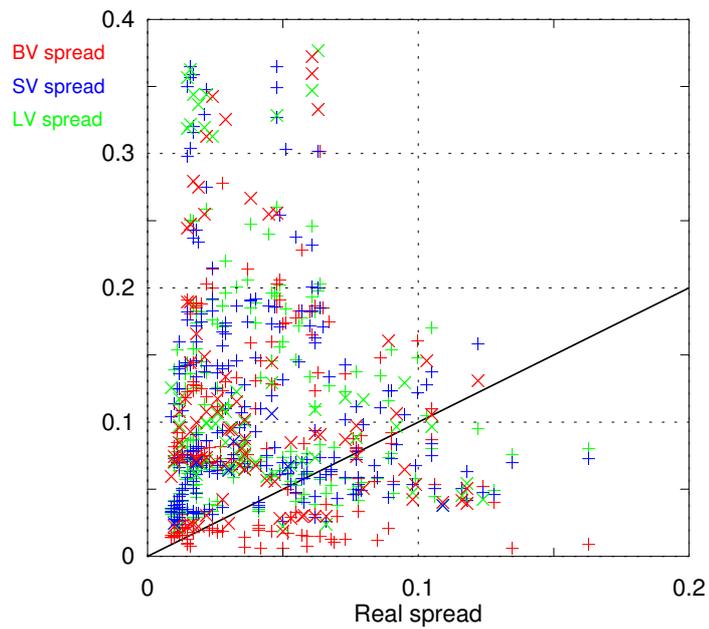


Figure 5.12: A scatterplot of the ratios of CV spread to real spread for SV (blue), BVA (red) and LV (green), for 250 initial conditions, calculated using a perfect model of the Marzec-Spiegel system. Whether the CV perturbation includes the verification (plus sign) or not (cross) is indicated. The line $y = x$ (black solid line) is shown for reference.

the verification appears to be correlated to large evolved spread for BV and LV, but not for SV.

5.2.2.3 Projection of dream perturbations onto CV subspaces

The dream perturbation represents an orientation for perturbations from the analysis onto the local manifold. While such an orientation is by no means unique, good projection of the dream perturbation onto the CV subspace suggests that the CV ensemble may include members which lie on the manifold. Figures 5.13 and 5.14 are scatterplots of the projection of dream perturbation into the CV basis, for $CV_{\max} = 1, 2$ respectively, comparing all pairs of (SV, LV, BVA). For clarity, results from only 250 initial conditions are illustrated in the scatterplots; average values of dream perturbation projection into CV basis are given in table 5.3. Results shown in the

CV_{\max}	DP·SV	DP·LV	DP·BVA/BVB
1	0.42	0.41	0.25
2	0.79	0.79	0.26

Table 5.3: Projection indices of dream perturbation (DP) onto CV subspaces for Marzec Spiegel system.

table show that projection of the dream perturbation onto SV and LV subspaces is similar, and significantly higher than that onto BV subspaces. Varying perturbation magnitude has little effect on the results. Increasing the number of CV used to two almost doubles projection indices for SV and LV, suggesting that the second SV and LV directions capture components of the dream perturbation not explored by the primary CV direction, yet the projection index onto 2d BV subspaces changes little; use of SVD ensures that this is *not* due to near linear dependency between the two BV.

The scatterplots of figures 5.13 and 5.14 provide the additional information that, when only one CV is used, the projection index is usually very small or quite large, but does not often taken values near 0.5. When two CV are used high projection indices are more common (at least for SV and LV). Correlation between SV and LV results is indicated by the high density of DP·LV vs. DP·SV points near the diagonal; high projection indices onto LV or SV subspaces has no obvious correlation with projection onto the BV subspace.

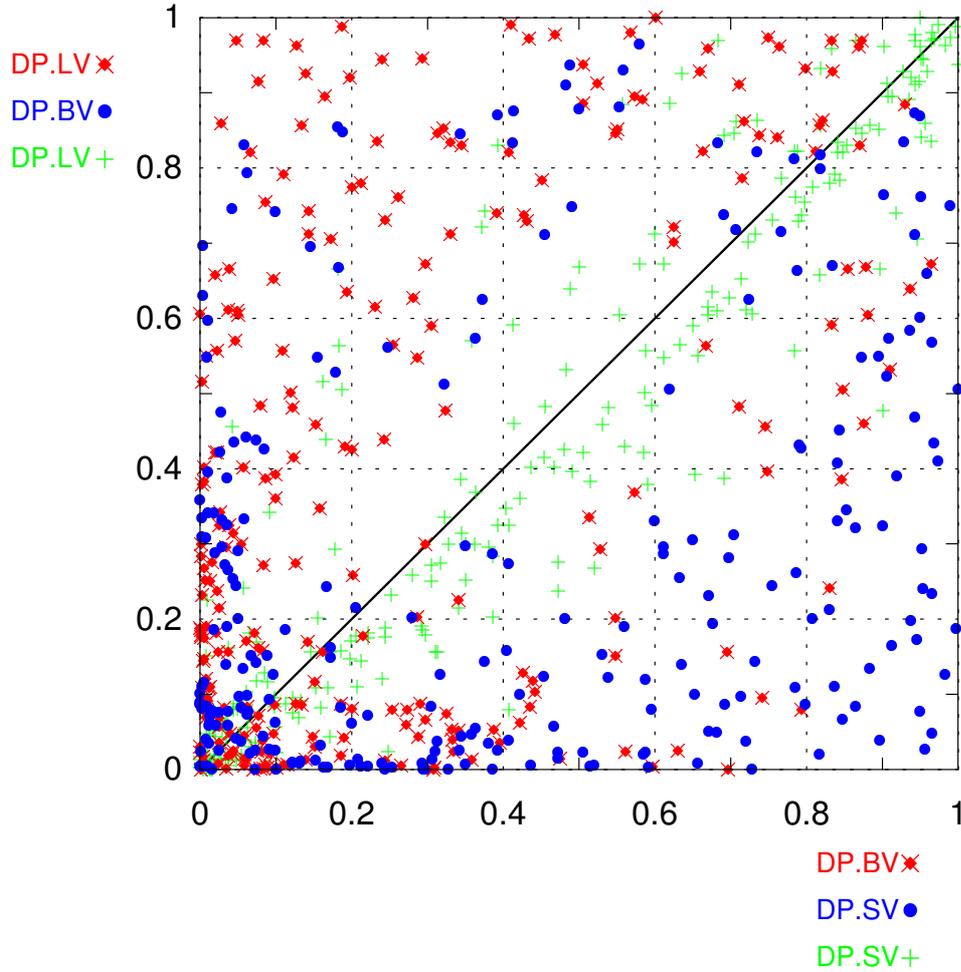


Figure 5.13: A scatterplot of the projection index of dream perturbation (DP) onto 1d CV subspace for all pairs of (SV, LV, BVA), for 250 initial conditions, calculated using a perfect model of the Marzec-Spiegel system: LV vs. BVA (red diamond cross), BVA vs. SV (blue bullet) and LV vs. SV (green plus sign). The line $y = x$ (black solid line) is shown for reference.

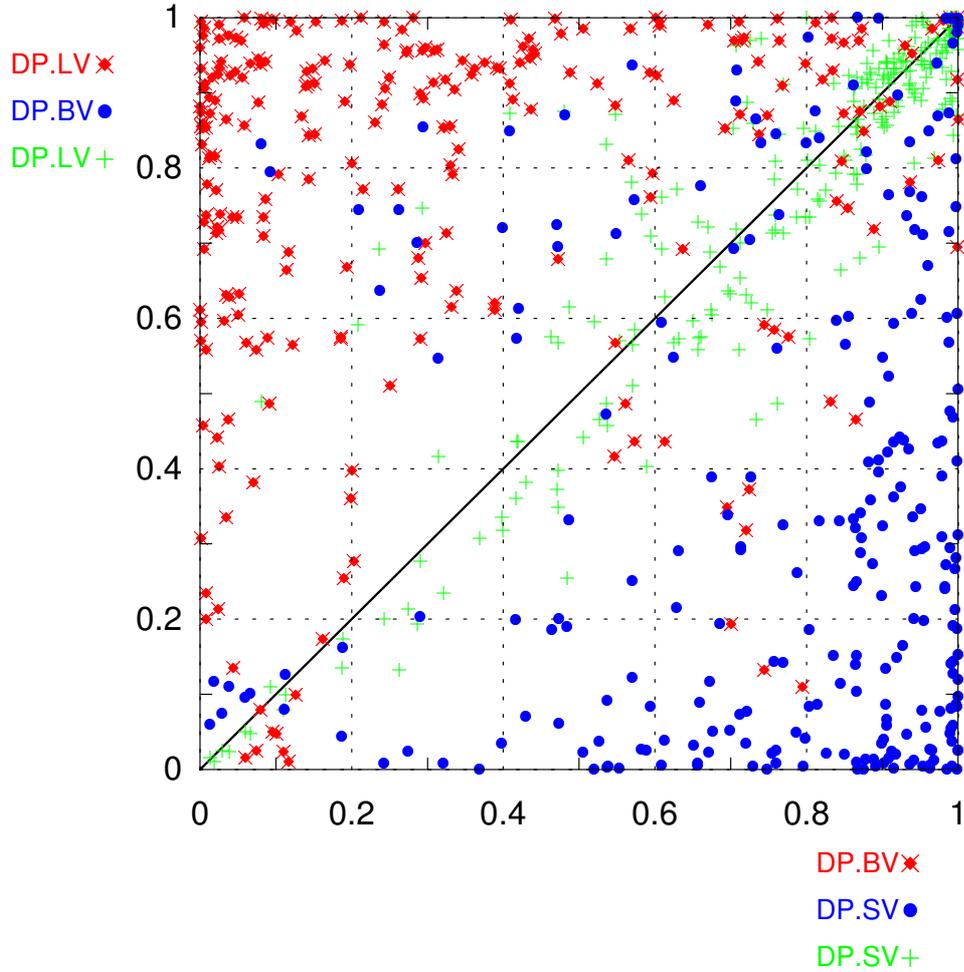


Figure 5.14: A scatterplot of the projection index of dream perturbation (DP) onto 2d CV subspace for all pairs of (SV, LV, BVA), for 250 initial conditions, calculated using a perfect model of the Marzec-Spiegel system: LV vs. BVA (red diamond cross), BVA vs. SV (blue bullet) and LV vs. SV (green plus sign). The line $y = x$ (black solid line) is shown for reference.

5.2.2.4 Projection of CV onto local manifold

Projections of CV are taken onto the CV_{\max} directions of the local manifold which capture the most variance, as described in section C.2 and [9]. The projections of CV onto the leading local manifold direction are summarised in figure 5.15 for the standard parameter values. Similar to the dream perturbation projection results,

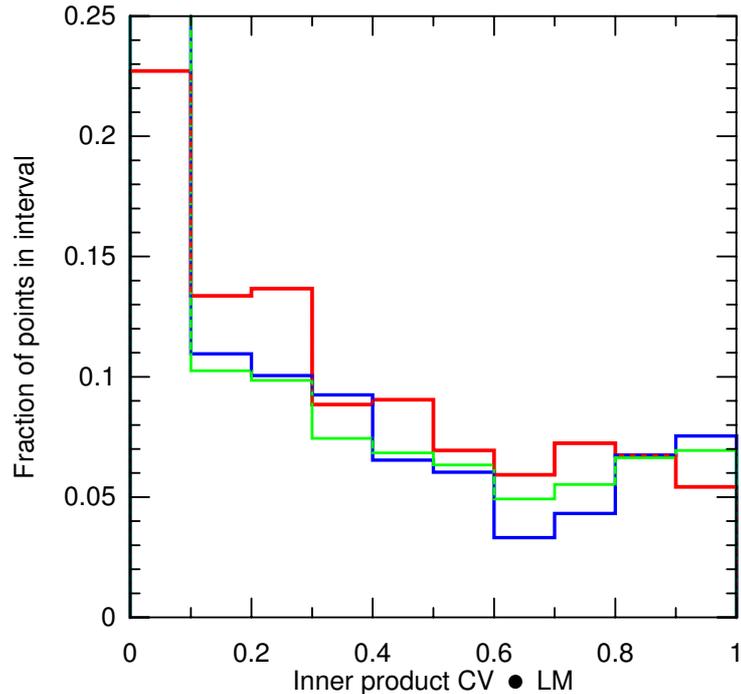


Figure 5.15: A histogram of the projection index of constrained vector subspace onto leading local manifold directions for SV (blue), BVA (red) and LV (green), calculated using a perfect model of the Marzec-Spiegel system.

increasing CV_{\max} to 2 (and projecting CV basis into the first two leading local manifold (LM) directions) almost doubles the SV·LM and LV·LM projections, while the BV·LM projection is comparable to that for $CV_{\max} = 1$.

5.2.2.5 Summary of CV performance

From the results above we can evaluate the performance of CV ensembles for the Marzec-Spiegel system by considering how well the CV fulfill the aim for which they were designed, and by considering the validity of the assumptions made for the perturbations used. SV perturbations usually capture the extent of the real spread providing that the perturbation magnitude is comparable to that of the analysis error; they almost always encompass the corresponding verification. The assumption that the SV approximate the SSV is valid.

Dream perturbations do not project particularly well onto BV subspaces, especially in comparison to the projection of dream perturbation onto other CV subspaces. There is considerable similarity between BV and LV subspaces. BV perturbations underestimate real spread for perturbation magnitudes comparable to that of the analysis error, and drastically overestimate it if the perturbation magnitude is increased. Regardless of the perturbation magnitudes however, the BV perturbations do encompass the verification in about half the cases considered; the BV perturbations always encompass the verification less frequently than the SV perturbations. The two methods of BV construction yield almost identical subspaces for perturbation magnitudes upto that of the analysis error, but differ when the perturbation magnitude is increased further, indicating that the linearity assumption does not hold for the latter case. Method B BV are marginally more similar to LV for larger perturbation amplitudes, while method A BV perturbations consistently encompass the verification more frequently; projection indices of the dream perturbation onto the two BV subspaces defined by the different methods are indistinguishable.

5.2.3 Annulus ensemble evaluation

Observations of the annulus are assumed to have quantisation uncertainty representing a predictive tolerance of $\mathcal{U}_5(0.2)$. Using a (long time scale) RBF model, constructed as described in section 4.3, ι -shadowing is used to define analysis values with $\tau_a = 4$ where such a trajectory exists. Exact values are not available, but are approximated by the value (at $t = 0$) of a shadowing trajectory which shadows for at least $t \in [-4, 4]$. The local manifold explores the five dimensions of the model space for almost all of the 369 cases considered. Measures were calculated using the following parameters, unless otherwise stated: $CV_{\max} = 3$, $\epsilon_{CV} = 0.05$, $\tau_{LV_{\text{opt}}} = \tau_{SV_{\text{opt}}} = 4$, $\tau_{BV_{\text{init}}} = 10$, $\tau_{BV_c} = 1$. Measures were also calculated using six alternative parameter values, namely (1) $\epsilon_{CV} = 0.001$; (2) $\epsilon_{CV} = 0.1$; (3) $\tau_{SV_{\text{opt}}} = 2$; (4) $\tau_{BV_{\text{init}}} = 5$; (5) $\tau_{BV_{\text{init}}} = 20$; and (6) $CV_{\max} = 1$, in order to examine the parameter value's effect, or lack thereof, on the measures.

5.2.3.1 Similarity between CV bases

Varying the optimisation time between two and four has no effect on the similarity between the 3d SV and SSV bases in 5d model space, which is 0.8 for both cases. When a single singular vector is considered the similarity is 0.5; although this is relatively high given two vectors in a 5d space, the validity of the assumption that

the SV approximate the SSV is questionable. Table 5.4 shows the similarity indices between the subspaces of the BV constructed using different methods and between BV and LV subspaces; varying the pre-breeding time has almost no effect on the similarities, hence only the results for $\tau_{\text{BV,init}} = 10$ are given. The similarity between

$\epsilon_{\text{CV}}, \text{CV}_{\text{max}}$	BVA·BVB	BVA·LV	BVB·LV
0.05, 3	0.74	0.45	0.43
0.05, 1	0.75	0.28	0.29
0.001, 3	1.00	0.45	0.45
0.1, 3	0.57	0.44	0.42

Table 5.4: Similarity indices between BV and LV subspaces for the annulus.

the subspaces arising from the two different methods of construction varies with perturbation magnitude, decreasing from being almost always identical for very small perturbation magnitudes, to having a similarity of 0.57 for perturbations of magnitude 0.1; certainly for magnitudes of 0.05 and larger the linearity assumption is not valid. The number of BV used appears to have little effect on the similarity between the alternative subspaces. In contrast, similarity between the BV and LV subspaces varies with the number of CV used, but is affected little by the magnitude of perturbations employed; while the similarity is not insignificant, it does not suggest a close alliance.

An index of 0.01 denotes a definite lack of similarity between SV and both BV subspaces when only one CV is used. Given that two 3d subspaces in a 5d space must have a similarity of at least 0.2, the similarity index of 0.34 indicates that the SV and BV subspaces remain relatively distinct even when three CV are used.

5.2.3.2 Spread

As for the RBF model of the annulus used to generate unconstrained ensembles in section 5.1.1 above, many initial conditions evolve out of the observed interval, $T \in [22^\circ\text{C}, 25^\circ\text{C}]$, under the model used here. Comparing the CV spread to the real spread when either the perturbed CV initial condition or the initial analysis evolve out of the observed interval makes little sense; the results below are for the (101) cases when this does not occur. Figure 5.16 a shows the results for the different CV spread to real spread ratios. All peak at values between one and two, with considerable distribution for values greater than two; the CV perturbations usually overestimate the real spread. Even when the perturbation magnitude is very small, the SV spread overestimates the real spread, as illustrated in figure 5.16 (b). In all these cases, the verification is almost always encompassed by the CV perturbations; in the few cases

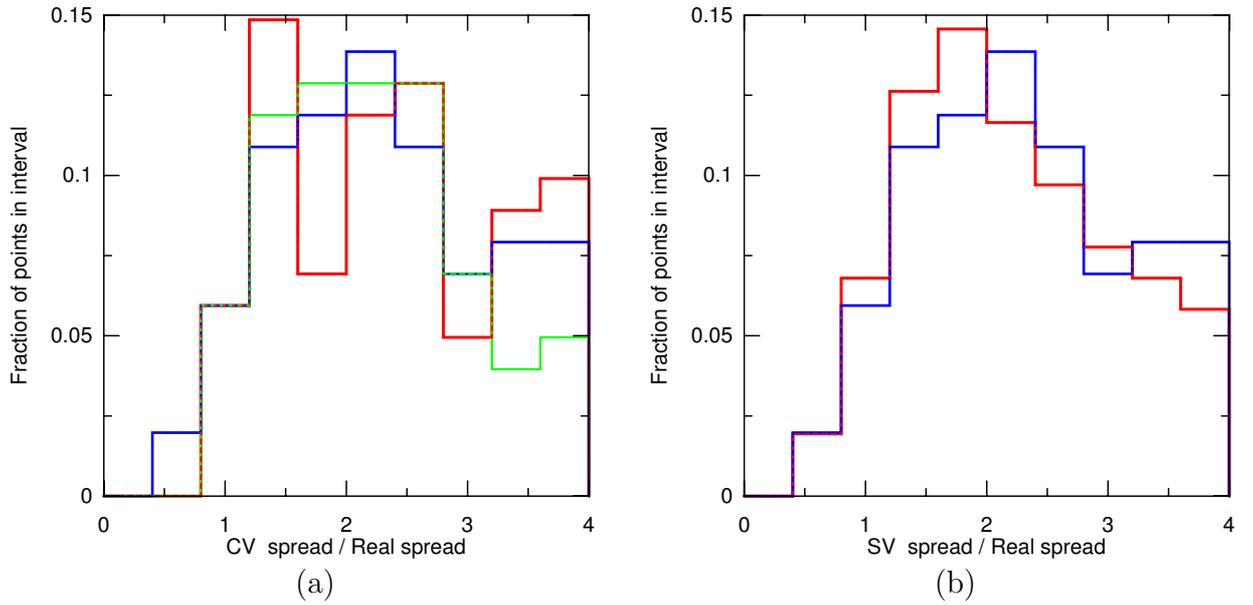


Figure 5.16: A histogram of the ratios of constrained vector spread to real spread (a) for SV (blue), BVA (red) and LV (green), all with $\epsilon_{CV} = 0.05$, and (b) for SV using a perturbation magnitude of 0.001 (red) and 0.05 (blue); results are calculated using a RBF model of the annulus.

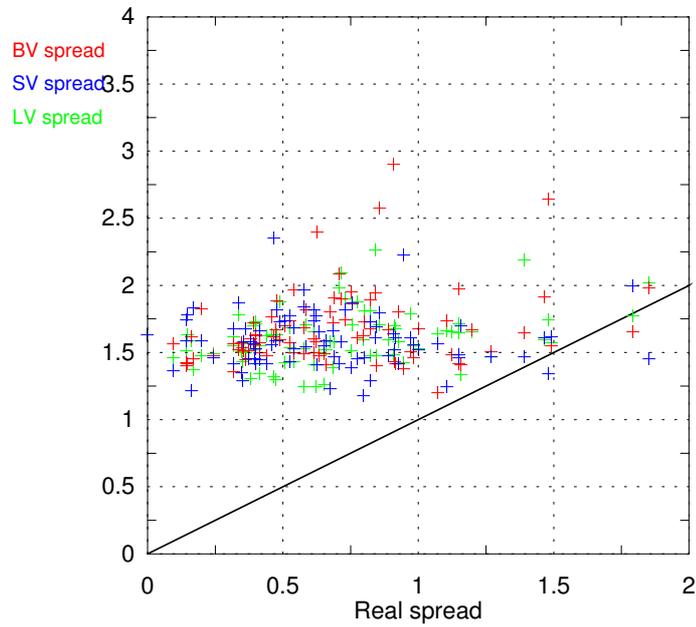


Figure 5.17: A scatterplot of the ratios of CV spread to real spread for SV (blue), BVA (red) and LV (green), for 250 initial conditions, calculated using a RBF model of the annulus. The line $y = x$ (black solid line) is shown for reference.

where the verification is not encompassed, no correlation with spread ratio is evident, as illustrated in figure 5.17.

5.2.3.3 Projection of dream perturbations onto CV subspaces

Scatterplots of the projection of the dream perturbation into the CV basis, comparing all pairs of (SV, LV, BVA), are given when three and one CV are used, see figures 5.18 and 5.19 respectively. For clarity, only 250 results are used in the scatterplots;

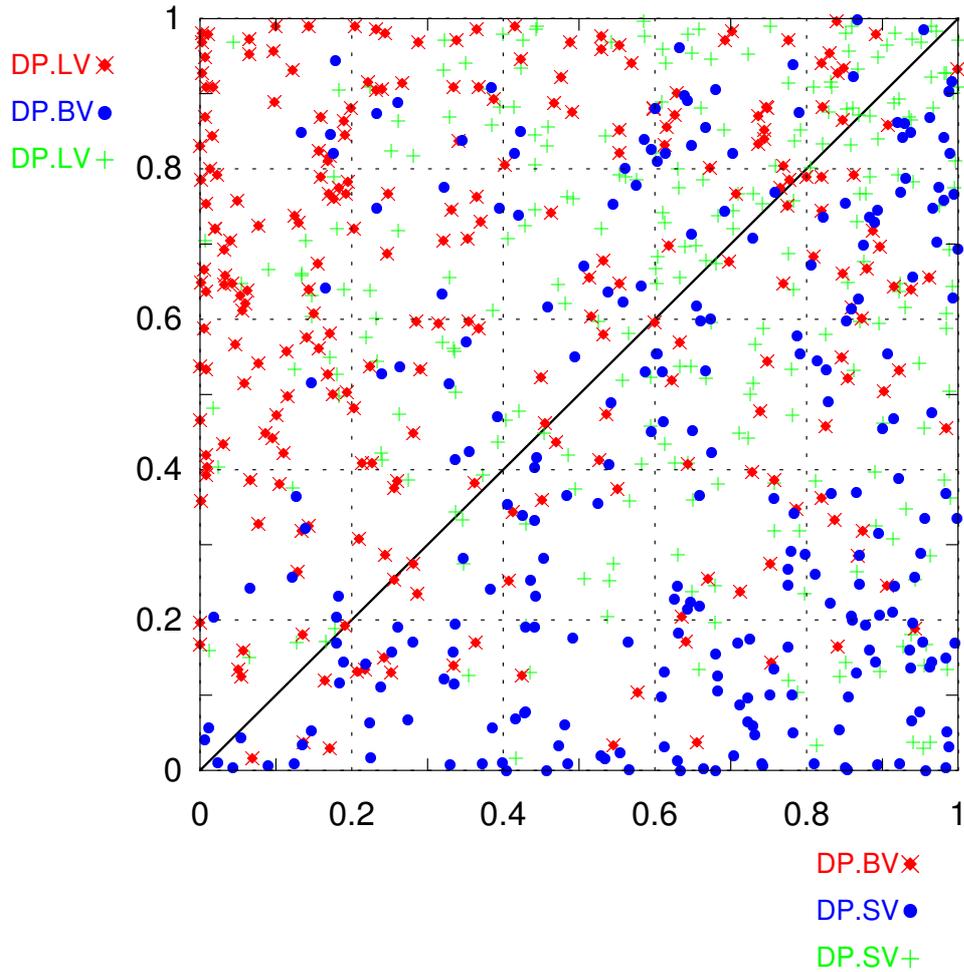


Figure 5.18: A scatterplot of the projection index of dream perturbations onto 3d CV subspaces for all pairs of (SV, LV, BVA), for 250 initial conditions, calculated using a RBF model of the annulus: LV vs. BVA (red diamond cross), BVA vs. SV (blue bullet) and LV vs. SV (green plus sign). The line $y = x$ (black solid line) is shown for reference.

average values of the projection indices over all initial conditions are given in table 5.5.

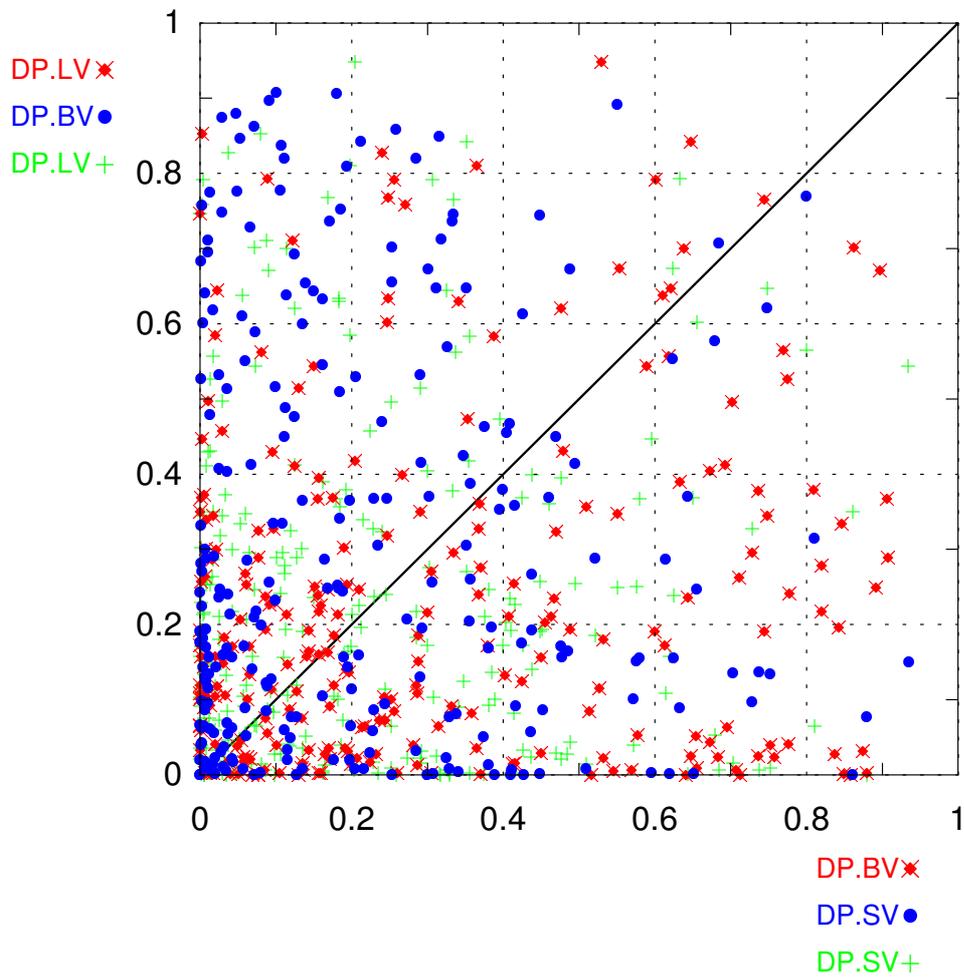


Figure 5.19: A scatterplot of the projection index of dream perturbation onto 1d CV subspace for all pairs of (SV, LV, BVA), for 250 initial conditions, calculated using a RBF model of the annulus: LV vs. BVA (red diamond cross), BVA vs. SV (blue bullet) and LV vs. SV (green plus sign). The line $y = x$ (black solid line) is shown for reference.

CV_{\max}	DP·SV	DP·LV	DP·BVA	DP·BVA
3	0.62	0.62	0.34	0.29
1	0.21	0.23	0.28	0.27

Table 5.5: Projection indices of dream perturbations onto CV subspaces for an RBF model of the annulus.

When only one CV is used then the dream perturbation projects better onto the BV subspace than that of the SV, especially if only the significant projections (those with indices greater than 0.8) are considered. (Counting the blue bursts above the 0.8 horizontal and to the right of the 0.8 vertical in figure 5.19, there are 15 cases for which $DP \cdot BV > 0.8$ as compared to only four cases with $DP \cdot SV > 0.8$.) These results are in agreement with preliminary results obtained using longer times to define both the analysis and shadowing ‘truth’, see figure 10 in *Smith & Gilmour* [84]. Increasing the number of BV has little effect on the projection of dream perturbation onto the BV subspace; the additional directions which the BV explore do not capture directions of the dream perturbation not explored by the first BV direction. In contrast, trebling the number of SV (or LV) used almost trebles the projection of the dream perturbation onto the CV subspace; the projection indices for SV are now superior to those for the BV. The number of cases for which significant projection (> 0.8) occurs for BV has increased from 15 to 37, while for SV an increase from four to 83; for about one third of cases studied the projection of the dream perturbation onto the SV subspace is significant (> 0.8) when 3 CV are used. Varying perturbation magnitude has little effect. The comparative performance of BV and SV, in terms of dream perturbation projection onto CV subspace, is unequivocally reversed in changing from single CV usage to the use of the multiple CV.

While the averages of the projection indices suggest that projection of the dream perturbation onto BVA is slightly better than that onto BVB, more information may be obtained by forming a scatterplot comparing the projections onto the two subspaces, see figure 5.20. For $\epsilon_{BV} = 0.001$ the subspaces are very similar (see section 5.2.3.1); it follows that the dream perturbation projection indices almost all lie on the diagonal, as seen in the scatterplot. As the perturbation magnitude increases, the subspaces begin to differ: while many points still lie on the diagonal, the majority which do not lie in the lower right half plane, denoting that projection onto the method A BV subspace is better than that onto the method B BV subspace. In those cases for which projection of dream perturbation onto BVB is greater than 0.8, projection is also significant onto the BVA subspace; significant projection onto the BVA subspace has no similar correlation with projection onto the BVB subspace.

5.2.3.4 Projection of CV onto local manifold

The projections of CV onto the leading local manifold direction are summarised in figure 5.21 for the standard parameter values. Similar to the dream perturbation

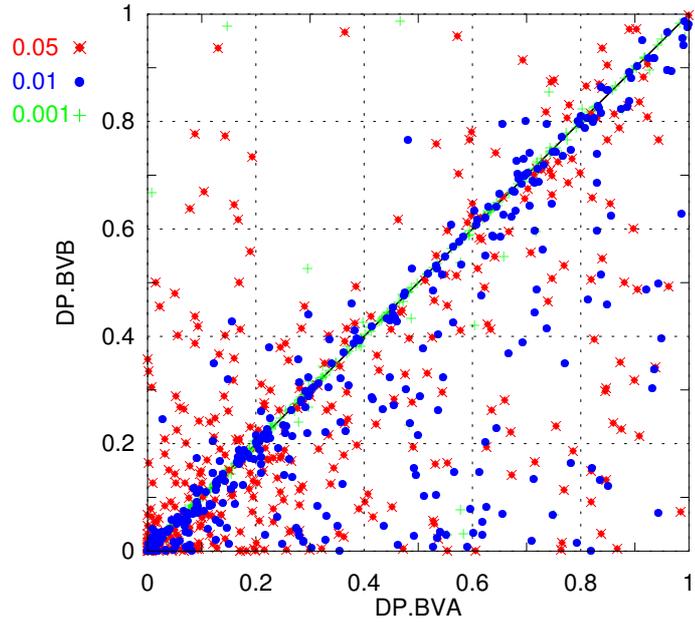


Figure 5.20: A scatterplot of the projection index of dream perturbation onto 3d BVB vs. BVA subspaces, for 250 initial conditions, calculated using a RBF model of the annulus: perturbation magnitude 0.05 (red diamond cross), 0.01 (blue bullet) and 0.001 (green plus sign). The line $y = x$ (black solid line) is shown for reference.

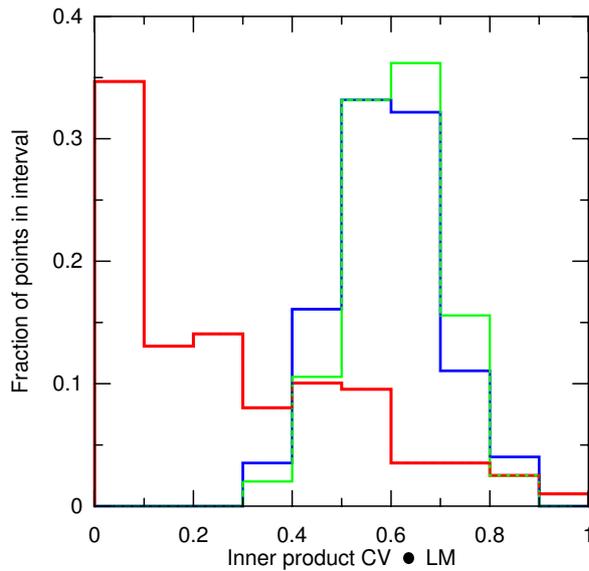


Figure 5.21: A histogram of the projection index of constrained vector subspace onto leading local manifold directions for SV (blue), BVA (red) and LV (green), calculated using a RBF model of the annulus.

projection results, increasing CV_{\max} from one to three almost trebles the SV·LM and LV·LM projections, while the BV·LM projection is comparable to that for $CV_{\max} = 1$.

5.2.3.5 Summary of CV performance

Evaluation of CV performance from the Marzec Spiegel results is simplified by the absence of model error. Results obtained using a RBF model of the annulus do contain the effects of model error, though the extent to which this influences the results is unknown. SV and SSV subspaces are quite similar, but are seldom identical; the assumption that SV approximate SSV is not valid. All CV perturbations usually capture the real spread, and almost all encompass the verification.

Dream perturbations do not project very well onto any of the CV subspaces when only one CV is used, but projection indices for both BV subspaces are higher than those for SV. If only cases with significant projection, taken to be those with indices > 0.8 , are considered then BV again perform better than SV, although none have frequent occurrence of significant projection. The results when three CV are used are very different. The projection indices for BV change little, while those for SV (and LV) almost treble; SV perform better than BV, both on average and when only significant projection is considered. Approximately one third of the cases considered have significant projection of the dream perturbation onto the SV subspace.

While the two methods of BV construction yield similar subspaces for very small perturbation magnitudes, differences appear as the magnitude is increased, indicating the invalidation of the linearity assumption. The similarity between BV and LV is not very high regardless of perturbation magnitude. The results for the different methods are very similar, with method A usually giving marginally better results; the only distinct advantage of method A is when occurrence of significant dream perturbation projection onto the BV subspaces is considered.

5.3 Summary

In Chapter 5 we have introduced probabilistic prediction and discussed both ensemble construction and performance evaluation. At the end of Chapter 4 it was mentioned that, even with the use of assimilation techniques, the initial condition remains uncertain. In this chapter we have considered how the effects of that uncertainty on predictions of the future system state may best be described using ensembles of initial conditions. Unconstrained and perfect ensembles were introduced and compared,

and the effects of model error on their construction considered. Computational limitations for high-dimensional models then motivated discussion of constrained vector (CV) ensembles and the initialisation procedures were noted. The questions of (1) how to evaluate the performance of CV ensembles and (2) how to test whether the assumptions made in their definition were valid, were addressed and the resulting evaluation measures applied to CV ensembles constructed using both a perfect model of the Marzec-Spiegel system and a RBF model of the annulus.

Comparison of unconstrained and perfect ensembles under a perfect model of the Marzec-Spiegel system enabled the varying (and unknown) error in unconstrained ensemble predictions to be contrasted with the accountable error of the perfect ensemble predictions (see figure 5.2). Further, a bifurcation seen in the unconstrained ensemble was shown to be an artifact of including members which do not initially lie on the invariant manifold of the system, rather than a real characteristic of the system. An unconstrained ensemble for the annulus demonstrated the further complications of model error which not only causes the ensemble to evolve differently from the system, but also gives rise to the inclusion of yet more physically unrealisable initial conditions as ensemble members, thereby obfuscating p.d.f. interpretation. The computationally prohibitive size of the unconstrained ensembles in higher dimensions (if they are to provide any useful information), combined with the large return times of many high-dimensional systems (which prohibits the approximation of analogues by observation), motivated consideration of constrained vector subspaces.

The singular vector (SV) and breeding vector (BV) subspaces used in numerical weather prediction were defined and the (numerous) assumptions made in their definition discussed. In particular, the assumption of both SV and (operational) BV ensembles that the linear approximation is relevant for the magnitude of the perturbations (over the optimisation time and the cycle time respectively) was noted. Dream perturbations were defined as the perturbation from the analysis to the best model representation of the system state. Twin pair perturbations, as used in operational ensembles, were introduced.

The question of how to evaluate ensemble performance was posed, and divided into two issues: how successful are the ensembles at achieving the aims for which they were defined, and are the assumptions made in defining the ensembles valid? It was acknowledged that the extent to which these questions may be addressed for observed high-dimensional physical systems is limited; assessment of the relevance of the linear approximation in NWP is the subject of Chapter 6. A variety of measures,

which may be evaluated when the system state is known or can be approximated, were introduced.

Similarity and projection indices to compare different subspaces were defined, enabling the assumption that the SV of the analysis approximate those of the ‘truth’ to be tested. A method of evaluating the spread of CV ensembles was introduced so that it could be compared with the magnitude of the dream perturbation, taken as the ‘real’ spread, so as to evaluate whether the SV ensembles fulfilled their aim of capturing the ‘worse case scenario’. The ability of BV to fulfill their aim of capturing the directions in which the analysis error may grow, and thus include in the ensemble the initial condition closest to the ‘truth’, was examined in two ways. Firstly, the similarity between the BV and Lyapunov vector subspaces was to be measured since, if the model is perfect and the perturbation infinitesimal, the BV evolve towards the first (finite time) Lyapunov vector. Secondly, projection of the dream perturbation onto the BV subspace should be considered, so as to evaluate how well the BV capture the analysis error.

All of the above measures were then evaluated for a perfect model of the Marzec-Spiegel system and a RBF model of the annulus; detailed summaries of the results for each are given in sections 5.2.2.5 and 5.2.3.5 respectively. Comparing the results for the perfect model of the Marzec-Spiegel system with those for the imperfect RBF model of the annulus, SV fulfill their aim in so far as they capture spread (provided the perturbation magnitude is comparable to that of the analysis error) in both cases. While it is valid to assume that the (model) SV approximate the system SV (SSV) for the Marzec-Spiegel system, the assumption is not valid for the annulus. BV do not capture the dream perturbation well in either system, especially when compared to the dream perturbation projection onto multiple SV and LV subspaces. BV subspaces show some similarity to those of the LV, although they are not as similar as SV subspaces are to SSV subspaces. SV and BV subspaces tend to show little similarity. The different methods of BV construction yield similar subspaces for small perturbations, but as might be expected, they differ as the perturbation magnitude increases, indicating the breakdown of the linearity assumption. A fuller and more direct evaluation of the validity of the linearity assumption for both BV and SV ensembles is presented in the following chapter.

Chapter 6

Testing the internal consistency of Ensemble Prediction Schemes

Here we introduce a new statistic, calculated from the nonlinear evolution of twin pair perturbations, and constructed in order to quantify the error made in assuming that linearity holds. Such a measure allows a verification, or otherwise, of the internal consistency of SV and BV ensembles with respect to the linearity assumption made in defining the subspaces; importantly, it may be evaluated for operational Numerical Weather Prediction (NWP) Ensemble Prediction Schemes (EPS) since they employ twin perturbations. In contrast to the more common methods of assessing forecast ability, the internal consistency of ensembles does not appear to have been (systematically) researched. Some results are however contained within literature motivated by different aims; a brief overview is presented in section 6.1 below. In order to gain insight into how nonlinear processes affect perturbation evolution, two other measures are introduced: q -pling times which describe the time taken for a perturbation to amplify by a factor of q ; and nonlinear amplification rates which, for SV perturbations, may be compared with their linear equivalents given by the singular values. Linearity results are presented for CV ensembles evolved under a perfect model of the Marzec-Spiegel system, and under a RBF model of the annulus. The relatively inexpensive cost of running these low-dimensional models enables several sets of parameters defining the CV perturbations (perturbation magnitude, optimisation time, pre-breeding time, ...) to be assessed. The results obtained motivate the calculation of statistics for operational NWP ensembles. Before presenting results from NWP models, a brief description of the practical details of CV ensemble formation and implementation specific to these models is given.

Figures 6.14 and 6.15 summarise the results of evaluating ECMWF and NCEP ensembles and are the main finding of this chapter: operational perturbations are

commonly held to evolve approximately linearly for at least two days; the results presented here show this *not* to be true. Further, the implication of the results is that neither ECMWF nor NCEP operational ensembles are consistent with assumptions made in their definition. A modification to the SV ensembles is suggested by analysing the results for the Marzec-Spiegel and annulus systems; collaboration with ECMWF enabled such a modification to be implemented using the operational NWP model and its effects investigated; the results are presented in section 6.7.

Determining the distribution of times for which the linear approximation is relevant for current atmospheric models is one of a number of questions to be addressed on a list compiled by those gathered at the Newton Institute in Cambridge to discuss atmospheric predictability in 1996; see the ‘Summary of Addressable Questions’ in the overview of the meeting [81]. The results obtained using twin pair perturbations in operational models presented here provide an answer to this question. Consideration of linearity results for individual SV perturbation pairs, of q -pling times and of nonlinear amplification rates enables discussion of two other questions listed, namely:

- Does the linear approximation break down in a systematic manner? In particular, is the image of the direction corresponding to the largest singular value (for a relevant optimisation time) the first in which the linear approximation fails?
- Many linear-dynamics-based intuitions are violated even in low-dimensional nonlinear systems, like the Lorenz 1963 model; yet these NWP models appear to behave consistently with these intuitions. Is there some principle which indicated that there are pathologies which happen only in low order systems. Do these occur “Even In, Or Only In” (EIOOI) low order systems?

The latter is an iteration of the question of whether results from low-dimensional systems ‘really apply to the atmosphere’, in the words of *Lorenz* [45]; all are discussed in light of the results presented.

6.1 Review of literature on the validity of linear approximations in numerical weather models

Computational limitations have fuelled interest in the tangent linear models (TLMs) of numerical weather models by reason of their potential application in two domains: sensitivity analysis of the model to varying parameters, and the assimilation of observations to yield analysis values [63] (as discussed in section 4.4.2). The validity

of the TLM obviously depends on how nonlinear the evolution is, for the magnitude of perturbation considered, for the given model. It should be noted however, that there are two other sources of error, as discussed by *Errico et al.* [23]: the exclusion of some model terms in the linearisation process of the model due to their non-differentiability; and the approximation of the true system basic state about which the linear approximation is formed, which depends on the frequency with which the basic state is updated. If the evolution of a perturbation under the full nonlinear model is compared with that under the TLM, then all sources of inaccuracy will be incorporated. Other investigations utilise the nonlinear evolution of twin (equal and opposite) pairs of perturbations to determine when nonlinear processes may not be neglected (for perturbations of given magnitude and orientation), thereby providing an upper bound for TLM validity. While these investigations provide an estimate of the time for which the linear approximation is valid, most use random perturbations and simplified models; no results could be found for SV or BV perturbations in operational NWP models.

Lacarra & Talagrand [43] use a barotropic (f -plane shallow water) model to investigate for how long the TLM is a good description of the temporal evolution of errors, with the aim of employing the TLM in data assimilation. They conclude that, for errors with amplitude comparable to that of forecast errors found in data assimilation, the TLM is a good approximation for ‘ranges of up to about 48 hours’. In discussing the applicability of their results to atmospheric dynamics the authors emphasise that baroclinic instabilities have a stronger influence than barotropic instabilities .

Vukićević [99] aims to validate the use of the TLM in sensitivity analysis. By evolving perturbations, given by initial data errors, under a primitive equation limited-area model (which includes both baroclinic and barotropic components) and its linear counterpart, the TLM is found to give a good approximation for 1-1.5 days.

Poor performance of (1992) operational assimilation methods in the presence of strong baroclinic instabilities led *Rabier & Courtier* [67] to investigate the validity of the TLM in such cases. Evolving a perturbation, whose initial amplitude is ‘far from being negligible’, both nonlinearly using a low resolution (T21L19) primitive equation model, and under its associated TLM, the authors show that the evolution of the eddy part is essentially linear for a range of 1-2 days, while the total evolution of zonal and eddy parts is less linear.

Errico et al. [23], using the same mesoscale model as *Vukićević* [99], studies evolution to 72 hours of random perturbations with magnitudes comparable to analysis errors. The correlation between nonlinearly and linearly evolved perturbations is

calculated to quantify the accuracy of the TLM in estimating the direction of perturbation evolution. A summer case and a winter case were considered; for both the correlations remain high out to 72 hours. The authors note that the boundary conditions imposed artificially constrain perturbation growth; the rate of error ‘removal’ by the boundary conditions is comparable to the rate of error growth in 24 hours for the winter case considered.

Houtekamer & Derome [37] do not investigate the validity of the linear approximation directly, but use a quasi-nondivergent global spectral model of low resolution (T21) to examine whether the mean of evolved twin pair perturbations yields a higher quality forecast than the control (run at the same resolution). Obviously the BV perturbations must evolve nonlinearly for the mean of the evolved twin pair perturbations and the evolved control to differ; differences are first noticeable between 24h to 48h when perturbations have magnitudes comparable to analysis errors.

Buizza [12] investigates the time after which nonlinear processes can not be neglected by considering the nonlinear evolution of twin pair perturbations and the corresponding control using an ECMWF primitive-equation model (with T63 resolution); the perturbations used are given by the SV, optimised for 36 hours, of a lower resolution (T21) model. From the three (twin pair and control) nonlinear evolutions, Buizza quantified the contribution of nonlinear processes in two ways. The first way considers a truncated Taylor expansion to describe the perturbation evolution and uses algebraic manipulation to compare the amplitude of nonlinear terms with those of the linear terms; this method differs from that of *Lacarra & Talagrand* [43] and of *Rabier & Courtier* [67] since it uses nonlinear evolution of twin pair perturbations rather than comparing nonlinear and linear evolution of a single perturbation. The second way calculates the correlation between the pair of evolved perturbations, as suggested by *Houtekamer & Derome* [37] in their study of BV perturbation evolution; this method is discussed in further detail below, in section 6.2. Again, this differs from the method of *Errico et al.* [23] in that it uses the nonlinear evolution of twin pair perturbations rather than comparing nonlinear and linear evolution of a single perturbation. When the initial perturbation amplitudes are comparable to analysis errors both methods suggest that nonlinear processes become important after 2-2.5 days.

Park & Droegemeier [63] investigate the validity of the TLM in a moist convective cloud model by comparing the magnitude of nonlinearly and linearly evolved perturbations. The results were found to be very sensitive to perturbation magnitude, and

also to the frequency with which the basic state, used in calculating the TLM, was updated.

The times for which the linear approximation or TLM may be considered to be a valid approximation to the nonlinear dynamics varies between investigations from 24 hours to at least 3 days. This variation is unsurprising since the magnitude and orientation of perturbations differ, the model considered is not unique, and the definition of validity varies between studies. *Palmer et al.* [60] state that linearity holds for operational perturbations and models for approximately 3 days; no reference or data is however given. When (anti-)correlation between twin pair perturbations is used to measure the validity of the linear assumption, a value of 0.7 is taken as a threshold for linearity; interpretation of this threshold is discussed below and results presented for comparison, for the systems considered here.

6.2 Quantifying the relevance of the linear assumption

Having established that we wish to better quantify the error made in assuming linearity, we now derive an alternative measure based only on the evolution of a control and a twin pair of perturbations, *i.e.* data available operationally. Approximating the extent for which the linear approximation is relevant obviates the need to use a tangent linear model, with its implicit approximations when implemented numerically using observational data. In addition, the internal consistency of BV ensembles depends upon the relevance of the linear approximation rather than that of the TLM. SV ensembles assume that the TLM is relevant for the optimisation time; relevance of the linear approximation is a necessary condition for relevance of the TLM, and will therefore provide an upper limit of the extent for which SV ensembles may be internally consistent. *Buizza* [12], after *Houtekamer & Derome* [37], quantifies the point at which nonlinearities dominate the system as that at which the (anti-)correlation between the evolved twin pair perturbations crosses the value of 0.7. If the control trajectory, initiated from $\mathbf{A}_0(0)$, is taken as the fiducial trajectory, and the deviations of a positive (negative) perturbation from the control in evolution are denoted by $\boldsymbol{\delta}^+(t)$ ($\boldsymbol{\delta}^-(t)$), then the *anti-correlation*, ℓ , is given by

$$\ell = -\frac{(\boldsymbol{\delta}^+(t), \boldsymbol{\delta}^-(t))}{\|\boldsymbol{\delta}^+(t)\| \cdot \|\boldsymbol{\delta}^-(t)\|}. \quad (6.1)$$

where the inner product, (\cdot, \cdot) , defines one of several possible metrics, $\|\cdot\|$. Note that the correlation measure reflects *only* the orientation of the evolved perturbations

and is blind to their relative magnitudes; if one perturbation grows while the other shrinks, with no change in orientation, the correlation will remain 1.

Ideally we wish to define a measure which both quantifies the error made in assuming linearity and is sensitive to both magnitude and orientation. While growth is linear $\delta^+(t) + \delta^-(t) = \mathbf{0}$, thus a measure of the error in assuming linearity is the sum of the evolved perturbations; it is sensitive to both the relative magnitudes and orientations of the evolved perturbations, as illustrated in figure 6.1. Scaling by the

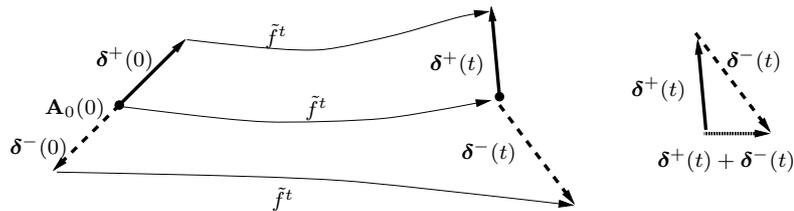


Figure 6.1: Defining Θ : Equal and opposite perturbations at $t = 0$, $\delta^\pm(0)$, evolve so as to be no longer symmetric at time t : the error in assuming linear dynamics, $\|\delta^+(t) + \delta^-(t)\|$, is scaled by the average magnitude of the evolved perturbations to give the test statistic Θ .

average magnitude of the evolved perturbations defines a suitable measure, Θ , given by

$$\Theta(\hat{\delta}, \|\delta\|, t) = \frac{\|\delta^+(t) + \delta^-(t)\|}{0.5\{\|\delta^+(t)\| + \|\delta^-(t)\|\}} \quad (6.2)$$

(where $\|\cdot\|$ is an appropriate norm); Θ will, of course, vary with the $\mathbf{A}_0(0)$ of initialisation. As long as the perturbations are growing linearly, $\delta^+ = -\delta^-$ and $\Theta = 0$. Thus $\Theta = 0$ implies that the linearised dynamics *may*¹ be exact, while $\Theta = 0.5$ implies that the error made in assuming linearity will be at least 50% of the average magnitude of the evolved perturbations. In comparison to the correlation measure, note that $\ell = 0.7$ corresponds to the perturbations deviating by $\sim 45^\circ$ from anti-parallel; the corresponding error in assuming linearity is at least $\sim 75\%$ of the mean magnitude of the evolved perturbations (*i.e.* $\Theta > 0.75$). There is no inconsistency here since each of the statistics is a necessary but not sufficient condition for linearity. Discrepancies in magnitude between twin pair perturbations may be detected by comparing Θ and correlation curves; for several cases they are explicitly plotted, see section 6.6.

Defining a global characteristic time for which the linear approximation is relevant for the CV ensemble necessitates averaging over a collection of initial conditions, distributed over the attractor, due to the variation of Θ with initial condition. *van*

¹It is crucial to remember that $\Theta = 0$ is only a necessary condition from linearity; one can contrive examples (*e.g.* cubic nonlinearities) where $\Theta = 0$ for some perturbations and yet the linear approximation is wildly inaccurate.

den Dool [96] suggests that the return time of the atmosphere is on the order of 10^{30} years, longer than the life of the universe; return times of the models and the time taken to process NWP data render averaging over the attractor infeasible for NWP models. Results presented here are averages over 25 days (7 days for the comparison of SV optimisation times). Where more than one twin pair is used, the average is also taken over all pairs. In the case of SV perturbations, the twins are ordered according to their corresponding singular values so that the first twin pair is that which is designed to have grown the most at optimisation time (under the linear model); one might expect the first twin pair to be the orientation in which the linear approximation fails soonest. Results comparing the values of Θ for individual SV twin pairs are considered to see if this is in fact the case. Similar behaviour is not expected for BV twin perturbations since their ordering does not depend on any such physical values.

Once the extent of linearity has been determined, it is of interest to ask *how* the linear approximation breaks down: do the finite perturbations grow at a constant exponential rate, rendering the linear approximation gradually less accurate with time, or is there some ‘event’ or region of phase space near which the perturbation grows rapidly causing the linear approximation to fail within a short period of time? *Smith et al.* [85] discuss several possible ways to quantify local growth rates which assume exponential growth rate of perturbations. Finite time Lyapunov exponents provide information about the local growth rate of infinitesimal perturbations, but such information is not relevant for finite perturbations. Instead of estimating the growth rate over a certain time interval, one may consider the time, τ_q , taken for a perturbation to amplify by a given factor of q ; similarly the time taken for a perturbation to amplify by a factor of q^{2n} is denoted $\tau_{q^{2n}}$. If the growth rate is uniformly exponential in time, then τ_{q^2} will be twice τ_q (and $\tau_{q^{2n}} = 2\tau_{q^n}$). If $\tau_{q^{2n}} \neq 2\tau_{q^n}$ then growth is not exponentially uniform.

In the following sections results for $q = 2$ are presented, and give some indication of the uniformity, or otherwise, of growth of CV perturbations in the Marzec-Spiegel and annulus models. Scatterplots provide an idea of the range of values covered and indicate the average values, while histograms of the ratios of $\tau_{2^{2n}}/\tau_{2^n}$ provide information on the density of distribution. These statistics are calculated only for initial perturbations which q -ple within the given time interval, and as such are limited in the information they contain. Data for the numerical weather models is also restricted in this way, and more so since ensemble perturbations are only sampled

every 12 hours; the usefulness of doubling times as a measure of perturbation growth is limited in this case.

Of particular interest is the comparison of SV perturbation growth with their corresponding singular values, for perturbations constructed from SV optimised over both 48 hours and 24 hours. Given the sampling rate of 12 hours, we revert to considering amplification factors over a certain time interval, known exactly. Interpretation of the results is aided by plotting line segments with gradients proportional to the amplification factor or singular value; a full description is given in section 6.7.

6.3 Linearity evaluation of the Marzec-Spiegel system

As discussed above, the relatively inexpensive cost of running a low-dimensional model of the Marzec-Spiegel system enables constrained vectors defined using a variety of parameter sets to be evaluated. Computation of analysis values and constrained vectors is as in section 5.2.2; the majority of analysis values are in error by less than 0.02. The standard parameter values used are: $CV_{\max} = 1$, $\epsilon_{CV} = 0.02$, $\tau_{SV_{\text{opt}}} = 2$, $\tau_{BV_{\text{init}}} = 10$, $\tau_{BV_c} = 1$. Four alternative parameter values, namely (1) $\epsilon_{CV} = 0.1$; (2) $\tau_{SV_{\text{opt}}} = 4$; (3) $\tau_{BV_{\text{init}}} = 5$; and (4) $CV_{\max} = 2$, are used in order to examine the parameter value's effect, or lack thereof, on the results. In each case the results are averaged over a large number of initial conditions, distributed uniformly with respect to the natural measure of the system.

6.3.1 Evaluation of linearity measures

Linearity results for CV perturbations in the Marzec-Spiegel system vary considerably with the magnitude of the initial perturbation, but are affected little by other variations in parameter values. Results for SV perturbations are shown in figure 6.2; relative nonlinearity, described by Θ , increases gradually (for perturbations of magnitude $\epsilon_{SV} = 0.02$), to a maximum of 0.3 at 30 time units, with no discernible difference between perturbations optimised over 2 or 4 time steps. Similar behaviour is seen in the correlation results; the (anti-)correlation remains high for all perturbations of magnitude 0.02. Using either correlation or relative nonlinearity as the measure, the assumption that linearised evolution provides a good approximation to the nonlinear evolution of the SV perturbations over the optimisation time is valid.

When two SV perturbations are used, their average relative nonlinearity is slightly greater than that of the single SV perturbation until about 12 time units, when the single SV relative nonlinearity becomes larger. Given that the single SV is that which

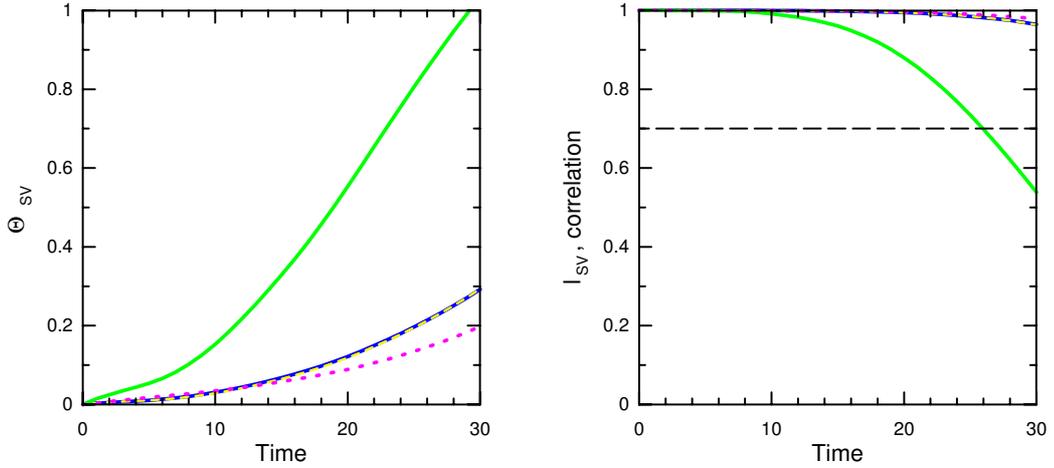


Figure 6.2: Twin experiment linearity results of SV perturbations for a perfect model of the Marzec-Spiegel system. Plots of average Θ (left) and average correlation (right) as a function of time are shown for the standard parameters ($\epsilon_{CV} = 0.02$, blue solid line), and alternatively with $\epsilon_{CV} = 0.1$ (green solid line), $CV_{\max} = 2$ (magenta dotted line), $\tau_{SV_{\text{opt}}} = 4$ (yellow dotted line).

initially has linearly grown the most at the optimisation time, one may expect it to be the direction which is first affected by nonlinearities, in which case the above result is counter-intuitive. Plotting results for the individual SV perturbations explicitly, as in figure 6.3, it becomes evident that, on average, the relative nonlinearity of the second

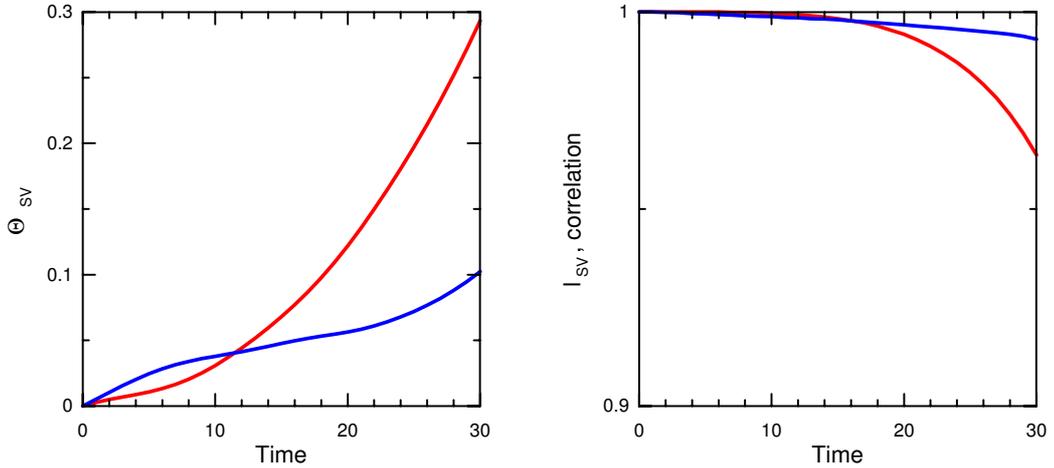


Figure 6.3: Linearity results of SV perturbations for a perfect model of the Marzec-Spiegel system using the standard parameters except that $CV_{\max} = 2$. Panels are as in figure 6.2 with results for the first SV (red) and the second SV (blue).

SV initially grows faster than that of the first. Similar behaviour is reflected in the correlation. On average the nonlinearity of the second SV perturbation is larger at small times, but details vary with initial condition.

Increasing the initial perturbation magnitude to 0.1 increases the relative nonlinearity drastically. As discussed in section 6.2 above, a correlation value of 0.7 is equivalent to relative nonlinearity of ~ 0.75 if the magnitudes of the evolved twin pair are similar, and to a larger relative nonlinearity the more disparate the twin magnitudes. Given that the relative nonlinearity for initial perturbations of magnitude 0.02 is ~ 0.85 when the correlation is 0.7, we conclude that the relative nonlinearity is due to differences in both orientation and magnitude between the twin pairs. Again, whichever measure is used, linearity holds at this optimisation time.

Results for BV perturbations are qualitatively similar, as shown in figure 6.4 : increasing the initial perturbation magnitude increases relative nonlinearity and decreases correlation while other parameter variations have little effect. The method of BV construction also appears to have little effect on linearity results, as does the use of one or two BV perturbations. This is to be expected given that the BV perturba-

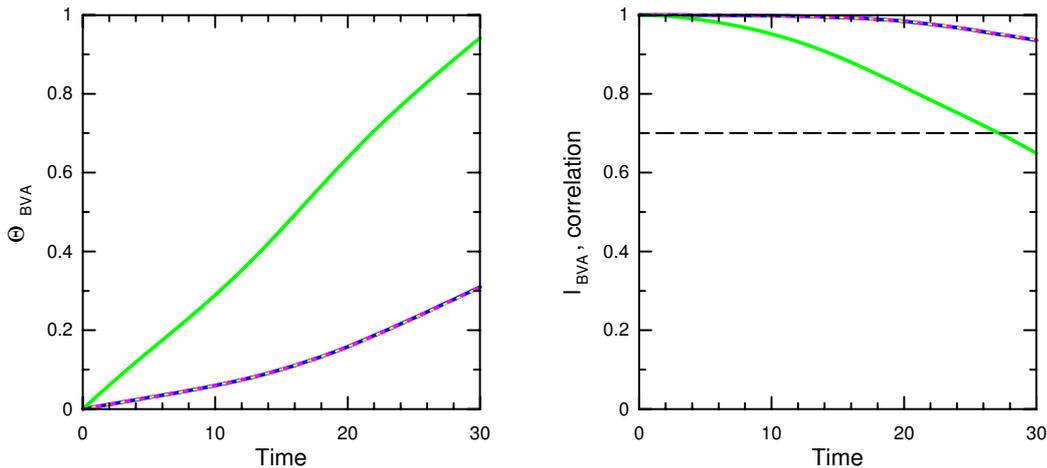


Figure 6.4: Linearity results of BV perturbations, generated using method A, for a perfect model of the Marzec-Spiegel system. Panels are as in figure 6.2 with results for the standard parameters (blue solid line), and alternatively with $\epsilon_{CV} = 0.1$ (green solid line), $CV_{\max} = 2$ (magenta dotted line), $\tau_{BV_{\text{init}}} = 5$ (yellow dotted line).

tions are not ranked in any way, but simply evolved from different initial orientations. As for the SV perturbations, the value of relative nonlinearity when the correlation is 0.7 suggests that nonlinearity is due to disparities in both magnitude and orientation for perturbations with initial magnitude 0.1. For the parameter values considered here, linearised evolution provides a good approximation to the nonlinear evolution over the cycle time for BV perturbations. In comparison with initial magnitude 0.02 SV perturbations, the correlation for BV perturbations decreases more rapidly while

relative nonlinearity results for the SV and BV perturbations are similar; disparate evolved twin magnitudes are more pronounced in SV perturbations.

6.3.2 Doubling times

The time taken for CV perturbations to double and quadruple in magnitude in the Marzec-Spiegel system is relatively long for perturbations of initial magnitude $\epsilon_{CV} = 0.02$ or 0.1 . Within the time span of $[0, 30]$, none of the SV or BV perturbations grew by a factor of 16, and very few ($< 10\%$) of the BV perturbations quadrupled in magnitude. For initial magnitude 0.02 the SV perturbations were evolved out to $t = 100$, in order to capture the τ_{16} vs. τ_4 behaviour; even so only $\sim 35\%$ of the perturbations grew by a factor of 16 within this time. Results are presented for those perturbations which did grow sufficiently, with averages taken over those perturbations. Both sets of results for the SV perturbations (figure 6.5) and the results for the BV perturbations (figure 6.6) are distributed near to the $y = 2x$ line, with the averages lying on the lines. From these results we conclude that for the majority

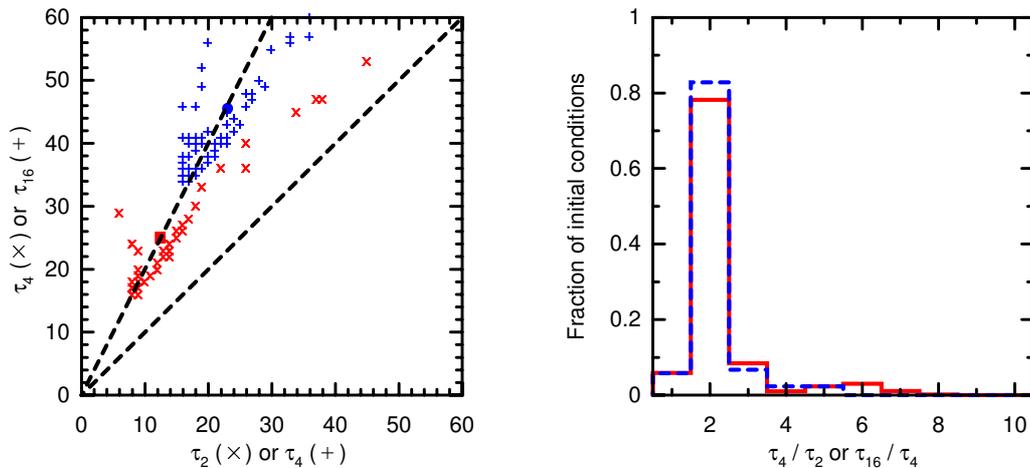


Figure 6.5: q -pling times of SV perturbations for a perfect model of the Marzec-Spiegel system. A histogram of the ratio of quadrupling to doubling times (right) gives information obscured by overplotting in the scatterplot (left) of quadrupling times vs. doubling times (crosses, with average shown by solid square) and of the time to grow by a factor of 16 vs. quadrupling times (blue plus signs, with average shown by solid circle). The lines $y = x$ and $y = 2x$ are also shown (dashed lines).

of perturbations which grow sufficiently within the given time interval, perturbation growth is exponentially uniform (up to a growth factor of 16 for SV perturbations, and a factor of 4 for BV perturbations). This result is supported by the histograms, which show that for the vast majority of initial perturbations $\tau_{2^{2n}} \simeq 2\tau_{2^n}$ ($n = 1, 2$

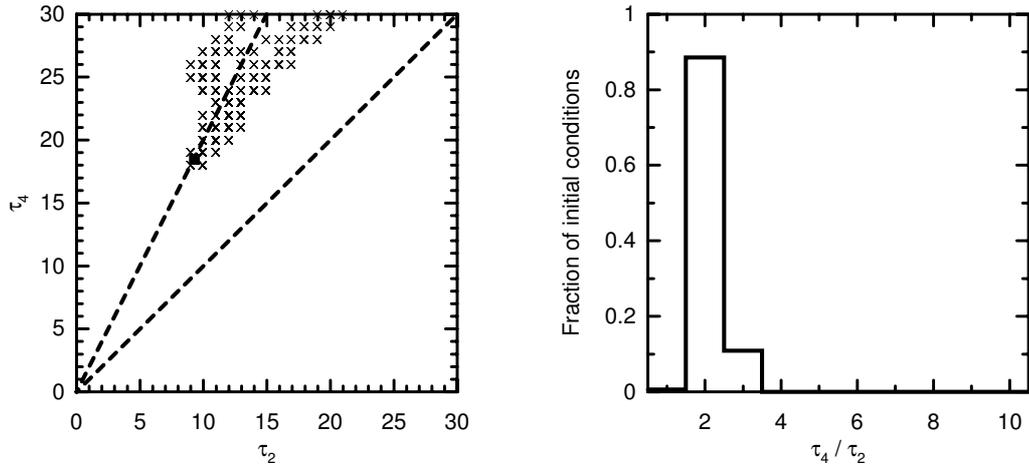


Figure 6.6: q -pling times of BV perturbations for a perfect model of the Marzec-Spiegel system, with panels as in figure 6.5, but showing only doubling and quadrupling times.

for SV perturbations; $n = 1$ for BV perturbations). Results for $\epsilon_{CV} = 0.1$ are very similar, the average time for each growth factor to occur being marginally shorter.

6.3.3 Summary of linearity evaluation measures

In summary, SV and BV ensembles constructed for each of the parameter value sets considered here are internally consistent with the assumption that linearity is a good approximation for the duration of the optimisation time or cycle time respectively. Perturbation growth, to a factor of 16 for SV and to a factor of 4 for BV, appears to be exponentially uniform. When two SV perturbations are employed, that which is initially more nonlinear varies with initial condition, but on average it is the second SV perturbation which is initially more nonlinear.

6.4 Linearity evaluation of the annulus

Having considered a low-dimensional system for which the perfect model is known, we now consider the physical system of the rotating fluid annulus for which a relatively good low-dimensional model may be constructed from observations. Computation of analysis values and constrained vector ensembles is as in section 5.2.3 unless otherwise stated. Results are calculated for the following standard parameter set: $CV_{\max} = 3$, $\epsilon_{CV} = 0.05$, $\tau_{SV_{\text{opt}}} = 4$, $\tau_{BV_{\text{init}}} = 10$, $\tau_{BV_c} = 1$. They were also calculated using six alternative parameter values, namely (1) $\epsilon_{CV} = 0.001$; (2) $\epsilon_{CV} = 0.01$; (3) $\epsilon_{CV} = 0.1$; (4) $\tau_{SV_{\text{opt}}} = 2$; (5) $\tau_{SV_{\text{opt}}} = 8$; (4) $\tau_{BV_{\text{init}}} = 5$; (5) $\tau_{BV_{\text{init}}} = 20$; and (6) $CV_{\max} = 1$, in order to examine the parameter value's effect, or lack thereof, on the results. Results

are averaged over a large number of initial conditions, distributed uniformly with respect to the natural measure of the system.

6.4.1 Evaluation of linearity measures

Similar to results for the Marzec-Spigel system, linearity results vary with initial perturbation magnitude, but are little affected by changes in other parameter values. Figure 6.7 provides a different perspective from those above; it shows the fraction of initial conditions considered for which the SV perturbations have relative nonlinearity less than 0.2 at each time step. Considering $\Theta = 0.2$ as a threshold below which the linearity approximation may be considered relevant, this perspective allows the fraction of initial conditions satisfying the linearity approximation to be determined as a function of time. Results for the positive and negative twin members are plot-

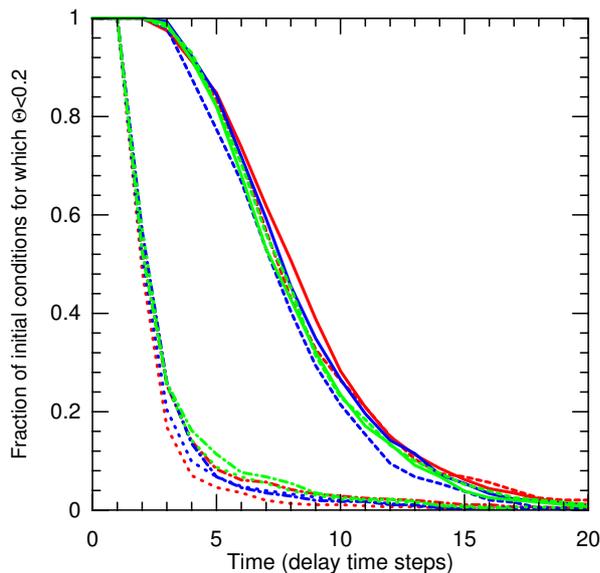


Figure 6.7: Illustration of the fraction of SV perturbations, to distinct initial conditions, which exceed a given linearity threshold as a function of time for a RBF model of the annulus. Results are for averages for positive perturbations, $\epsilon_{SV} = 0.01$ (solid line), $\epsilon_{SV} = 0.1$ (dot-dashed line) and negative perturbations, $\epsilon_{SV} = 0.01$ (dashed line), $\epsilon_{SV} = 0.1$ (dotted line). At each magnitude there are 3 sets of results, obtained by varying the optimisation time, t_{opt} , from 2 (red) to 4 (blue) or 8 (green) time steps.

ted separately, with little difference between them. Similarly, variation between SV optimised over different times and single or multiple SV perturbations is minimal. In contrast, the results are divided into two distinct sets according to whether the perturbations have initial magnitude of 0.01 or 0.1; decreasing the initial perturbation magnitude increases the fraction of initial conditions below the threshold value, thus

increasing the likelihood that the SV ensemble will be internally consistent if other parameter values remain constant². Results for BV perturbations differ little between the methods of construction; Θ values are consistently slightly larger than those for SV perturbations of the same perturbation magnitude.

Considering the smaller initial perturbation magnitude, only $\sim 30\%$ of initial SV perturbations optimised over 8 time steps have a relative nonlinearity less than 0.2 at the optimisation time. The considerable independence of results on optimisation time (and general decrease of the number of initial conditions with relative nonlinearity below the threshold value as time increases) implies that the internal consistency of the SV ensemble may also be increased by keeping the initial perturbation constant and decreasing the optimisation time. Over 90% of initial SV perturbations (of magnitude 0.01) optimised over 4 time steps have $\Theta < 0.2$ at optimisation time, and almost all initial SV perturbations optimised over 2 time steps satisfy this criterion.

Returning to the display of results as the average relative nonlinearity and average correlation as a function of time, figures 6.8 and 6.9 show the same dependence of results on initial (SV and BV, respectively) perturbation magnitude; varying other parameters again makes little difference. The levels of relative nonlinearity and

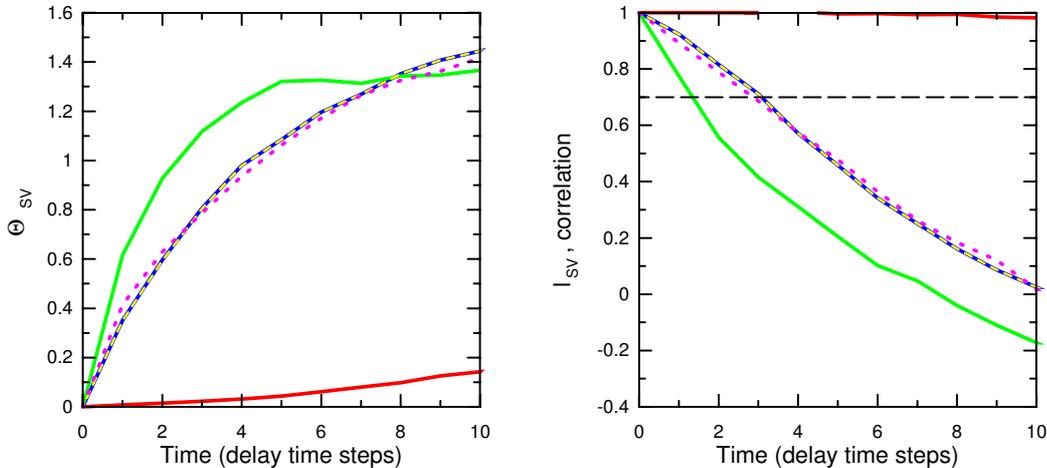


Figure 6.8: Linearity results of SV perturbations for a RBF model of the annulus. Panels are as in figure 6.2 with results for the standard parameters (blue solid line), and alternatively with $\epsilon_{CV} = 0.1$ (green solid line), $\epsilon_{CV} = 0.001$ (red solid line), $CV_{max} = 1$ (magenta dotted line), $\tau_{SV_{opt}} = 2$ (yellow dotted line).

(anti-)correlation show much more variation than for the Marzec-Spiegel system, with some signs of saturation of relative nonlinearity at values above 1.2. SV perturbations

²This assumes that the qualitative nature of the results does not vary with threshold function; results for other threshold values (not shown) support this assumption.

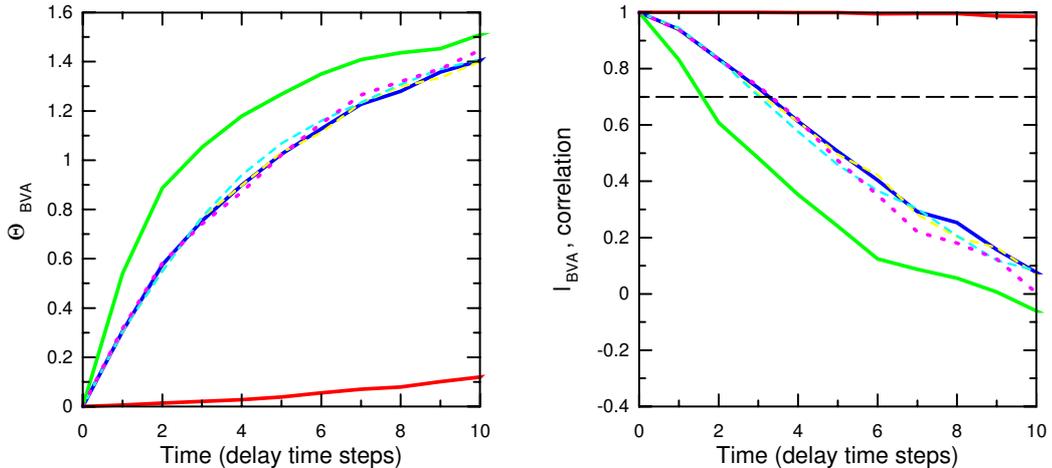


Figure 6.9: Linearity results of BV perturbations, generated using method A, for a RBF model of the annulus. Panels are as in figure 6.2 with results for the standard parameters (blue solid line), and alternatively with $\epsilon_{CV} = 0.1$ (green solid line), $\epsilon_{CV} = 0.001$ (red solid line), $CV_{max} = 1$ (magenta dotted line), $\tau_{BV_{init}} = 5$ (yellow dotted line), $\tau_{BV_{init}} = 20$ (cyan dashed line).

optimised over 4 time steps and with an initial magnitude of 0.05 have an average relative nonlinearity value of ~ 0.9 , *i.e.* the error in assuming linearity at the optimisation time is 90% the average evolved perturbation magnitude; such an SV ensemble *can not* be considered to be internally consistent under the z500 norm used here. The corresponding correlation value is ~ 0.7 , indicating that nonlinearity arises from disparities in both magnitude and orientation. For BV perturbations the average values of Θ and correlation for initial magnitudes upto 0.05 support internal consistency; those for $\epsilon_{BV} = 0.1$ are marginal.

From figure 6.8 we may also see that the single SV perturbation grows more non-linearly for the first three time steps. While individual results show all combinations of (1, 2, 3) in initial nonlinear ranking, the average is 1, 2, 3 for the first three time steps, independent of initial perturbation magnitude, see figure 6.10. No significant differences can be identified in a similar plot for BV perturbations (not shown).

6.4.2 Doubling times

q -pling times for SV and BV perturbations evolved under a RBF model of the annulus show considerable variation, see figures 6.11 and 6.12 respectively for results for perturbations with initial magnitude 0.05. While a significant fraction of the perturbations initiated grow by a factor of 4, only a small fraction grow by a factor of 16; results are shown for those which do (at least 250 initial conditions). While the

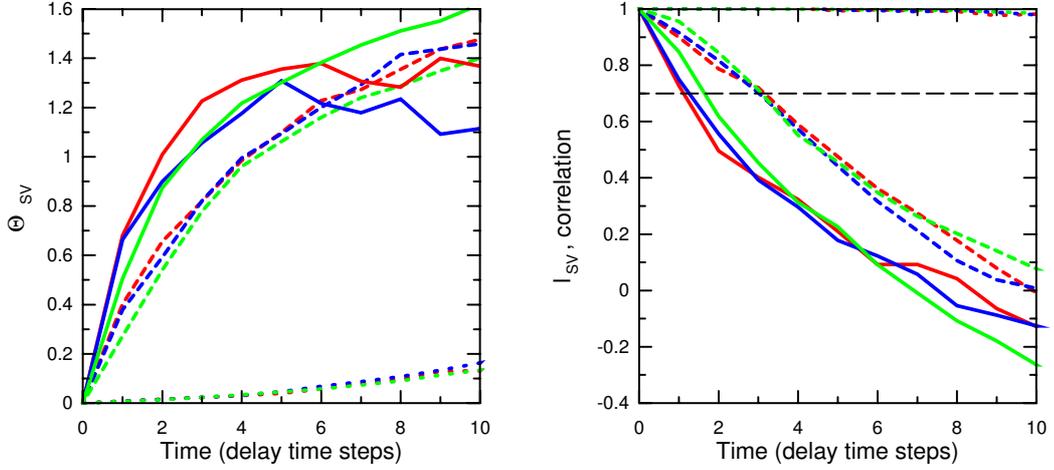


Figure 6.10: Linearity results of SV perturbations for a RBF model of the annulus using the standard parameters ($\epsilon_{SV} = 0.05$, dashed lines), and alternatively with $\epsilon_{SV} = 0.001$ (dotted lines) and $\epsilon_{SV} = 0.1$ (solid lines). Panels are as in figure 6.2 with results for the first SV (red), the second SV (blue), and the third SV (green).

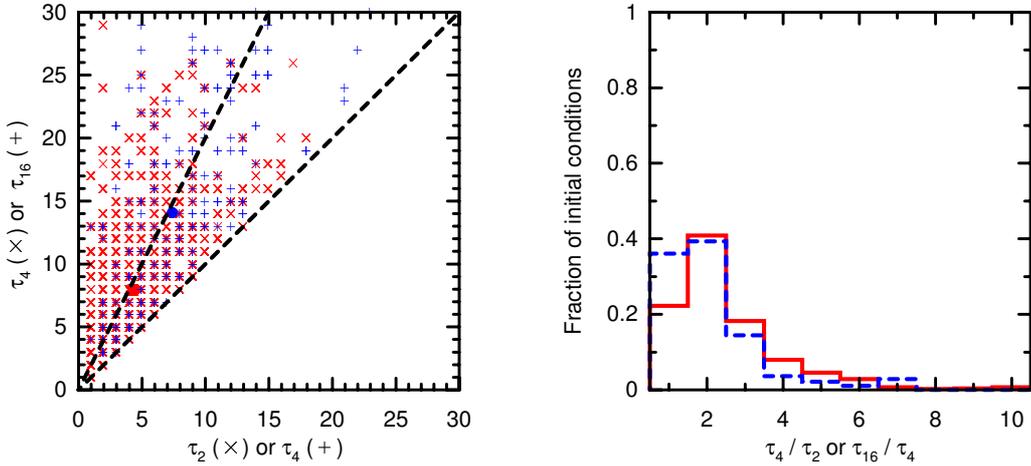


Figure 6.11: q -pling times of SV perturbations for a RBF model of the annulus. A histogram of the ratio of q^2 -pling to q -pling times (right) gives information obscured by overplotting in the scatterplot (left) of quadrupling times vs. doubling times (red crosses, with average shown by solid square) and of the time to grow by a factor of 16 vs. quadrupling times (blue plus signs, with average shown by solid circle). The lines $y = x$ and $y = 2x$ are also shown (dashed lines).

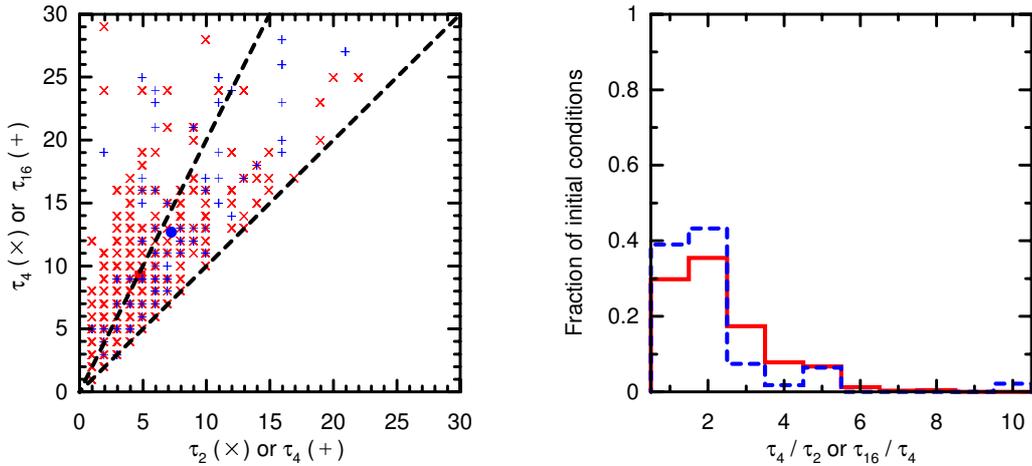


Figure 6.12: q -pling times of BV perturbations for a RBF model of the annulus, with panels as in figure 6.11.

average values all lie near the line $y = 2x$, the range of distribution of the scatterplots and the ratios of $\tau_{2^{2n}}/\tau_{2^n}$ summarised by the histograms show that the growth of CV perturbations is not exponentially uniform out to a factor of four, let alone a factor of 16.

6.4.3 Summary of linearity evaluation measures

For the RBF model of the annulus, BV ensembles with perturbations of initial magnitude 0.05 are internally consistent for a cycle time of 1. SV ensembles with perturbations of the same magnitude and optimised over 4 time steps are not internally consistent; if the optimisation time is reduced to 2 time steps the error in assuming linearity at the optimisation time is significantly reduced, to $\sim 30\%$ of the average evolved perturbation magnitude. The growth of SV and BV perturbations is not exponentially uniform. When multiple SV perturbations are employed, on average the one corresponding to the largest singular value is initially the most nonlinear; on average, but not always.

6.5 Details of ensemble employment in NWP

While the aims of numerical weather prediction were outlined in section 5.1.3 to explain the motivation behind CV ensembles, practical details of their implementation in NWP were omitted there, and are now presented below. The term NWP model is used to describe a variety of models, differing in resolution, used within each operational centre, *e.g.* within the European Centre for Medium-range Weather Forecasting

(ECMWF) or the American National Centre for Environmental Prediction (NCEP). The vertical resolution of each model is specified by the number of layers used in the model; a 28 layer model is denoted L28. Grid points over the Earth are usually distributed on a Gaussian grid and the number of points used is related to the level of spectral triangular truncation (for details see section 2.3.4); if 42 modes are retained the model is described as T42. It is possible to use a linear grid, which uses less grid points than the Gaussian grid while retaining as many modes; a linear T_l159 model has the same number of grid points as a T106 model, but keeps 53 more modes [14]. Switching between the different models is possible by interpolating the data (more details can be found in *James* [40] and references therein); we shall focus on aspects of the state space dynamics which do not require understanding of the interpolation procedures.

At ECMWF, three models are run operationally: (1) T42L31, with which the SV are determined and the ensemble formed as detailed below; (2) $T_l159L31$, used to evolve the ensembles (the analysis is also evolved at this resolution to give a low resolution ‘control’); and (3) T213L31, under which the analysis is also evolved to give a high resolution control forecast. There also exists a tangent linear model (TLM), and its adjoint, for the T42L31 model which includes the tangent version of the adiabatic part of the model, linearised horizontal diffusion and simple vertical diffusion and surface drag; among the processes not included are radiation, convection and gravity wave drag [13, 63]. The 25 twin pair SV perturbations which form the ECMWF SV ensemble are calculated as follows (further details may be found in *Buizza et al.* [14] and *Molteni et al.* [50]):

- Each day at 1200UTC, the TLM is started from the T42L31 resolution value of the analysis; all the singular vectors, optimised over 48 hours are calculated by the Lanczos algorithm [88].
- Two sets of 25 SV, \mathbf{v}_j ; $j = 1, 25$, are then selected, one over the Northern Hemisphere and another over the Southern Hemisphere (both extra Tropics, *i.e.* 30°N/S-90°N/S). The first 4 SV for each set are always selected. Each subsequent SV is selected if at least half of its total energy is outside the localised regions of the SV already chosen.
- Both sets of SV are then independently rotated such that the resulting perturbations \mathbf{p}_j have the same globally average energy as the singular vectors but

smaller local maxima and more uniform spatial distribution. Note that rotation of singular vectors implies that none of the perturbations need to be in any ‘optimal’ direction [98].

- The SV in each set are then rescaled by choosing α_{jn} such that $\mathbf{p}_j = \sum_{k=1}^{25} \alpha_{jk} \mathbf{v}_k$ and $\|\mathbf{p}_j\| \sqrt{R_n} \leq \|\mathbf{a}_e\|$, where \mathbf{a}_e represents the approximated analysis error and R_n is taken to be 0.6. This allows initial perturbations to be larger over oceans, where data is sparse, than over more densely sampled land [14, 50].
- The corresponding members of the NH and SH sets of SV are then added together, and each resulting perturbation both added to and subtracted from the T42L31 analysis value.

Each ensemble member is then interpolated from T42L31 to T_l159L31 resolution and evolved at this resolution out to 10 days.

Construction of NCEP BV ensembles follows the procedure described previously as method B BV construction [94, 95]. BV ensembles consist of 8 twin pair perturbations, which are both constructed and evolved at T62L28 resolution. The analysis value is evolved both at this resolution and also under a higher resolution, T126L28, model. Initial BV perturbations are derived from 8 independent breeding cycles. These perturbations are added to and subtracted from the current analysis, and evolved at T62L28 resolution for a cycle time of 24 hours. (The evolution is continued out to 10 days to give the ensemble forecast.) The difference between each evolved twin pair is then scaled, using the kinetic energy norm, to a magnitude equal to the estimated seasonal analysis error; this scaled perturbation is then added to and subtracted from the latest analysis to form the new twin pair.

Evaluation of NWP ensemble prediction schemes uses a variety of norms. Ideally we would consider the total energy norm as it encompasses many variables and is considered to be the best norm in which to measure error [61]; unfortunately this requires many variables and leads to a quantity of data which is prohibitively large given the computational resources available in Oxford. Results for total energy computed in section 6.7 were made possible by the generous provision of computational resources by ECMWF. If the total energy norm is not feasible then, given the number of parameters used in the models, there is a vast choice of norms. The 500hPa geopotential height (z_{500}) is considered to capture the atmospheric dynamics and is traditionally used in prediction studies [89]. Geopotential heights, z , are the surfaces defined by the height above sea level on which the geopotential, Φ , is constant; they

are related by the equation $z = \frac{\Phi}{g_\phi}$, where g is the acceleration due to gravity and ϕ denotes the dependence of acceleration due to gravity on latitude. For the purposes of calculating geopotential height it is assumed that $g_\phi \simeq g_{38} = 9.80 \text{ ms}^{-2}$ [20]; a typical value for z500 is 5000m. In order to weight area equally, calculations in the z500 norm are weighted by the cosine of the latitude. Since z500 is considered to be the most linear norm [93] we also consider the 500hPa horizontal wind velocity (uv500) norm when the data is available; this is also weighted by the cosine of the latitude in order to have equal area weighting.

6.6 Linearity evaluation of operational NWP ensembles

In this section we examine ensembles initiated and evolved as part of the operational procedure for numerical weather prediction at the European Centre for Medium-range Weather Forecasting (ECMWF) and at the American National Centre for Environmental Prediction (NCEP). The forecasts of 500hPa geopotential height obtained from the ECMWF ensemble of September 1, 1998 are given in figure 6.13. (Ensemble members are typically initial perturbations of $\sim 5000\text{m} \pm 2\text{m}$ in the z500 norm.) Results for ECMWF SV ensembles are calculated for the 500hPa geopotential height over the Northern Hemisphere excluding Tropics (22.5°N-90°N) for the 25 twin pairs for each of 25 different cases, giving 625 twin SV pairs in total; the SV were initiated at 1200UTC on each day between December 11, 1996 and January 4, 1997 inclusive. Results for NCEP BV ensembles are presented using both the 500hPa geopotential height (z500) and the 500hPa horizontal wind velocity (uv500), over the Northern Hemisphere excluding the Tropics (22.5°N-90°N), unless specified otherwise. In the case of z500, results are averaged over the 8 twin pairs for each of 25 different cases, giving 200 twin BV pairs in total; the BV considered were initiated at 1200 on each of October 2, 4-6, 8-17, 20-22, 24-27, and November 21, 23, 25, 26, all in 1997. For uv500, 10 cases were considered (80 BV pairs in total), namely BV initiated on December 9-17 and December 19, 1997. Discontinuity in BV data reflects difficulties in accessing and down loading the data by hand from the ftp server.

6.6.1 Evaluation of linearity measures

As shown in figures 6.14 and 6.15, the average relative nonlinearity, (anti-)correlation and difference in magnitude of positive and negative evolved perturbations as a fraction of their sum, calculated in the z500 norm, are almost indistinguishable for the SV and BV ensembles over the entire period to 240 hours. The extent of the SV

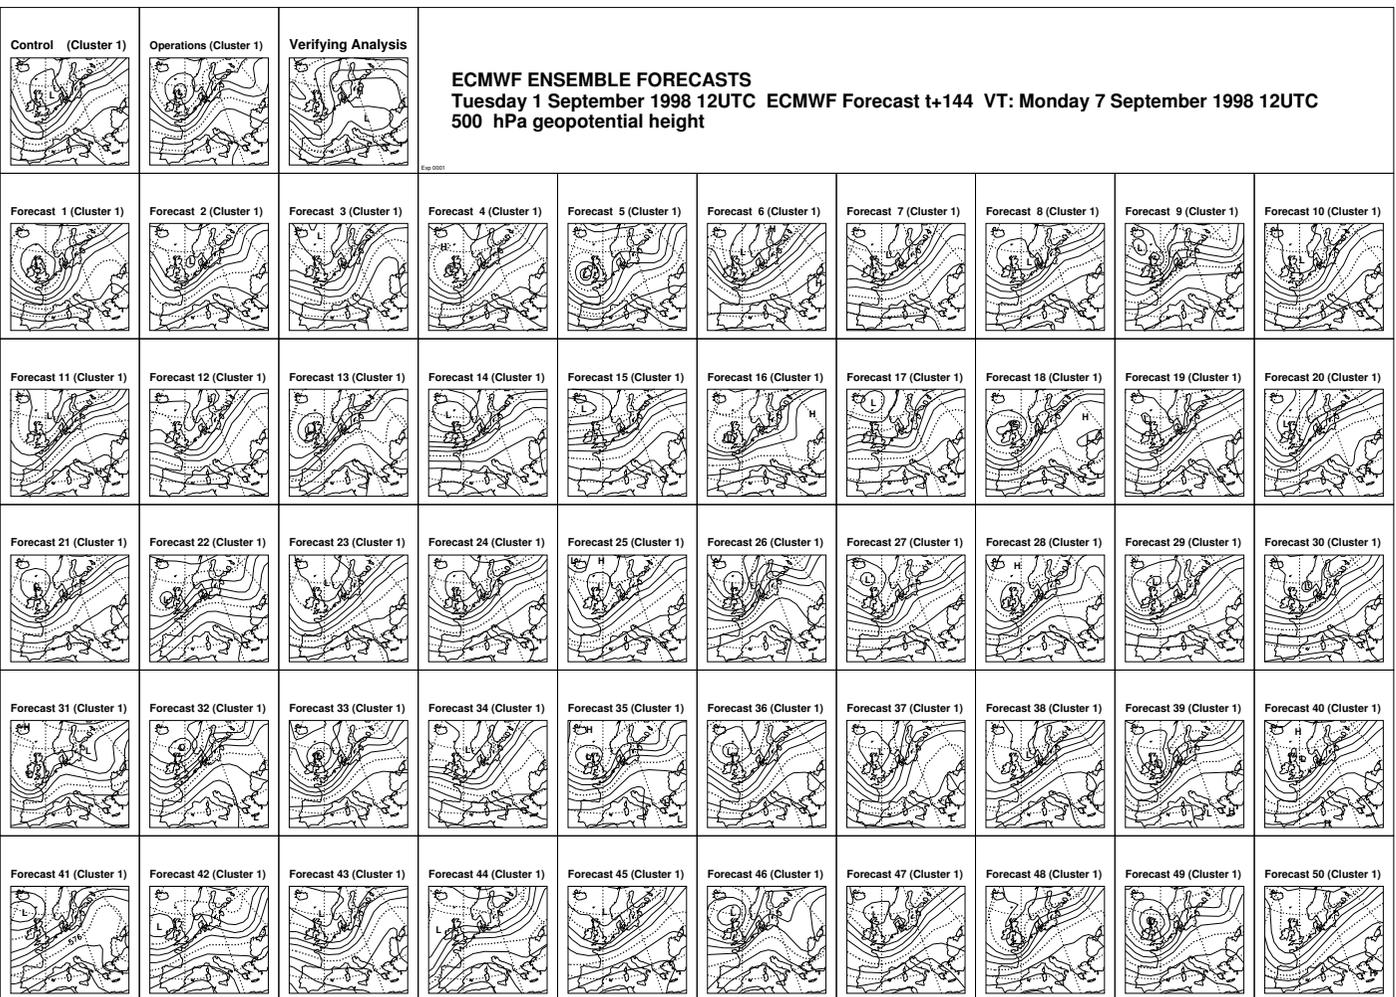


Figure 6.13: ECMWF ‘stamp map’ showing the 50 ensemble member predictions of the 500hPa geopotential height field for 7.ix.98 at 1200UTC. The ensemble members were initiated on 1.ix.98 at 1200UTC by adding perturbations, given by the SV, to the operational reference (whose evolved field is shown at the top left as ‘Operations’.) Also shown, in the top left, are the evolved high resolution ‘Control’ field and the ‘Verifying Analysis’; the latter is derived, using data assimilation, from the field actually observed on 7.ix.98 and nearby (in time) observations. Courtesy of ECMWF.

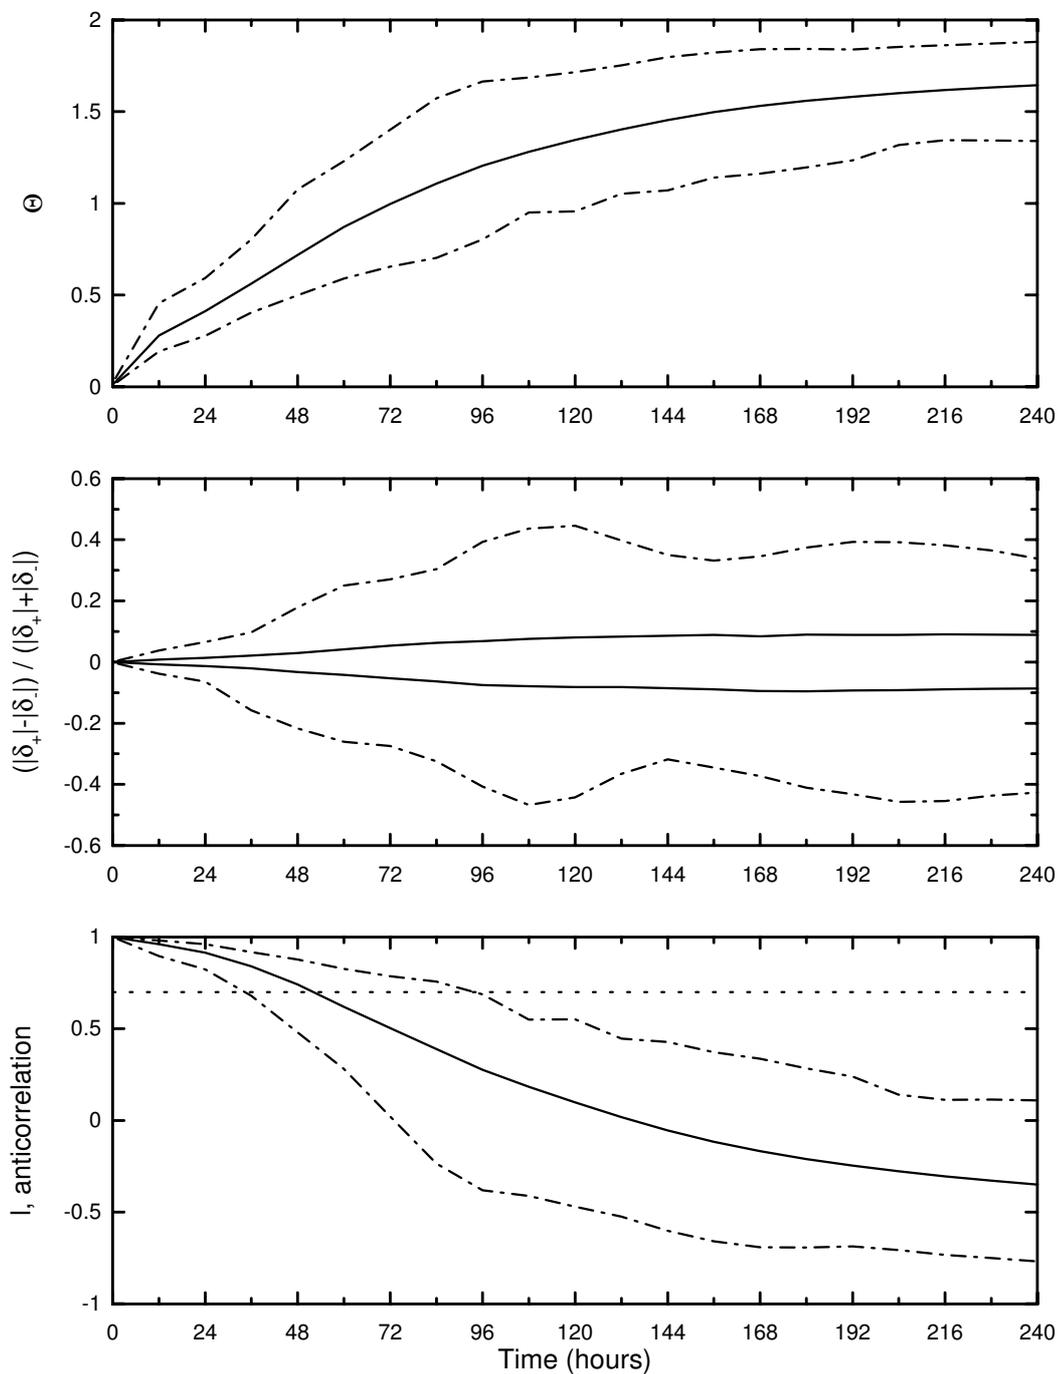


Figure 6.14: Linearity results for ECMWF operational twin SV perturbations ($\tau_{\text{opt}} = 48$ hours), calculated in the 500hPa geopotential height norm over the Northern Hemisphere excluding the Tropics and averaged over 25 days. The panels show, from top to bottom, the mean (solid line) and extent (dot-dashed lines) of error made in assuming linearity as measured by Θ , the difference in the magnitude of the positive and negative evolved perturbations as a fraction of their sum, and the (anti-)correlation between twin pairs.

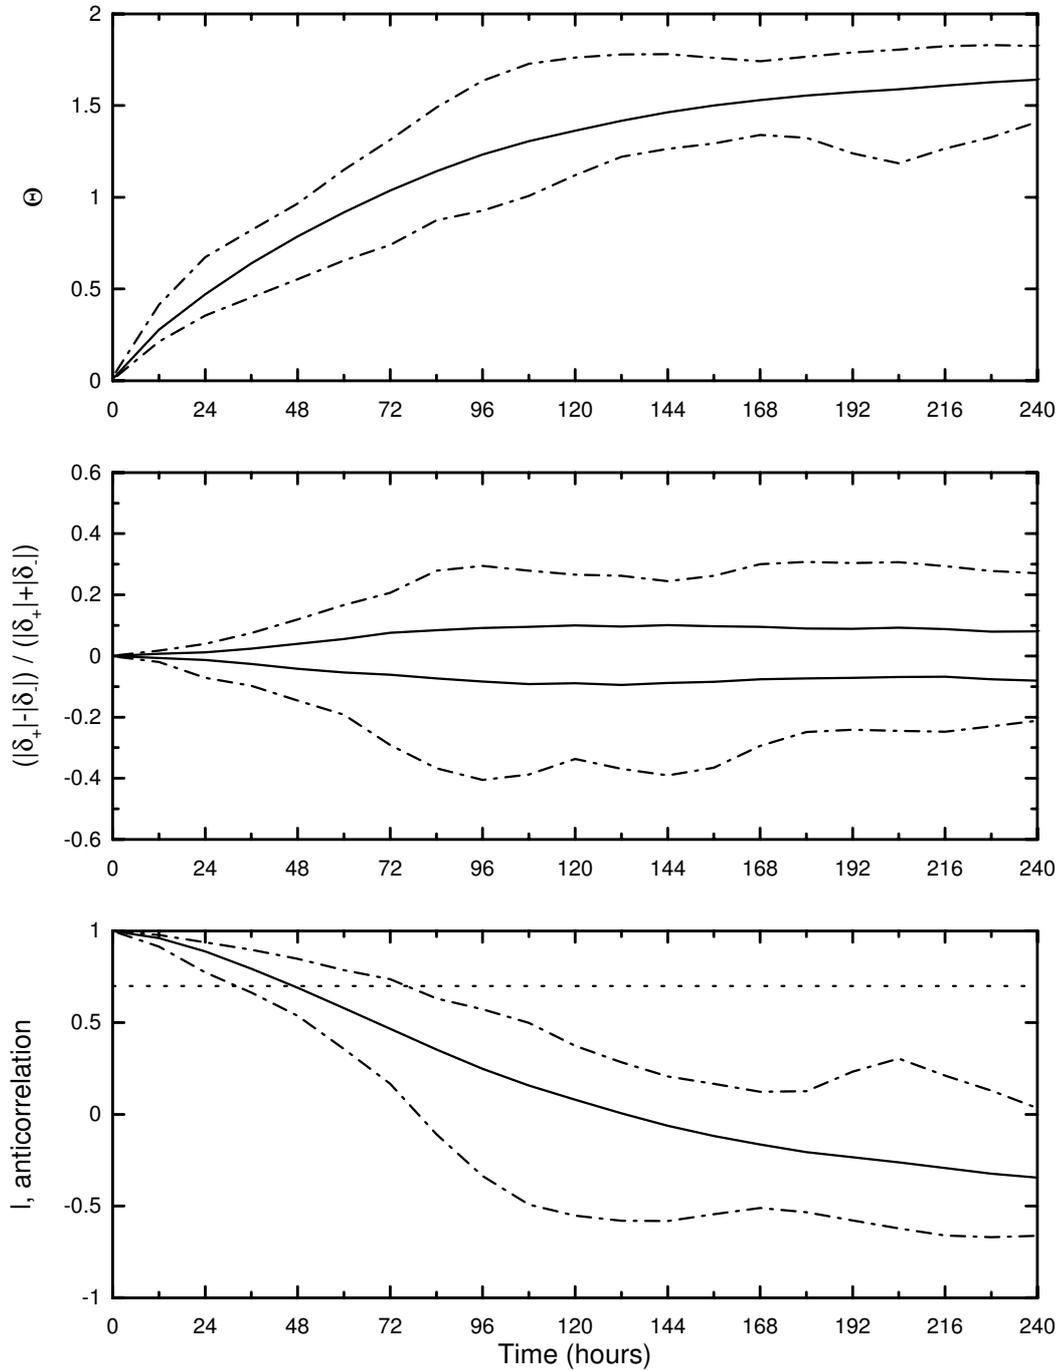


Figure 6.15: Linearity results for NCEP operational twin BV perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere excluding the Tropics and averaged over 25 days. The panels are as in figure 6.14.

results is generally larger than that of the BV results. At the optimisation time of 48 hours, the error in assuming linearity ranges from 50% to 108% the average magnitude of the evolved perturbations, with an average of 72%. From the second panel in figure 6.14 we can see that on average the difference in magnitude between the positive and negative perturbations is less than 5% of their summed magnitudes, so the nonlinearity is almost all due to the perturbations evolving in different directions. This corroborates the average correlation of 0.74 at 48 hours; as explained above a correlation of 0.7 represents 75% relative nonlinearity if the perturbations are of the same magnitude, more if they differ. A relative nonlinearity of 0.72 at optimisation time, and an (anti-)correlation of 0.74 suggest that *SV ensembles*, for SV optimised over 48 hours, *are not internally consistent*.

For operational BV ensembles, the assumption in their definition is that evolution of BV perturbations is approximately linear for the breeding cycle time, 24 hours. The relative nonlinearity has an average value of 0.47, and the correlation an average of 0.89, at 24 hours; again the nonlinearity is almost all due to the perturbations evolving in different directions. An error of almost 50% in assuming linearity is not considered to be acceptable here, hence *BV ensembles are not internally consistent*.

Having established that the SV perturbations evolve nonlinearly, it is of interest to ask if the SV perturbation corresponding to the largest singular value displays nonlinear growth first³. Figure 6.16 presents the results for specific SV twin pairs for 1 day. In this example the effect of nonlinear evolution is similar for the first 12 hours for the 6 SV twin pair perturbations shown; after 12 hours it is the twentieth SV twin pair which grows most nonlinearly as a result of discrepancies in both magnitude and orientation. Variation in nonlinear ranking of the SV twin pairs can be seen as they all evolve to a similar saturating value of relative nonlinearity and correlation; the variation in magnitude discrepancies within each pair shows no tendency to converge to a single value.

If the 500hPa geopotential height fields corresponding to each SV perturbation were drawn, they would each be a pattern of synoptic structures, *i.e.* of ‘high’ and ‘low’ pressures. It is of interest to ask whether the SV which are most nonlinear at a given time combine in such a way as to reflect the synoptic structures seen in the corresponding control forecast; do the regions of ‘strong’ nonlinearity growth correlate with the forecast synoptic structures? In order to quantify the local nonlinearity, we

³When the SV perturbations are rotated during initialisation it is possible to calculate the modified singular values for the perturbations, hence the first rotated SV is still that with largest singular value.

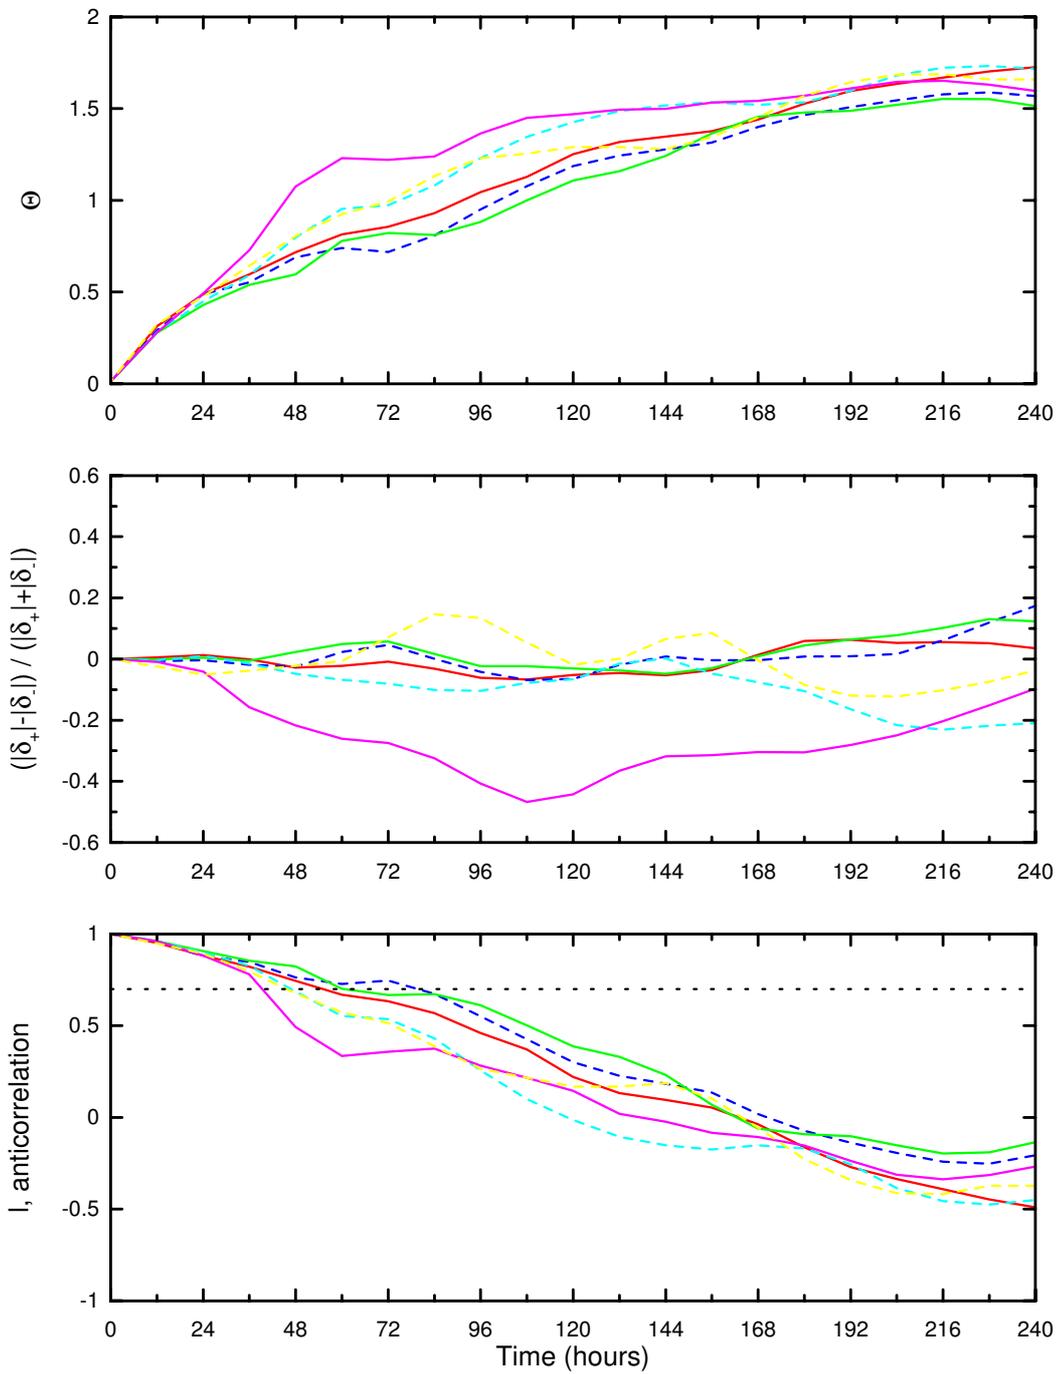


Figure 6.16: Linearity results for ECMWF operational twin SV perturbations ($\tau_{\text{opt}} = 48$ hours); shown are results for the first (red, solid line), fifth (blue, dashed), tenth (green, solid), fifteenth (light blue, dashed), twentieth (magenta, solid) and twenty-fifth (yellow, dashed) twin pairs. All results are calculated in the 500hPa geopotential height norm over the Northern Hemisphere excluding the Tropics for 19.xii.96. The panels are as in 6.14.

may calculate Θ at each grid point. Since the z500 norm is one dimensional, $\Theta = 2$ whenever both of the twin perturbations lie on the same side of the control. In the top panel in figure 6.17 the colour at each grid point is determined by the number of twin pairs for which there is extreme nonlinearity, taken as when $\Theta = 2$. When compared with the corresponding forecast, some correlation of regions with many nonlinear pairs and synoptic structures is evident (*e.g.* over Greenland and Western Canada); statistically the correlation is significant. It may be hypothesised that the correlation occurs because of ‘saturation’ of one of the twin pair perturbations in regions of synoptic structures: the perturbation which is towards a more extreme state (*e.g.* a positive perturbation in the pressure near a high pressure structure) may have its growth curbed, or it may evolve so that both perturbations lie on the side of the control away from the extreme state, but is unlikely to continue growing at the same rate in the direction of the extreme state.

Results for SV ensembles were only available over the Northern Hemisphere for z500; results for BV ensembles are available Worldwide for a variety of norms and enable investigation of the effect of both choice of region over which norms are calculated and choice of norm itself. Figure 6.18 shows relative nonlinearity and correlation results for the z500 norm calculated over different regions. Once again results only differ significantly in their spread at any given time: the Tropics (22.5°S-22.5°N) have greatest spread, followed by the Northern Hemisphere and then Worldwide. Choice of region over which to calculate the results appears to matter little. At 48 hours the error made in assuming linearity is at least 50% in all the BV perturbation cases considered.

Choice of norm has more impact on the results. Figure 6.19 shows linearity results for BV perturbations calculated in the 500hPa horizontal wind velocity norm (uv500), as well as those for ECMWF SV and NCEP BV calculated in the z500 norm for comparison. The increase in relative nonlinearity, and decrease in correlation, is more rapid, with an averages of 0.71 and 0.75 respectively at the cycle time of 24 hours. When the uv500 norm is used, linearity results corroborate the result that operational BV ensembles are not consistent with the linearity assumption made in their definition.

6.6.2 Doubling times

All SV perturbations grew by a factor of 16 within the time interval of 10 days for which data was available, and while all BV perturbations grew by a factor of 4, none grew by a factor of 16 within this period. Results for SV suggest roughly uniform

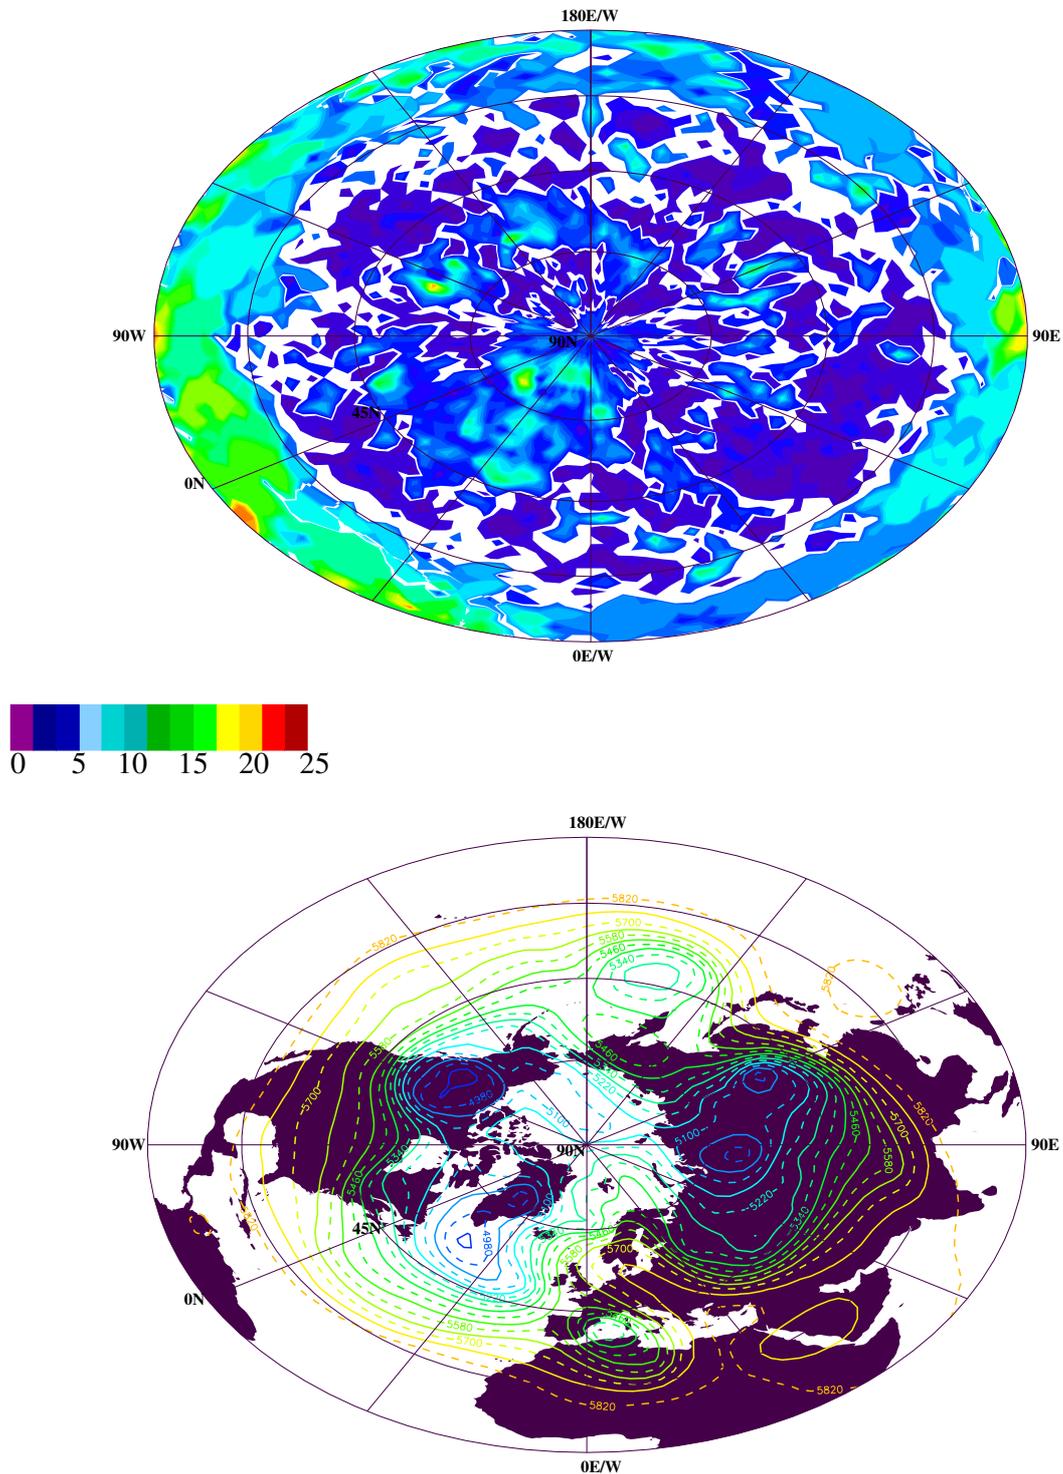


Figure 6.17: Lambert equal area projections of the Northern Hemisphere of ECMWF 36 hour forecasts initiated at 1200 on 19.xii.96. The number of pairs of SV twin perturbations for which the local nonlinearity is maximal, *i.e.* for which $\Theta = 2$ at that grid point, is denoted by the shading (as shown by the colour bar) (top). Comparing regions of strong nonlinearity with the synoptic structures ('highs' and 'lows') seen on the low resolution control forecast (bottom), some degree of correlation is evident, *e.g.* over Greenland and Western Canada.

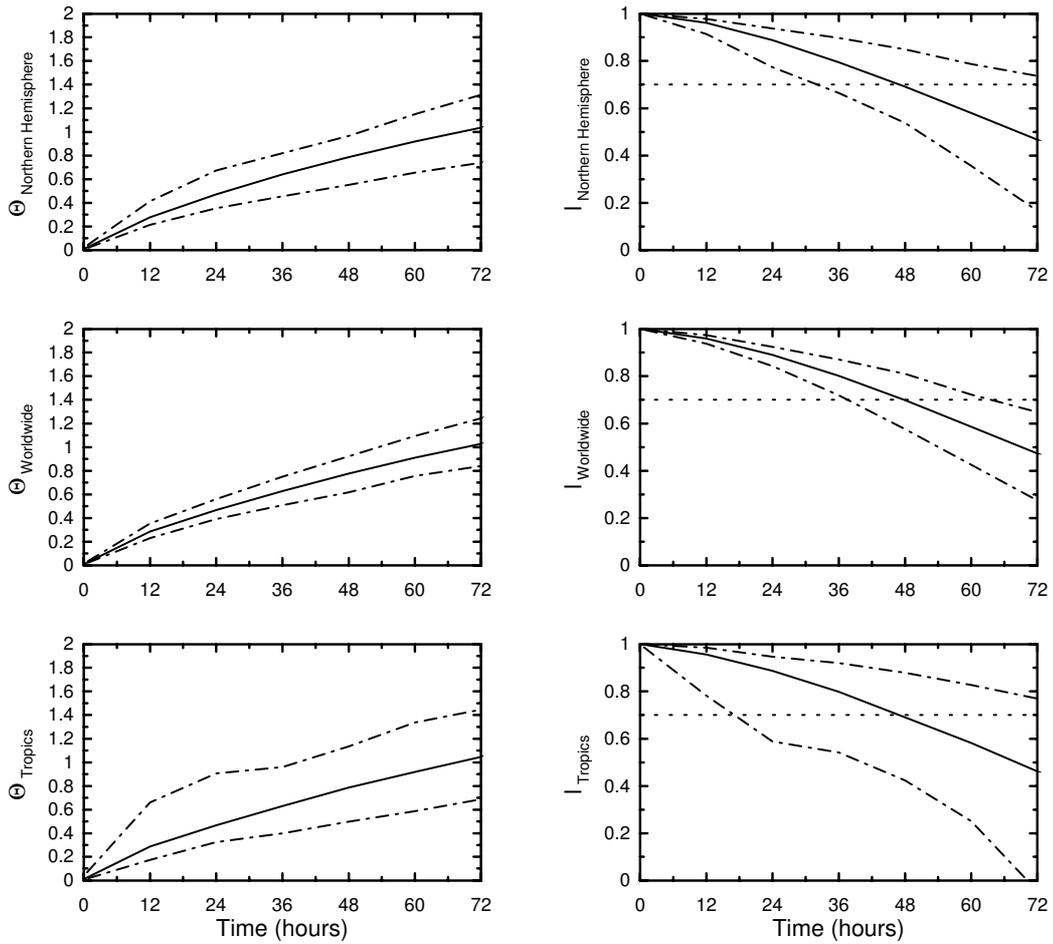


Figure 6.18: Linearity results for NCEP operational twin BV perturbations, calculated in the 500hPa geopotential height norm over different regions of the world, and averaged over 25 days. The panels show the mean (solid line) and extent (dot-dashed lines) of values of Θ (left) and the (anti-)correlation (right), as defined in the text, calculated over the Northern Hemisphere (top), Worldwide (middle), and over the Tropics (bottom).

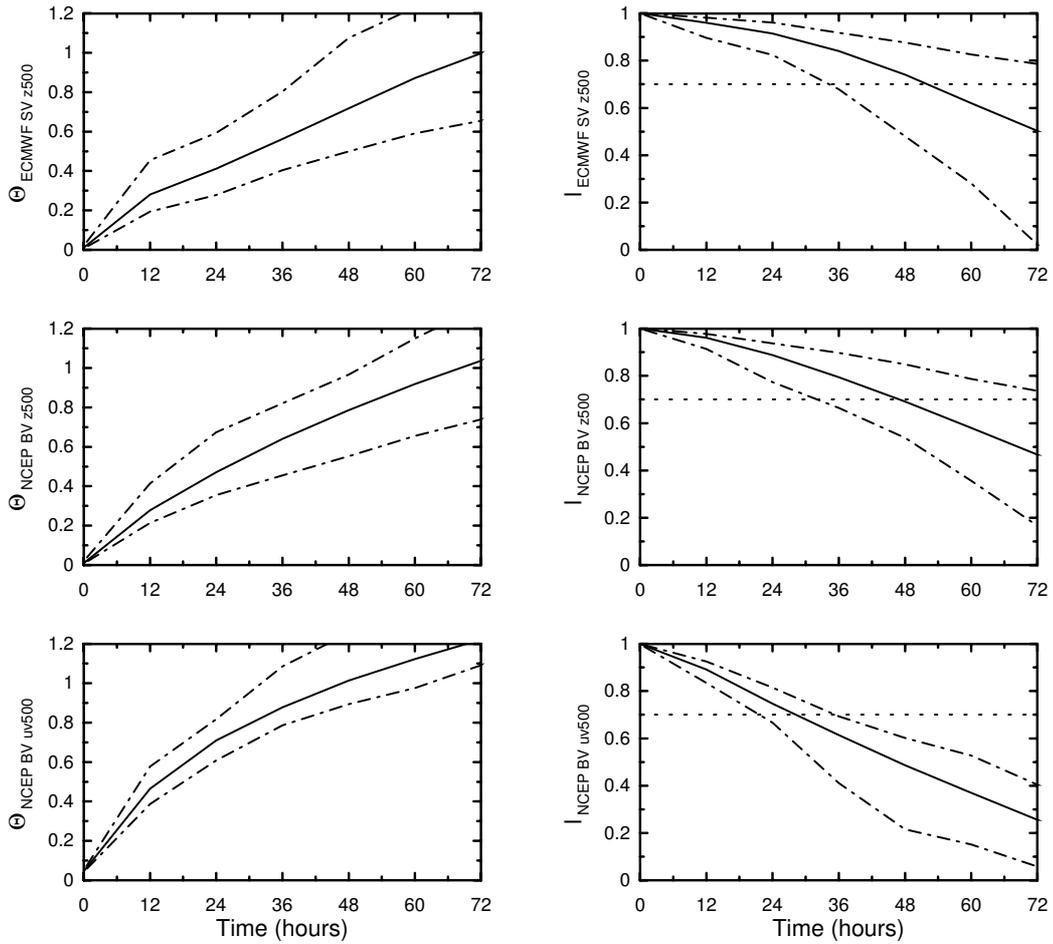


Figure 6.19: Linearity results, calculated over the Northern Hemisphere, for operational twin ECMWF SV perturbations calculated in the 500hPa geopotential height norm (top), for NCEP BV perturbations calculated in the 500hPa geopotential height norm (middle), and for NCEP BV perturbations calculated in the 500hPa horizontal wind velocity norm (bottom). The panels are as in figure 6.18.

exponential growth until the SV have grown by a factor of 4; the average value lies on the line $y = 2x$ and the histogram has over 85% of initial conditions with a quadrupling:doubling ratio of 2, see figure 6.20. All doubling and quadrupling times

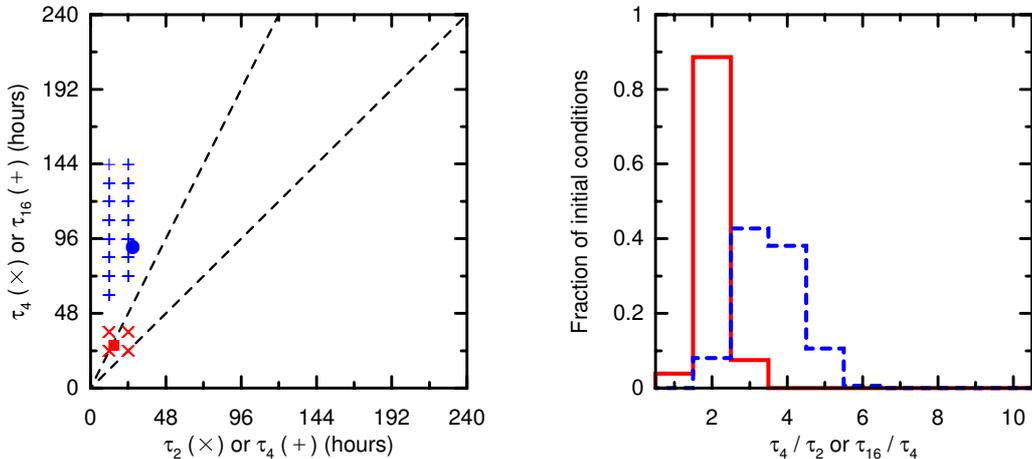


Figure 6.20: q -pling times of ECMWF SV perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere. A histogram of the ratio of q^2 -pling to q -pling times (right) gives information obscured by over-plotting in the scatterplot (left) of quadrupling times vs. doubling times (red crosses, with average shown by solid square) and of the time to grow by a factor of 16 vs. quadrupling times (blue plus signs, with average shown by solid circle). The lines $y = x$ and $y = 2x$ are also shown (dashed lines).

are less than 36 hours however, and the data is only available at 12 hour intervals; results may differ if more regularly sampled data is used. The larger distribution of times at which the SV perturbation have grown by a factor of 16 leave no doubt that, even if growth is uniformly exponential out to 36 hours, SV perturbations do not continue to grow uniformly exponentially after 36 hours. Considerable variation is seen in both the scatterplot and histogram for doubling and quadrupling times for BV perturbations, see figure 6.21; growth of BV perturbations is not uniform.

6.6.3 Summary of linearity evaluation measures

Operational SV ensembles appear not to be internally consistent with the assumption that the TLM is valid out until the optimisation time; the nonlinear growth of perturbations is due mainly to discrepancies in their orientations. Linearity results from the both the Marzec-Spiegel and annulus suggest that if internally consistency is to be improved, either the initial perturbation amplitude must be decreased and/or the optimisation time must reduced. Ideally, the perturbation magnitude needs to reflect the analysis error [50], while the optimisation time is less constrained by other consid-

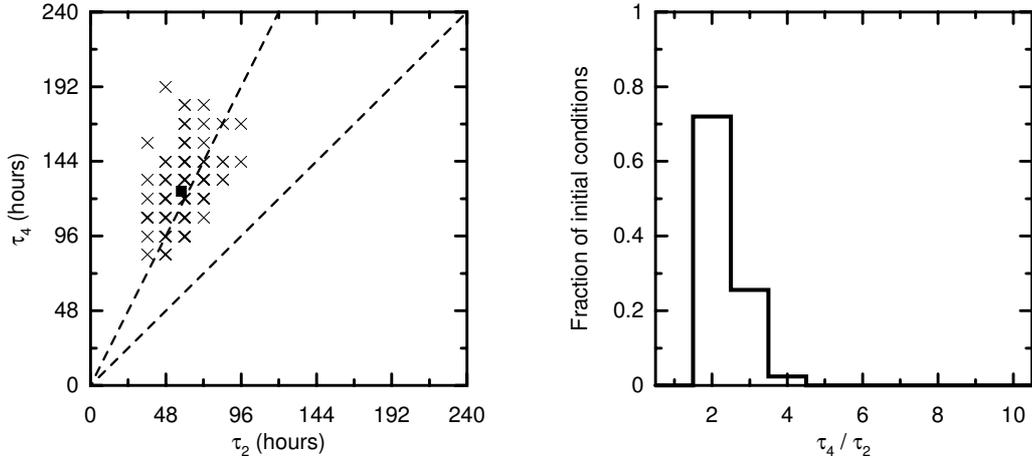


Figure 6.21: q -pling times of NCEP BV perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere, with panels as in fig. 6.11, but showing only doubling and quadrupling times.

erations⁴. For a 48 hour optimisation time, the first (rotated) SV does not always have the most nonlinear initial growth. There is evidence that nonlinear growth is a local phenomena, and that centres of this nonlinearity correlate to synoptic structures. It would be of interest to calculate the correlation between the forecast regions of nonlinearity and the verifying analysis, (related to) the observations actually recorded at the time of interest, and to perform these calculations for multiple cases. The sampling time may be too large to provide an accurate portrayal of doubling times; an alternative method using amplification factors (*i.e.* the growth over a given time rather than the time for a given growth) is presented in the following section for optimisation times of both 24 hours and 48 hours. Results presented are all calculated using the 500hPa geopotential norm over the Northern Hemisphere; results from BV perturbations suggest that the choice of region is unimportant, but that the norm chosen may affect results. Due to the availability of more data and more computational resources at ECMWF, some results in the total energy norm are presented below.

From the results calculated both in the 500hPa geopotential height and the horizontal wind velocity norms, BV perturbations are not internally consistent with the assumption that the BV perturbations evolve linearly over the cycle time of 24 hours. Comparison of results in the two norms corroborates the view of those working with NWP models that the z500 norm is the most linear [93], and suggests that the op-

⁴Collaboration with ECMWF has enabled the formation of SV ensembles, with SV optimised over 24 hours. Discussion of the construction of such an ensemble, and presentation and discussion of results comparing SV optimised over both 24 and 48 hours are in the following section.

erational BV perturbations do not lie in the directions as described in the formal definition of breeding vectors. BV constructed using methods A and B (for details see section 5.1.3.4) for either the Marzec-Spiegel system or the annulus show little difference in either ensemble performance measures (see sections 5.2.2 and 5.2.3). Note that the cases for which no significant difference is seen are those for which BV ensembles are also internally consistent. The internal consistency of BV perturbations of initial magnitude 0.1 evolved under the RBF model of the annulus is marginal; in this case projection of method A BV perturbations onto the dream perturbation is slightly better than for method B BV perturbations. Operational (method B) BV are not internally consistent and it would therefore be of interest to see how BV ensembles constructed using method A perform in the operational model, since method A of BV construction does not make any assumption about linear evolution of BV perturbations, in contrast to method B.

6.7 Comparing 48h SV and 24h SV ensembles

Linearity results for operational ECMWF SV ensembles, optimised over 48 hours, suggest that they are not internally consistent with the assumption that the evolution of perturbations is well approximated by the tangent linear model used in their definition. *Barkmeijer* [6], using a low resolution model, recognises that growth of SV perturbations is nonlinear prior to the optimisation time; he offers an alternative method of SV perturbation construction, which is effectively an iterative method using a shorter optimisation time. The simplest response to the lack of internal consistency in the operational SV ensembles is either to reduce the perturbation magnitude or to reduce the optimisation time over which the SV are formed; the perturbation magnitude needs to reflect the analysis error [50], while the optimisation time is less constrained by other considerations. Collaboration with ECMWF has enabled access to their operational models to construct and evolve SV ensembles optimised over both 48 hours and 24 hours for 7 days, using otherwise identical model parameters. It is these results which are presented in this section. In addition, the enhanced computational facilities enabled calculations in the total energy norm, and the evolution of ensembles in both low and medium resolution models.

Referring the reader back to the description of SV ensemble formulation in section 6.5, two sets of SV ensembles were constructed for January 1, 1998: the first used SV optimised over 48 hours (48h SV) while the second used SV optimised over the shorter time of 24 hours (24h SV); in both cases SV were only selected over the Northern

Hemisphere extra Tropics and the same rescaling factor, $R_n = 0.6$, was used to scale the singular values to be comparable to the analysis error. In addition to the two rotated sets of perturbations, two sets of perturbations were generated simply by rescaling the SV without rotating them, in order to allow thorough comparison of the (pure) 48h and 24h SV. The singular value spectra of the 48h and 24h SV are shown in figure 6.22. The 48h SV singular values are larger than those of the 24h SV; this is to be expected as the singular values represent the expected amplification factor at the optimisation time.

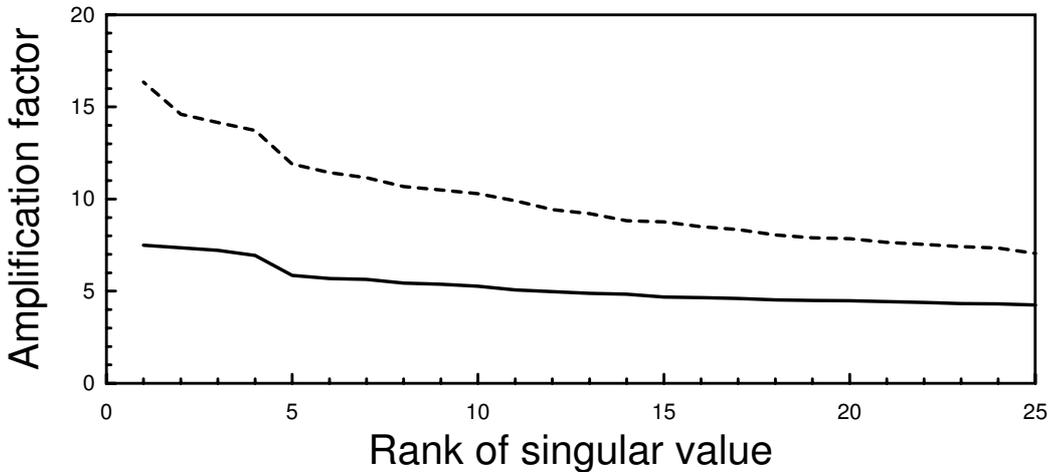


Figure 6.22: Singular value spectra for 48h SV (dashed line) and 24h SV (solid line) constructed, using the ECMWF T42L31 model, for 1.i.98.

In order to compare the SV subspaces generated using optimisation times of 48h and 24h, the averages of the first ten (rotated) SV perturbations for each case are shown in figure 6.23. The averages are calculated using the total energy norm and each contribution weighted according to its singular value (calculated from the proportions of each SV the perturbation includes); while there is considerable similarity between the two fields, they also differ significantly; the ‘similarity index’ (as defined in section 5.2.1) between the two subspaces (spanned by all 25 SV) is $\sim 40\%$.

6.7.1 Evaluation of linearity measures for one day

The four ensembles, 48h and 24h SV, rotated and non-rotated, were evolved forward under the $T_{l159L31}$ model to see both how they compared and also to see if this comparison was the same for non-rotated SV ensembles. The ensembles were also evolved under the T42L31 model in order to see if nonlinearities were present in the low resolution model and amplified by the higher resolution, or if the nonlinearities only appeared at higher resolution. From the linearity results for rotated

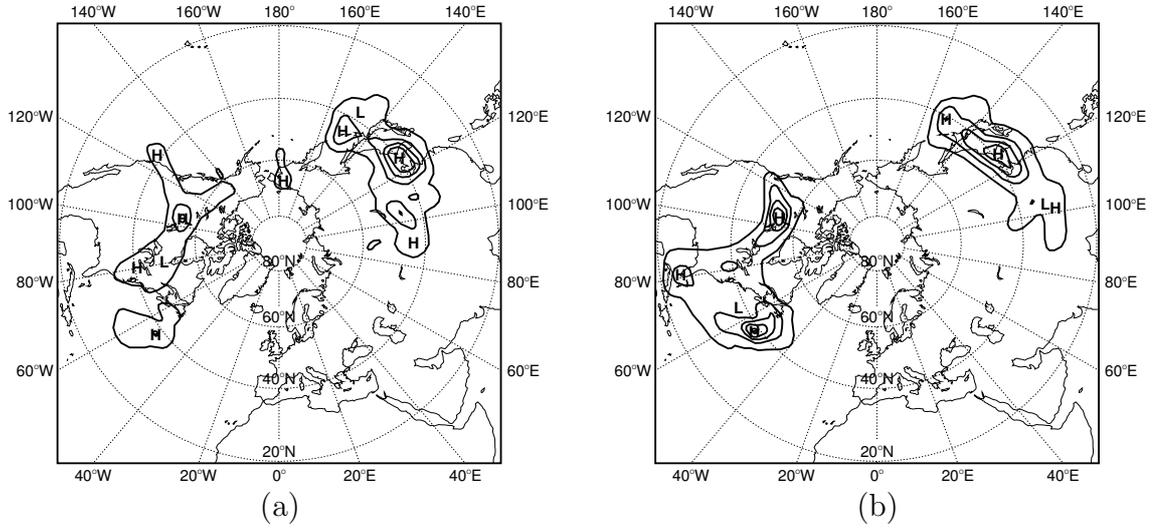


Figure 6.23: Total energy fields of 48h SV (left) and 24h SV (right) constructed for 1.i.98. The fields shown are an average over the first 10 perturbation fields included in the ensemble, weighted according to the corresponding singular values. Courtesy of ECMWF.

perturbations evolved at high resolution and non-rotated perturbations evolved at both resolutions (shown for 48h and 24h SV in figures 6.24 and 6.25 respectively), we can see that the 48h and 24h SV behave similarly; the results are only available for a limited number of cases since the calculations are in the total energy norm. Evolution of SV perturbations at higher resolution enhances, but is not the source of, the nonlinear perturbation growth. Rotation of the SV in the formation of SV perturbations generally, but not always, decreases the relative nonlinearity of growth. This result is corroborated by the relative nonlinearity in the first non-rotated SV pair which is usually, but not always (see T_l159 non-rotated 24h SV linearity results), one of the largest. Results for correlation portray the same onset of nonlinearity as the relative nonlinearity, *i.e.* Θ , results suggesting that, as for the 48h rotated SV twin pairs, nonlinearity of growth is mainly due to differences in orientation, rather than magnitude, between the twin perturbations.

As described in the previous section, the SV are rescaled so that the magnitudes of their local maxima are proportional to the magnitude to the estimated analysis error. The constant of proportionality, $\sqrt{R_n}$, is determined in order that the spread (or extent) of the ensemble usually contains the high resolution control forecast. Comparing the amplitude of 24h SV with those of 48h SV (at times from 48h to 84h) in the total energy norm suggests that for the 24h SV the constant of proportionality should be increased to $R_n = 1.0$, representing a 29% increase in the initial magnitude of 24h SV perturbations. For all subsequent results, the 24h SV perturbations are

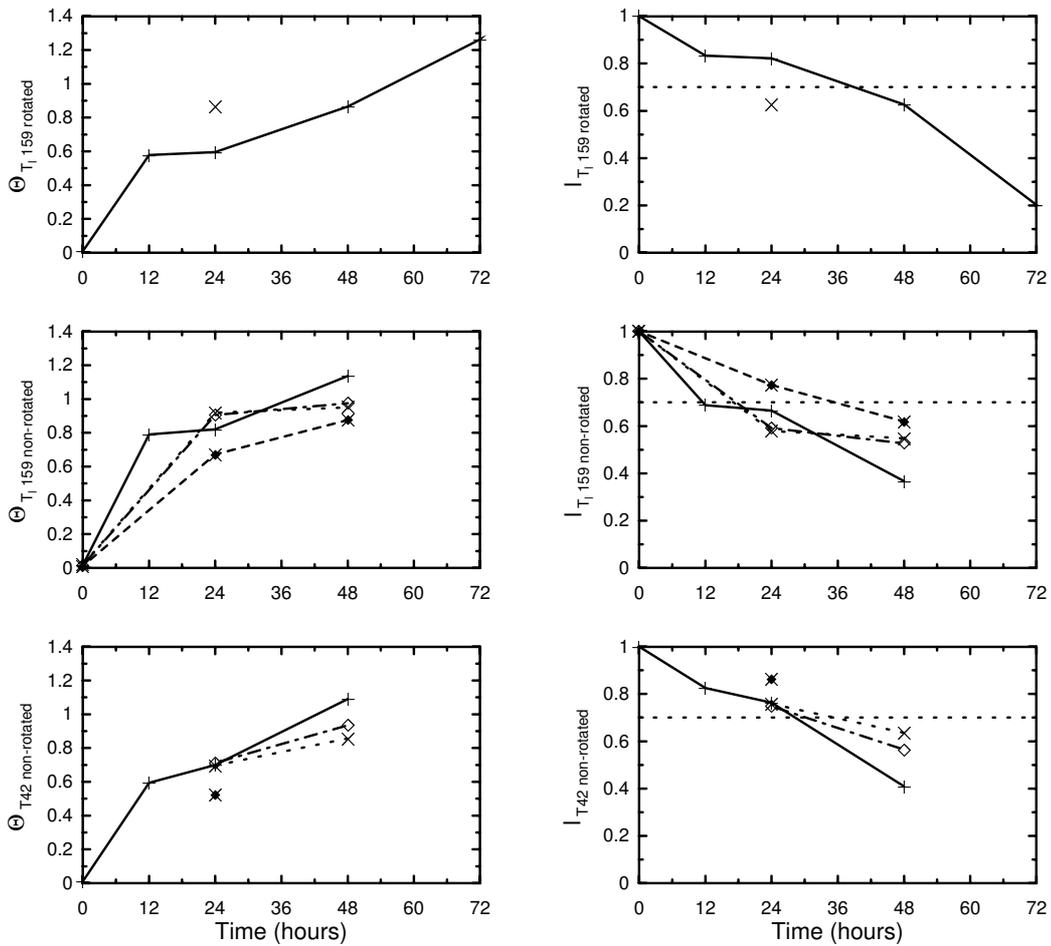


Figure 6.24: Linearity results for ECMWF twin SV perturbations, optimised over 48h, calculated in the total energy norm over the Northern Hemisphere for 1.i.98. The panels show the values of Θ (left) and the (anti-)correlation (right). Results for the first (pluses joined by solid line) and fifth (cross) rotated twin pairs are shown (top). Results from ensembles formed using non-rotated SV perturbations, evolved at $T_1 159$ (middle) and $T42$ (bottom) resolution are also shown for the first (plus, solid line), fifth (cross, dotted line), tenth (diamond, dot-dashed line) and twentieth (cross diamond, dashed line) twin pairs; all show that the error in assuming linearity at 48h is $> 80\%$.

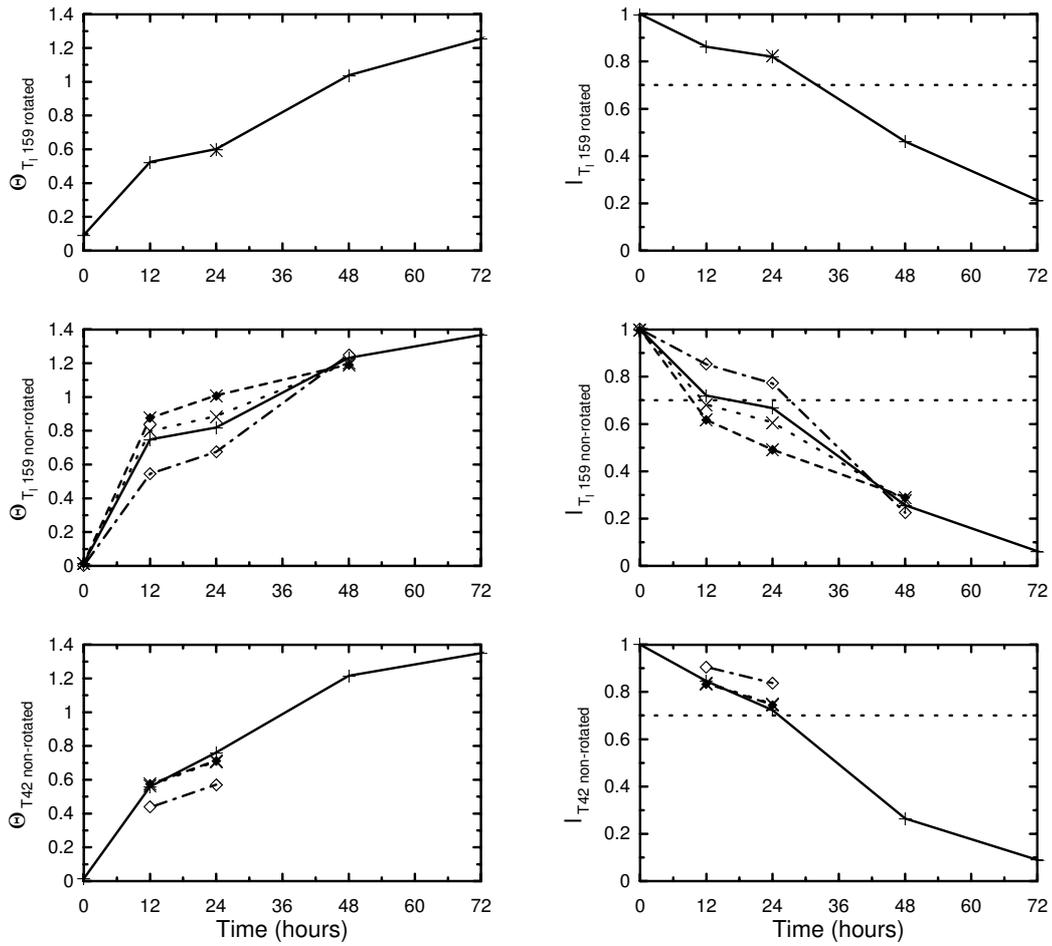


Figure 6.25: Linearity results for ECMWF twin 24h SV perturbations calculated in the total energy norm over the Northern Hemisphere for 1.i.98. The panels are as in figure 6.24 with the same notation, namely results are shown for the first (plus, solid line), fifth (cross, dotted line), tenth (diamond, dot-dashed line) and twentieth (cross diamond, dashed line) twin pairs.

24h SV increased magnitude perturbations. In order to easily compute results out to 10 days for all 25 perturbation pairs, the 500hPa geopotential height (z500) is used below. The change in initial 24h SV perturbation magnitude is implemented and the new 24h SV ensembles, formed using both rotated and non-rotated SV, are evolved under the usual (T_l159) resolution. The amplitude of evolved 24h SV perturbations is found to be similar to that of evolved 48h SV perturbations over at least 48h to 84h (in the z500 norm).

A complete set of linearity results are shown in figure 6.26. While the general trends are the same as for total energy, the nonlinearity growth is considerably *less* in the z500 norm than in the total energy norm. The averages over non-rotated perturbations continue to show enhanced nonlinear growth compared to the averages over the rotated perturbations. The wider range of values for the rotated perturbations may be due to their initiation in different orientations. 24h SV perturbations appear to display more rapid initial nonlinear growth than 48h SV.

6.7.2 Amplification factors for one day

We wish to compare the amplification factors of the perturbations with the expected linear growth over a given time interval. The linear growth over the optimisation time is given by the singular values; linear growths over other time interval may be calculated by scaling the singular values appropriately. The 24h SV singular values and the square root of the 48h SV singular values give the expected linear growth over the initial 24 hours; similarly for a 48 hour interval the square of the 24h SV singular values is used. Following *Lorenz* [46], the middle panel of figure 6.27 includes line segments whose gradient is proportional to these scaled singular values (with a fixed constant of proportionality). While the scaled singular values are at best expected to be relevant for the first 24 or 48 hours, they are drawn for each 24 hour interval as a reference. Also included are line segments whose gradient is proportional to the actual amplification of the nonlinearly evolved SV perturbations, *i.e.* $\frac{\|\mathbf{P}_1\|_{t=24h}}{\|\mathbf{P}_1\|_{t=0h}}, \frac{\|\mathbf{P}_1\|_{t=48h}}{\|\mathbf{P}_1\|_{t=24h}}, \dots$, for both the first 48h SV and first 24h SV non-rotated perturbations.

Consider the middle panel of figure 6.27, which compares singular values with (nonlinear) amplification factors over 24 hour intervals for the first non-rotated SV perturbation of both 48h SV and 24h SV ensembles. Line segments with gradients proportional to the 24h SV singular value and the square root of the 48h SV singular value are drawn for each 24 hour interval. The amplification factor from 0 to 24 hours is the calculated for both perturbations, and the line segments with gradients proportional to these factors drawn at $t = 0h$. The amplification factors (of both

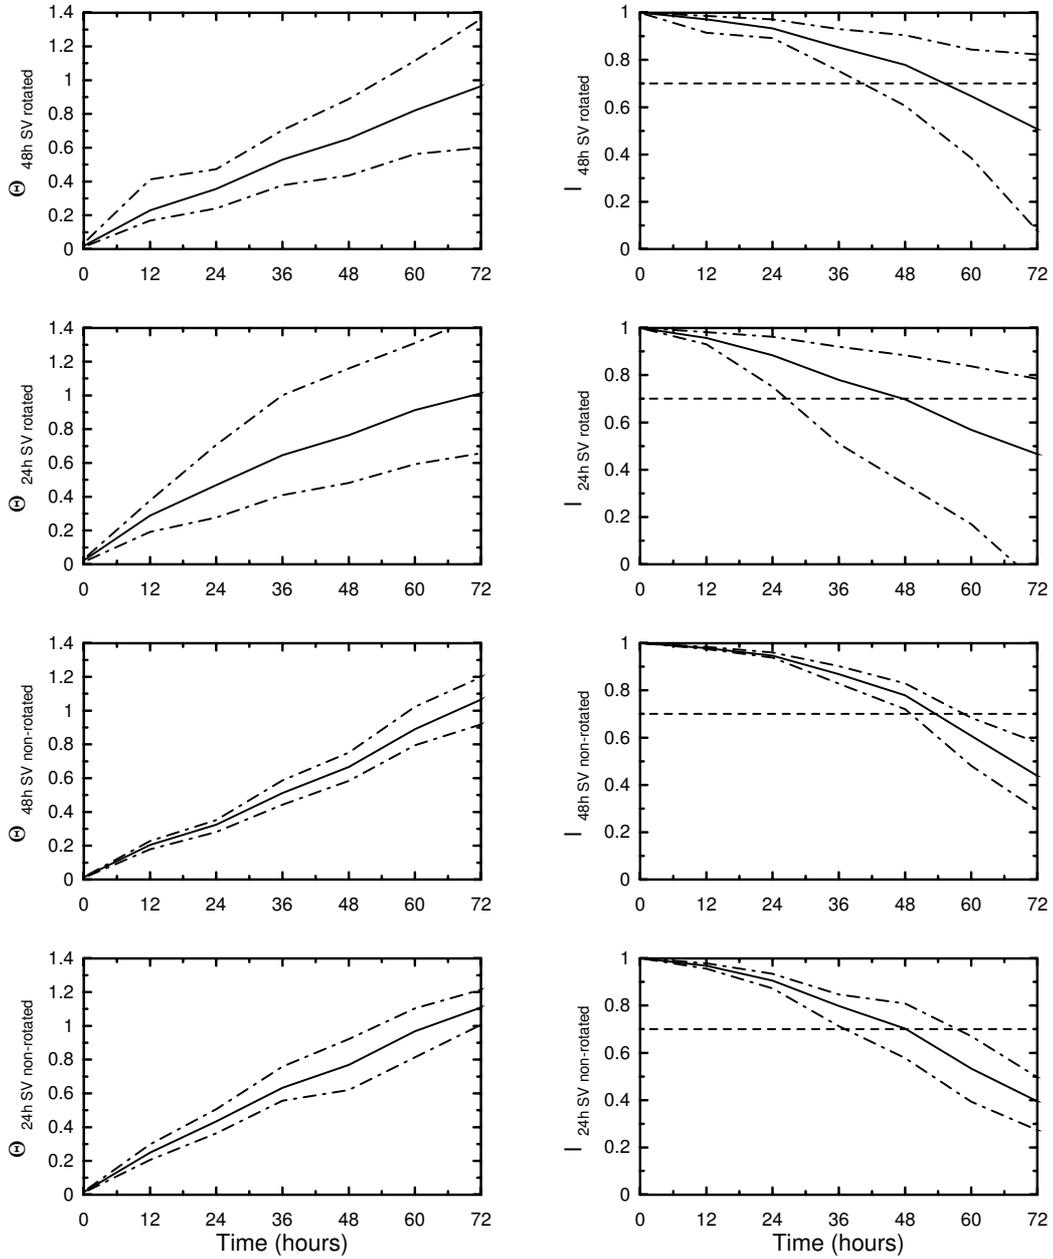


Figure 6.26: Linearity results for ECMWF twin rotated and non-rotated, 48h SV and 24h SV, perturbations calculated in the 500hPa geopotential height norm over the Northern Hemisphere for 1.i.98. The panels show the mean (solid line) and 4 and 96 percentiles (dot-dashed lines) of Θ (left) and the (anti-)correlation between twin pairs (right), as defined in the text. Results are for, from top to bottom, 48h SV rotated perturbations, 24h SV rotated, 48h SV non-rotated and 24h SV non-rotated.

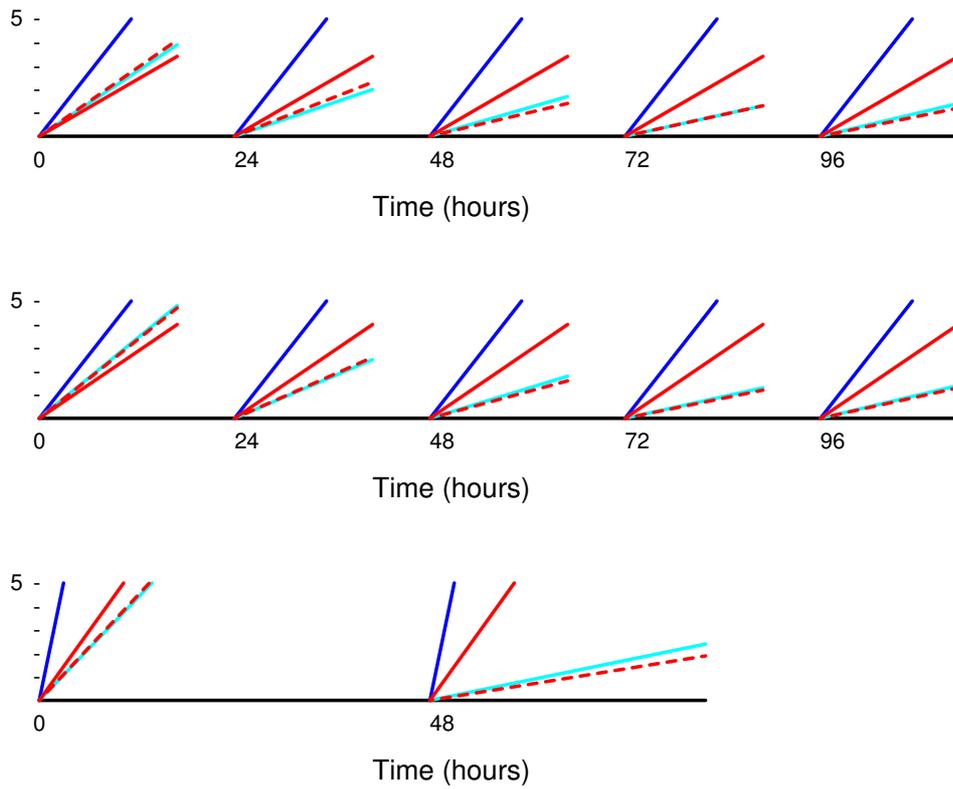


Figure 6.27: Amplification of ECMWF SV perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere for 1.i.98; the amplification factor of each segment, over the period given on the horizontal scale, is indicated by projection of the relevant line segment onto the vertical scale. For reference, line segments with gradients proportional to the 48h singular values (red solid lines) and 24h singular values (blue solid lines), are shown. (The line lengths have no relevance.) Results for the first rotated SV perturbation (24h interval, top) and first non-rotated SV perturbation (24h interval, middle; 48h interval, bottom) are given for 48h SV (red dashed lines) and 24h SV (cyan solid lines).

48h SV and 24h SV perturbations) over the first 24 hours are found to be very similar. On repeating this calculation over subsequent 24 hour intervals (to 120h), the amplification factors for the two perturbations continue to be very similar. From the middle panel of figure 6.27 we can see that, for the first 24h interval, the amplification factors are less than the 24h SV singular value, but greater than the (scaled) 48h SV singular value; for subsequent (24 hour) time intervals, the amplification factors are smaller than both (scaled) singular values. From this figure we may conclude that the amplification of the 24h SV perturbation over any 24 hour interval (upto 120 hours) is less than the growth suggested by the first (24h SV) singular value. In contrast, the 48h SV perturbation amplifies more rapidly than is suggested by the (scaled 48h SV) singular value over the first 24 hour interval, and then less rapidly thereafter. When 48 hour intervals are considered (see the lower panel in figure 6.27), the amplification of both the first 48h SV and 24h SV perturbations over the both 48 hour intervals is less than that suggested by either (scaled) singular value.

If it is known which SV are contained in the rotated perturbations, and with what weightings, the singular values (of the non-rotated SV) may be scaled accordingly to give a singular value for the rotated SV perturbation. This enables a similar plot, comparing singular values with amplification factors, for the first rotated SV perturbations, as shown in the top panel of figure 6.27. Amplification factors of both 48h and 24h SV rotated perturbations are still similar throughout the interval, but are smaller than those for corresponding non-rotated perturbations. As for the non-rotated perturbations, the amplification over 24 hour intervals due to nonlinear evolution is never as rapid as suggested by the 24h SV first singular value, and it exceeds the growth given by the scaled 48h SV first singular value only for the first 24 hours.

If this behaviour is typical of all 25 perturbation pairs, then the increased non-linearity of the 24h SV perturbations compared to that of the 48h SV perturbations (seen as higher values of Θ and lower values of correlation) is not due to the 24h SV perturbations growing more rapidly, but rather is caused by the 24h SV perturbations growing *more nonlinearly* than their 48h SV counterparts. Results calculated using the total energy norm for the fifth, tenth and twentieth SV twin pairs do show similar nonlinear amplification factors between 48h and 24h SV perturbations. (It should be noted that for the total energy norm results, the 48h and 24h SV perturbations used are of similar initial amplitude, in contrast to those used for the z500 norm results in which the 24h SV perturbations are larger than those constructed using 48h SV.)

The results so far in this section provide a thorough comparison of 48h and 24h SV perturbations initialised on January 1, 1998; it is not suggested that results for one case are conclusive, but the computational power required to investigate multiple cases in this detail lies beyond that available for this thesis. The comparison of linearity results, calculated in the z500 norm, for 48h and 24h SV for several days is feasible, and the results are presented below.

6.7.3 Evaluation of linearity measures, averaged over 7 days

Two sets of rotated SV perturbations, with optimisation times of 48 and 24 hours, were generated for each of 6 further days. It was mentioned above that the similarity between the two SV total energy fields was about 40% for January 1, 1998. Table 6.1 gives the similarity indices between 48h and 24h SV fields for each of the seven days considered; average similarity is again about 40%. In order to put this in context,

Initialisation date 1200 on d.i.98	Similarity index between		
	48h SV	24h SV	48h & 24h SV
1.i.98			36.7
2.i.98	10.7	7.5	39.6
3.i.98	8.8	6.7	34.7
4.i.98	13.6	8.8	35.7
5.i.98	14.8	9.5	37.7
6.i.98	15.0	9.5	41.9
7.i.98	8.8	7.5	42.2
Average	12.0	8.2	38.4

Table 6.1: Similarity indices (as defined in section 5.2.1) between SV on consecutive days for both 48h SV and 24h SV, and between 48h and 24h SV for the same day, for 7 days 1.i.98-7.i.98. Average values over the 7 days are also given.

the similarity between SV fields on consecutive days is also shown in the table. That for 48h SV is on average 12.0%, while the similarity between 24h SV on consecutive days is 8.2% on average; the 48h and 24h SV for the same day are much more similar than 48h (or 24h) SV on consecutive days.

The ensembles were evolved under the T_7159 model and linearity results calculated in the 500hPa geopotential height norm are shown in figure 6.28. The results from

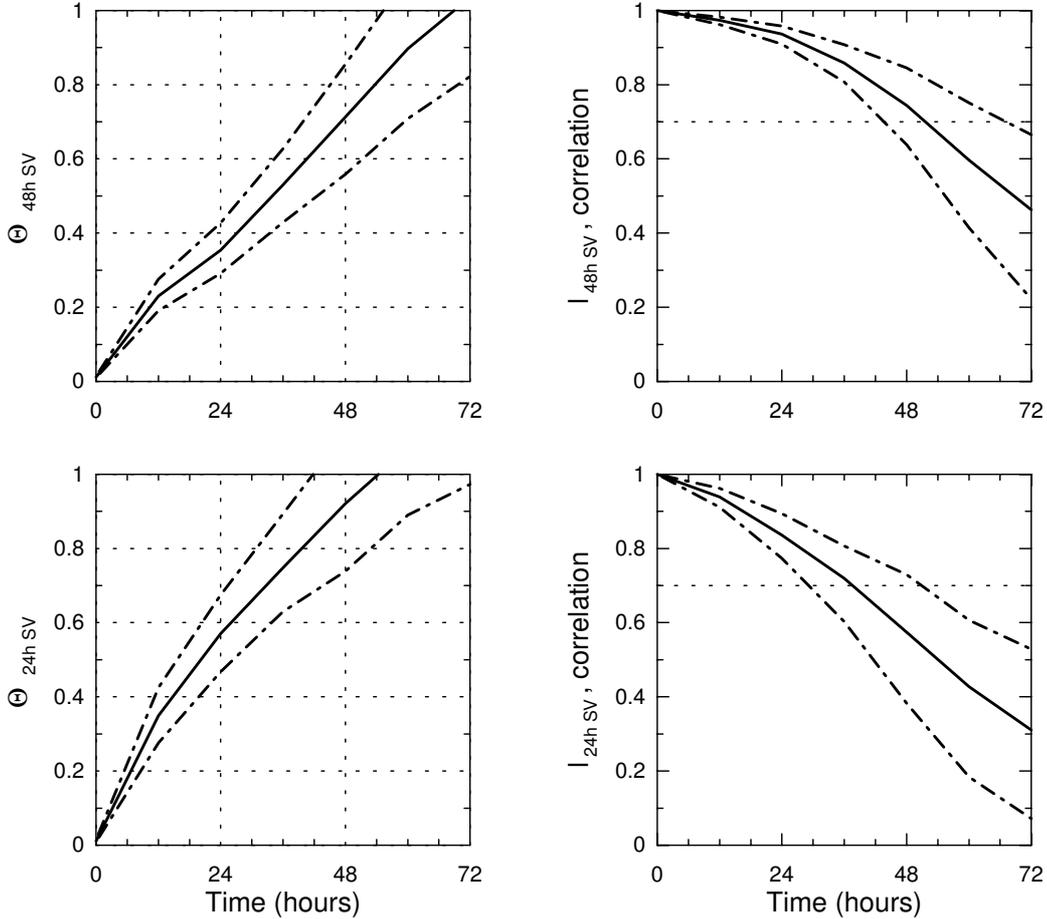


Figure 6.28: Linearity results for ECMWF twin (rotated) 48h SV (top) and increased amplitude 24h SV (bottom) perturbations, calculated in the 500hPa geopotential height norm over the Northern Hemisphere and averaged over 7 days, 1.i.98-7.i.98. The panels are as in figure 6.26.

January 1, 1998 again seem to be typical, and comparing the 48h SV results averaged over the week January 1-7, 1998, they are similar to those averaged over 25 days and presented in the previous section (see figure 6.19). The relative nonlinearity of the 48h SV at the optimisation time of 48 hours is ~ 0.7 , as is the correlation. From the results for the Marzec-Spiegel system and the annulus it appeared that changing the optimisation time had little effect on the linearity results. Had that been the case here, the relative nonlinearity for 24h SV perturbations at the optimisation time of 24 hours would have been less than ~ 0.4 , with a correlation greater than 0.9; such anticipated values would suggest that the assumptions of linear growth of the perturbations until the optimisation time were reasonable and that 24h SV are internally

consistent. However, unlike the lower dimensional models considered, changing the optimisation time of SV perturbations in the ECMWF numerical weather model does affect the linearity results. The average error in assuming linearity at the optimisation time of 24 hours is $\sim 55\%$ of the average evolved perturbation magnitude, with an average correlation of 0.84. While these values show that the (average) nonlinearity of 24h SV perturbations at optimisation time is definitely less than that of 48h SV perturbations, they are not sufficient to claim that reducing the optimisation time to 24 hours renders SV ensembles internally consistent. While it is not the purpose here to enter a discussion on the performance measures used to quantify the success of SV ensembles, it is of interest to note that, according to these measures, the 48h and 24h SV ensembles score almost identically for the skill and spread (or extent) measures averaged over the 7 cases considered; the 24h SV ensembles are more internally consistent, less expensive (computationally) and as successful as their 48h counterparts.

6.8 Summary

In Chapter 5 the evaluation of ensemble performance was divided into quantifying the success of the ensemble in fulfilling its aims and validating any assumptions made in defining the ensembles. This chapter addresses the evaluation of assumptions made in the definition of Ensemble Prediction Schemes used in Numerical Weather Prediction. Due to the expense of running the large numerical models, measures must use values operationally available. The use of twin (equal and opposite) pairs in numerical weather ensemble prediction schemes facilitates investigation of the assumption that CV perturbations evolve approximately linearly, for a given time, made in the definition of SV and construction of method B BV. The measure of linearity which exists in the literature [12] reflects only the orientation of the evolved perturbations and is blind to their relative magnitudes. A new measure, the relative nonlinearity (denoted as Θ), was defined so as to be sensitive to both magnitude and orientation.

Results were calculated for the Marzec-Spiegel system and for the annulus for a variety of parameter values. Variations in optimisation time, pre-breeding time and dimension of the CV subspace were found to have little effect on results, in contrast to the substantial effect of varying initial perturbation magnitude. Perturbation evolution was approximately linear for the optimisation times and cycle times for almost all of the SV and BV perturbations considered (BV perturbations of initial magni-

tude 0.1 in the annulus being the exception); the ensembles were usually found to be internally consistent with the linearity assumptions made in their definition.

Operational ECMWF (48h) SV ensembles were found to *not* be internally consistent with the assumption that perturbation evolution is well approximated by the TLM for the optimisation time (since being well approximated by linear evolution is a necessary but not sufficient condition for the TLM to provide a good approximation). Reducing the optimisation time did affect the linearity results, in contrast to the findings for the Marzec-Spiegel and annulus systems. As a result, the 24h SV ensembles constructed were more consistent, but it can not be said that perturbation is linear for 24 hours. Operational NCEP BV were also found *not* to be internally consistent. It would be of great interest to see how BV ensembles constructed using method A, which does not assume linearity, compare to the method B BV used at present. In the results presented, the breakdown of linearity was mainly due to discrepancies in the orientation of the twin perturbations, but differences in magnitudes were also found.

The SV which corresponds to the largest singular value specifies the direction in which an infinitesimal perturbation will have grown the most at the optimisation time; this first SV need not be that which linearly grows the most at any instant during the period $t \in [0, \tau_{\text{opt}})$, and being the first SV gives absolutely no information about the nonlinear evolution. Results from all of the systems show that the SV which is initially most nonlinear varies; the first SV is *not* always the direction in which the breakdown of linearity is first seen.

While doubling times showed exponentially uniform growth of SV and BV perturbations in the Marzec-Spiegel system, such behaviour was not found for perturbations evolved under the RBF model of the annulus. The 12 hour sampling interval and short doubling and quadrupling times of SV perturbations evolved under the ECMWF model prohibit conclusive results. The times for SV perturbations to increase by factors of 4 and 16 do not suggest uniform growth rates on average, neither do doubling and quadrupling times of BV perturbations. The difference in the results for the Marzec-Spiegel system and the annulus may be due to the relative simplicity of the Marzec-Spiegel system, to the complex nature of the annulus, to the model error of the RBF model of the annulus, or to a combination of these. It must be stressed that the orientation of the perturbation affects the doubling times, as discussed by *Smith et al.* [85]; the results presented here do not therefore contradict the results presented within the afore mentioned paper, which are for LV perturbations. Similarly, doubling times are sensitive to initial perturbation magnitude and model dynamics. An example of this sensitivity is the difference between doubling times

of at least 2 days found by *Lorenz* [47] for SV perturbations evolved under the 1982 operational ECMWF model and the doubling times of at most 24 hours found for current operational SV perturbations; model resolution and dynamics have changed in the meantime, as have the magnitude of analysis errors which specify perturbation amplitude. Uniformly exponential growth of SV and BV perturbations may be a phenomenon seen only in low-dimensional systems. In contrast the slower growth of BV perturbations compared to SV perturbations is seen in all systems considered here, but this result is to be expected.

From the results presented here, it is clear that the question of whether pathologies occur Even In, Or Only In (EIOOI) depends on the pathology. The independence of the linearity measures on parameters other than initial perturbation magnitude seen in low-dimensional systems is not seen in the ECMWF NWP model. In contrast it appears that results concerning the cause of the breakdown of linearity found in low-dimensional systems do ‘really apply to the atmosphere’, as do the results regarding which direction this breakdown first occurs in. It is not clear what makes some results unique to low-dimensional systems while others apply across the (dimensional) board.

Finally, in answer to the question ‘What is the distribution of linear times for current atmospheric models?’, the results presented here should dispel the belief that the evolution of perturbations used operationally is approximately linear for at least two days. If the linear time is taken as the time for which the error in assuming linearity is less than 30% of the average evolved perturbation magnitude (*i.e.* $\min t$ such that $\Theta(t) < 0.3$), and values are calculated using the 500hPa geopotential height norm over the Northern Hemisphere excluding the Tropics, then operational NCEP BV ensembles and operational ECMWF (48h) SV ensembles have linear times of less than 24 hours. In light of the results presented here ECMWF is currently considering modifying its operational ensemble.

Chapter 7

Conclusions

The aim of this thesis is to evaluate and compare the performance of models and ensemble prediction schemes. A summary of the main results are given in this chapter.

Optimal model order Information criteria are contrasted with prediction errors as a measure of model order optimality for use when exact data is known. Results for RBF models of the Moran-Ricker map (in section 4.2) suggest that, in general, the penalty applied by the (Akaike and Rissanen) information criteria for increasing the order (and hence complexity) of a model is far less than the improvement in accuracy gained; the information criteria are very similar to the one step in-sample prediction errors. If the model is constructed using noisy data, it is better to over-estimate the optimal model error than to under-estimate it.

Discrepancies of r.m.s. error measure Traditional root mean square error is limited as a model evaluation measure; when data is noisy it may reject the system in favour of an imperfect model as illustrated using models of the Ikeda map in section 3.3.1.

ι -shadowing as a model evaluation measure ι -shadowing time distributions are introduced as an alternative measure of model optimality suited to erroneous models and uncertain data. Using a model of the Ikeda map, ι -shadowing is shown to exploit model sensitivity while penalising model error (see section 3.3.2). In section 4.1 ι -shadowing is demonstrated to evaluate competing models of the logistic and Moran-Ricker maps and of the Lorenz equations, and to rank the models correctly, which may be verified in these cases since the system equations are known.

ι -shadowing to localise model error Using imperfect models of the Sinai map, the failure points (*i.e.* the points at which ι -shadowing trajectories fail to shadow) are shown to localise model error (see section 4.2). Failure points of an RBF model of the annulus are found in section 4.3, and used to construct a modified model which is shown to be an improvement (in the ι -shadowing sense).

ι -shadowing as an assimilation technique ι -shadowing is introduced as an assimilation technique in section 4.4.2; for a fixed assimilation window it will not always return an analysis value, but those returned will be consistent with the observations in the assimilation interval. (The failure of ι -shadowing to return an analysis value is likely to indicate regions of large model error.) Analysis values calculated by ι -shadowing assimilation are shown to yield better approximations to the true system states (than the observations) for the Marzec-Spiegel system in section 5.2.2; ι -shadowing assimilation is also applied to observations of the annulus in section 5.2.3.

Ensembles and the invariant manifold The importance of constraining ensemble members to lie on the invariant manifold, distributed according to the natural measure, is clearly illustrated for the Marzec-Spiegel system (see figure 5.2 on page 82). The error in the unconstrained ensembles is shown not to be accountable; specifically it is illustrated that bifurcations seen in unconstrained ensembles may be artifacts of including in the ensemble members which are not constrained to lie on the manifold, rather than characteristics of the system.

Constrained vector (CV) ensemble performance Evaluation of the CV ensembles may be considered in two parts:

1. How successful are CV ensembles at achieving the aims for which the subspaces were defined? Measures to assess the success of ensembles in achieving their aims (as described in section 5.2.1) require knowledge of (or a good approximation to) the exact system state and results are hence limited to CV ensembles constructed using a perfect model of the Marzec-Spiegel system and a RBF model of the annulus.

Do singular vector (SV) ensembles achieve their aim? The aim of SV ensembles is to predict the ‘worse case scenario’; success is measured by whether or not the SV ensembles captured the spread of the system (see section 5.2.1.1 for details).

- Using a perfect model of the Marzec-Spiegel system, the SV ensembles were found to capture the spread providing the initial perturbation magnitude was comparable to that of the analysis error (see section 5.2.2.5).
- Using a RBF model of the annulus, the SV ensembles usually over-estimate the spread for all of the initial perturbation magnitudes considered (see section 5.2.3.5).

Do breeding vector (BV) ensembles achieve their aim? The aim of BV ensembles is to include as a member the initial condition closest to the system value; success is measured both by how similar the BV and Lyapunov vector (LV) subspaces are, and by how well the ‘dream perturbation’ (the perturbation from the analysis to the system state) projects onto the BV subspace (see section 5.2.1.2 for details).

- Using a perfect model of the Marzec-Spiegel system, the BV ensemble subspaces were similar to those of the LV. The dream perturbation did not project well onto the BV subspaces (using either 1 or 2 BV), especially when compared to the projection onto the SV subspaces.
- Using a RBF model of the annulus, there was little similarity between BV and LV subspaces. Projection of the dream perturbation was found to be similar whether 1 or 3 BV were used. This was in contrast to the marked change in dream perturbation projection into the SV subspace as the number of SV used was varied; for 1 CV, projection into the BV subspace was better (on average), while for 3 CV projection into the SV subspace was better, again on average (see figures 5.19 and 5.18 on pages 114 and 113 respectively).

2. Are assumptions made in defining CV subspaces valid? Few of the assumptions are able to be verified (as described in section 5.2.1). If the exact system state is known (or a good approximation available) then the assumption that the model SV (formed about the analysis) approximate the system SV (formed about the ‘truth’) may be evaluated. The employment of twin (equal and opposite) pairs of perturbations in operational ensembles enables the relevance, or otherwise, of the linear approximation (assumed in the definition of SV ensembles and the construction of operational BV ensembles) to be assessed (details may be found in sections 5.1.3.1 and 5.1.3.4).

Do model SV approximate system SV? The relevance of the assumption is measured by the similarity between the model SV and system SV subspaces.

- Using a perfect model of the Marzec-Spiegel system, model SV are found to be very similar to those of the system, hence the approximation is relevant.
- Using a RBF model of the annulus, SV formed about the analysis are found not to be similar to those formed about the ‘truth’; the approximation is not relevant.

How relevant is the linearity approximation? For internally consistent SV ensembles the linear approximation must be relevant for the magnitude of the SV perturbations over the optimisation time. For operational (method B) BV ensembles the linear approximation must be relevant for the magnitude of the BV perturbations over the cycle time if the BV ensemble is to be consistent with assumptions made in its construction. (Method A BV do not require the linearity approximation.)

The measure of linearity which exists in the literature reflects only the orientation of the evolved perturbations and is blind to their relative magnitude. A new measure, the relative nonlinearity (defined in section 6.2), is sensitive to both magnitude and direction. Further, the threshold in the literature is shown to correspond to an error in assuming linearity which is $\sim 75\%$ of the mean magnitude of the evolved perturbations. q -pling times were also calculated to investigate whether or not perturbation growth is exponentially uniform.

Marzec-Spiegel Linearity results for the Marzec-Spiegel system suggest that both SV and BV ensembles are internally consistent with the linearity assumptions made in their definition, for the range of parameter values considered (see section 6.3). The second SV was usually more nonlinear initially. Perturbation growth was found to be approximately exponentially uniform (upto a factor of 16 for SV, 4 for BV).

Annulus For the annulus, BV ensembles were found to be internally consistent for all parameter values considered. The internal consistency of the SV ensembles depends on the optimisation time used: the linearity results did not vary much with varying optimisation time, so provided the optimisation time was short enough, the ensemble was consistent. The first SV was usually nonlinear initially, and growth of both BV and SV was found not to be exponentially uniform (see section 6.4).

Operational NWP ensembles Linearity results for operational ECMWF SV ensembles show that they are *not* internally consistent (see figure 6.14); the error in assuming linearity at the operational optimisation time of 48 hours is on average 72% of the mean magnitude of the evolved perturbations (in the z500 norm). Regions of extreme nonlinearity are found to correlate with synoptic structures (see figure 6.17). The first SV is not always the first to be nonlinear initially.

Operational NCEP BV ensembles are also found *not* to be internally consistent; the error in assuming linearity at the operational cycle time of 12 hours is on average 47% in the z500 norm, and increases to 71% if the uv500 norm is used (thereby corroborating the view that the z500 norm is the most linear). The region over which the norm is taken is found to have little effect, the only discernible difference being an increased spread when the norm is taken over the Tropics (see figure 6.18).

No conclusions may be made about the error growth as the sampling time of 12 hours is too large for reliable calculations to be made.

Comparing 48h SV and 24h SV ensembles For this thesis 24h SV ensembles were constructed using the operational ECMWF model (for details of see section 6.7). If the linearity results had been independent of optimisation time (as found for the Marzec-Spiegel system and the annulus) then the error in assuming linearity at 24h would have been $\sim 40\%$. This was not found to be the case and the error at 24h for the 24h SV is $\sim 55\%$; 24h SV ensembles are also not internally consistent.

On average the similarity between 48h SV and 24h SV subspaces for the same day was 40% (which is large compared to the similarity of consecutive day subspaces for either 48h SV or 24h SV). Calculations of the linearity measures in the total energy norm show the nonlinear growth to be more rapid than that in the z500 norm. The average nonlinearity of non-rotated perturbations was greater than that of rotated ensembles, but the spread of the results was greater for the the latter. 24h SV perturbations showed more rapid initial nonlinear growth than the corresponding 48h SV perturbations, but a 29% larger initial amplitude was needed for the 24h SV if the spread of the ensembles was to be comparable at 48h. Amplification factors were calculated (over 24h and 48h intervals) and found to be similar for both 48h SV and 24h SV perturbations; the enhanced nonlinearity

initially seen in the 24h SV perturbations is therefore not because they are growing more (in terms of magnitude), but because the growth is more nonlinear.

If a 30% error in assuming linearity is taken as a threshold of relevance then *operational NCEP BV ensemble perturbations and operational ECMWF (48h) SV ensemble perturbations evolve linearly for less than 24 hours* (calculated using in the z500 norm over the Northern Hemisphere excluding the Tropics). This answers the question regarding the distribution of linear times for current atmospheric models raised in the introduction.

Finally, the CV ensemble performance results also enable a response to the other two questions mentioned in the introduction:

- The linear approximation does not breakdown in a consistent manner. In particular, the image of the direction corresponding the largest singular value (for the operational optimisation time) is usually the first in which the linear approximation fails; usually, but not always.
- As to the question of whether pathologies occur Even In Or Only In low-dimensional systems, several answers may be deduced from the results presented in this thesis:
 1. Linearity measure results show dependence on the initial CV perturbation magnitude occurs even in low-dimensional systems.
 2. Independence of linearity results on all other parameters (*e.g.* optimisation time, cycle time) seems to occur only in low-dimensional systems. (Linearity results for the ECMWF SV ensembles are dependent on the optimisation time.)
 3. Extreme nonlinearities in localised regions as the cause of the breakdown of linearity occurs even in low-dimensional systems; results found in low-dimensional systems do ‘really apply to the atmosphere’.
 4. The inconsistency of which SV direction is the first in which the linear approximation fails occurs even in the low-dimensional systems considered.

It is not yet clear what makes some pathologies unique to low-dimensional systems while other apply across the (dimensional) board.

Appendix A

Anosov-Bowen ϵ -shadowing

In order to formally state the Anosov-Bowen ϵ -shadowing lemma [4, 7] we first introduce some notation. A system is *uniformly hyperbolic* if, at each point \mathbf{x} , the linear tangent space at that point, $\mathcal{T}_{\mathbf{x}}$, may be divided into 2 distinct subspaces [56]: a stable subspace, $E_{\mathbf{x}}^s$, and an unstable subspace, $E_{\mathbf{x}}^u$, so

$$\mathcal{T}_{\mathbf{x}} = E_{\mathbf{x}}^s \oplus E_{\mathbf{x}}^u,$$

i.e. the angle between the stable and unstable subspaces is uniformly bounded *away* from zero.

The linear tangent subspaces are defined in terms of the linear or tangent propagator $\mathcal{M}(\mathbf{x}, \Delta t)$ defined in section 2.1.3 [56]: there exist $K > 0$ and $0 < \rho < 1$ such that

$$\begin{aligned} \|\mathcal{M}(\mathbf{x}, n\tau_s) \cdot \mathbf{w}\| &\leq K\rho^n \|\mathbf{w}\| & \text{if } \mathbf{w} \in E_{\mathbf{x}}^s \\ \|\mathcal{M}(\mathbf{x}, -n\tau_s) \cdot \mathbf{w}\| &\leq K\rho^n \|\mathbf{w}\| & \text{if } \mathbf{w} \in E_{\mathbf{x}}^u \end{aligned}$$

Further, the decomposition of the tangent space must vary continuously in \mathbf{x} and the subspaces must evolve smoothly onto one another, *i.e.*

$$\begin{aligned} E_{\mathbf{x}_{k+1}}^s &= \mathcal{M}(\mathbf{x}_k, \tau_s) E_{\mathbf{x}_k}^s \\ E_{\mathbf{x}_{k+1}}^u &= \mathcal{M}(\mathbf{x}_k, \tau_s) E_{\mathbf{x}_k}^u \end{aligned}$$

Suppose we have such a uniformly hyperbolic system, with a true orbit $\mathbf{S} = \{\mathbf{s}_t\}_{t=a}^b$, where $\mathbf{f}(\mathbf{s}_t) = \mathbf{s}_{t+1}$, and a pseudo-orbit $\mathbf{P} = \{\mathbf{p}_t\}_{t=a}^b$ given by a model of the system, *e.g.* a numerical algorithm. Then \mathbf{P} is a δ *pseudo-orbit*, denoted \mathbf{P}_δ , for $a \leq t \leq b$ if

$$|\mathbf{f}(\mathbf{p}_t) - \mathbf{p}_{t+1}| \leq \delta$$

for every $a \leq t \leq b$ [42].

Given $\epsilon \geq 0$, a δ pseudo-orbit \mathbf{P}_δ is said to be ϵ -shadowed by \mathbf{x} , a point in the system, for $a \leq t \leq b$ if

$$|\mathbf{f}^t(\mathbf{x}) - \mathbf{p}_t| \leq \epsilon$$

for every $a \leq t \leq b$ [42].

A δ pseudo-orbit satisfies a set of single time step restrictions, while ϵ -shadowing involves satisfying an iterated restriction. Together they enable the formal statement of the lemma:

Lemma 1 (Anosov-Bowen Shadowing Lemma) *For a uniformly hyperbolic system, for every $\epsilon > 0$, $\exists \delta > 0$ such that each δ pseudo-orbit can be ϵ -shadowed by a point in the domain of the system.*

Appendix B

ι -shadowing: the algorithm

In order to determine the ι -shadowing time of a model for a trajectory, we need to find the ‘optimal’ initial condition which ι -shadows for the longest time. Conceivably, the simplest approach to finding the optimal initial condition is to randomly choose candidate points from the initial neighbourhood of observation, Ξ_0 , and iterate them to find their ι -shadowing time until all points in Ξ_0 have been evaluated. This is impractical due to the infinite set (or numerically finite, but computationally prohibitively large, set) of possible initial conditions. A more viable method assumes that the regions within Ξ_0 with the same ι -shadowing times are simply connected. For a 1d system it is then possible to consider successive ensembles, systematically distributed over nested intervals, each centred on the member of the previous ensemble which has the longest ι -shadowing time. In higher dimensions the size of ensemble required and time taken to evolve all of the ensemble members soon makes this methodical approach infeasible. A generalisation of the approach involving ensembles randomly distributed within nested hyper-cubes is described below.

Alternatively, we may seek to use the information of how a model trajectory fails to ι -shadow, feeding it back to give a suggested perturbation in the initial condition through the linear propagator. This assumes that the evolution of such a perturbation is well approximated by linearised dynamics for the duration of the ι -shadowing time. If such perturbations lie outside regions well described by linear dynamics then the linear propagator feedback method may fail, and we revert to the random distribution of trial initial conditions. The extent of these ‘linear’ regions may be determined using ‘twin experiments’ (as discussed in detail in Chapter 6) where equal and opposite perturbations ϵ and $-\epsilon$ are added to the initial condition and iterated for the ι -shadowing time. The magnitude of the perturbations, $\|\epsilon\|$, measured in some prescribed norm, is decreased until the iterated perturbed values agree within a given tolerance, *e.g.* 10%. In order to determine the region which has dynamics well

described by a linear approximation for all choices of direction of ϵ , the perturbations tested must be in the directions most rapidly growing for the chosen magnitude.

Neither of the algorithms described are optimal; they are intended to be relatively simple and functional. In practice a combination of the two is employed, the linear feedback method is used while the linearity assumption is a reasonable approximation and the random search method used otherwise.

B.1 Method of interval convergence

An ensemble, $\{x_0^k\}, k = 1, \dots, n_e$, is evenly distributed over the 1d interval $[y_0 - \zeta, y_0 + \zeta]$, and each member is evolved forwards, under the model, to find its ι -shadowing time, τ_ι^k . Often there are a range of ensemble members $[x_0^{\kappa_1}, x_0^{\kappa_2}]$, $\kappa_1 < \kappa_2$, each with ι -shadowing time of τ_ι^κ ; the new ensemble $\{x_0'^k\}$ is constructed to cover the interval $[x_0^{\max(\kappa_1-1, 0)}, x_0^{\min(\kappa_2+1, n_e)}]$. The ensemble size n_e is chosen to be fairly small initially. If greater resolution is needed due to a large interval of x with equal ι -shadowing times, the ensemble size n_e is doubled up to a prescribed maximum; specifically, if $|x_0'^{\kappa_1} - x_0'^{\kappa_2}| > 0.6|x_0^0 - x_0^{n_e}|$. Successive ensembles are evolved and refined until either the entire observed trajectory is ι -shadowed, or until a specified number of refinements have been made.

B.2 Random search method

The idea of converging intervals generalises to converging hyper-cubes in higher dimensions. While the computational cost of uniform ensembles increases as a power of the dimension of the space, the process remains tractable if ensembles are randomly, rather than uniformly, distributed throughout a given region. In practice the following two pass algorithm is implemented:

A hyper-cuboid, Θ , is defined by a centre, \mathbf{x}_0^0 , and the distances from the centre to the surfaces of the hyper-cuboid, $(\boldsymbol{\eta})_l, l = 1, \dots, m$. The centre is initially taken to be the initial observation \mathbf{y}_0 , and $\boldsymbol{\eta}$ is taken as the bound of the observational uncertainty, $\boldsymbol{\zeta}$; initially $\Theta = \Xi_0$, the neighbourhood of observation. A second region $\Theta_1 \subset \Theta$ is then defined, with deviation $\frac{1}{\xi} \boldsymbol{\eta}$ where ξ is a given constant, *e.g.* 1.005. The ensemble $\{\mathbf{x}_0^k\}, k = 1, \dots, 2n_e$ is formed by randomly distributing n_e points in each of Θ ($k = 1, \dots, n_e$) and Θ_1 ($k = n_e + 1, \dots, 2n_e$). The first n_e members are then evolved to find their ι -shadowing times τ_ι^k , and κ found such that

$$\tau_\iota^\kappa > \tau_\iota^i \quad \forall i \in [1, n_e].$$

Given that the observed initial condition ι -shadows for a time τ_l^0 , we test to see if

$$\tau_l^k > \alpha \tau_l^0$$

for given $\alpha > 1$, typically $\alpha = 2$; if so a substantial improvement has been found and the new centre \mathbf{x}'_0 is taken to be \mathbf{x}_0^k . If this condition is not satisfied, then each remaining ensemble member, \mathbf{x}_0^k , $k = n_e + 1, \dots, 2n_e$, is evolved *and* tested in turn, either until a member satisfies

$$\tau_l^k > \alpha \tau_l^0, \quad (\text{B.1})$$

or until all members have been tested, with τ_l^k updated to remain the maximum ι -shadowing time of evolved members. If equation B.1 is not satisfied by any member of the ensemble then we test for the less stringent condition

$$\tau_l^k > \beta \tau_l^0$$

for given β , where $\alpha > \beta > 1$, and, if this is satisfied, take \mathbf{x}_0^k as the new centre; otherwise \mathbf{x}_0^0 is unchanged.

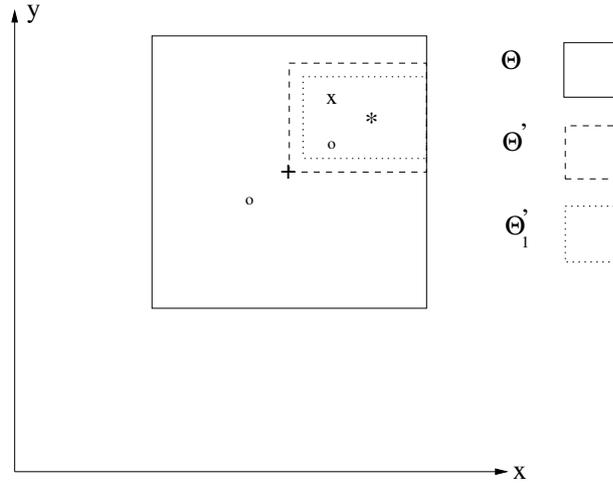


Figure B.1: Given an initial observation (+) and neighbourhood of observation, Θ , an ensemble is created and a new trial initial condition, \mathbf{x}'_0 (*), found. The new ensemble consists of two points (o) determined by the position of \mathbf{x}'_0 relative to \mathbf{x}_0 , and other points randomly distributed within the regions Θ' and Θ'_1 . The system value (x) is also shown.

Having chosen a new centre, the new region Θ' is defined to have sides with a distance from the centre given by

$$(\eta')_l = \max \left\{ \frac{(\eta)_l}{(\xi)_l}, |(\mathbf{y}_0 - \mathbf{x}'_0)_l| \right\}, l = 1, \dots, m,$$

so long as this is contained within the observational neighbourhood (see figure B.1). Inasmuch as we assume the ι -shadowing times to be simply connected, we wish to have a higher density of candidate points in the nearby vicinity of the new centre; a new region Θ'_1 with sides a distance $\frac{1}{\xi^4}\boldsymbol{\eta}$, $\xi > 1$ from the centre is created as described above. The new ensemble is constructed with the inclusion of two specified points if the centre has moved, namely

$$\mathbf{x}'_0 + \lambda(\mathbf{x}'_0 - \mathbf{x}_0), \quad \lambda = \frac{1}{2} \text{ and } \begin{cases} \frac{3}{2} & \text{if resulting point lies within } \Xi_0 \\ \frac{1}{4} & \text{otherwise} \end{cases} .$$

The procedure then continues as previously and successive evolutions and refinements are made until either the ι -shadowing time is the duration of the observed trajectory or until a given maximum number of refinements have been made.

B.3 Localised linear feedback

This approach assumes that the evolution of the suggested perturbation is well approximated by the linear propagator for the duration of the ι -shadowing time. As above, we consider a trial initial condition, \mathbf{x}_0 , originally taken to be the initial observation, \mathbf{y}_0 , and evolve it forwards until it fails to ι -shadow at some time $\tau_\iota = \nu\tau_s$, where τ_s is the data sampling time. We then seek to use the information on how the pseudo-trajectory fails to ι -shadow to determine a small perturbation to the initial condition which will increase the ι -shadowing time. Given the model $\tilde{\mathbf{f}}$, so $\mathbf{x}_i = \tilde{\mathbf{f}}^i(\mathbf{x}_0)$, with observations \mathbf{y}_i at times $t = i\tau_s$, and *unknown* system values \mathbf{s}_i which satisfy $\mathbf{s}_i = \mathbf{f}^i(\mathbf{s}_0)$, and assuming that the domain of the system is contained in the range of the model, we define

$$\boldsymbol{\Delta}_i = \mathbf{s}_i - \tilde{\mathbf{f}}^i(\mathbf{x}_0) \tag{B.2}$$

$$\boldsymbol{\gamma}_i = \mathbf{s}_i - \mathbf{y}_i \tag{B.3}$$

$$\boldsymbol{\epsilon}_i = \mathbf{y}_i - \tilde{\mathbf{f}}^i(\mathbf{x}_0) \tag{B.4}$$

as the *model error*, *error in observation* and *measurable error* respectively (see figure B.2 A).

The magnitude of each component of the error in observation, is bounded by the magnitude of the corresponding observational uncertainty distribution, *i.e.* $|(\boldsymbol{\gamma})_l| \leq |(\boldsymbol{\zeta})_l|, l = 1, \dots, m$. As long as the model ι -shadows the system then the model and measurable errors are also bounded by the observation error. The failure of a model

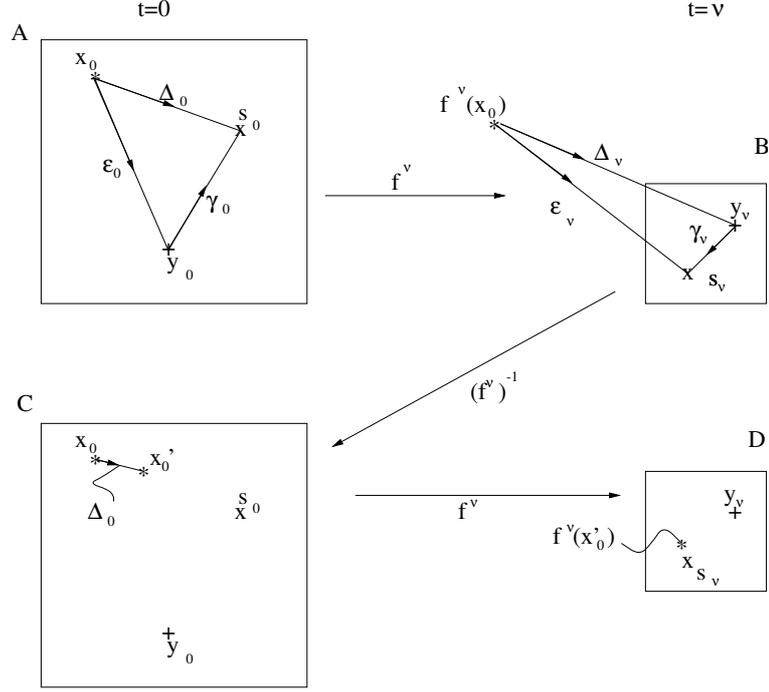


Figure B.2: Illustration of how error may be fed back to increase the ι -shadowing time of the trial initial condition. A: Given a trial initial condition \mathbf{x}_0 (*), initial system value \mathbf{s}_0 (x) and initial observation \mathbf{y}_0 (+), it is possible to define Δ_0 , γ_0 , ϵ_0 . All values lie within the neighbourhood of observation (solid region). B: The trial initial condition is found to fail to ι -shadow at time $t = \nu\tau_s$, with $\mathbf{f}^\nu(\mathbf{x}_0)$ lying far outside the neighbourhood of observation. C: Approximating ϵ_ν by Δ_ν , this information is fed back to give a suggested perturbation Δ_0 to the initial condition. D: Iterating this new trial initial condition, \mathbf{x}'_0 , ν times we find that the iterated value $\mathbf{f}^\nu(\mathbf{x}'_0)$ lies close to the system value \mathbf{s}_0 ; the new trial initial condition has a longer ι -shadowing time.

trajectory to ι -shadow, at time $\nu\tau_s$ say, is often due to it passing through a region of rapid error growth, hence $\tilde{\mathbf{f}}^\nu(\mathbf{x}_0)$ will lie far outside the observational neighbourhood and we may assume that

$$|(\epsilon)_l|, |(\Delta)_l| \gg |(\gamma)_l| \quad \forall l \in [1, m] \quad (\text{B.5})$$

(see figure B.2 B). Combining equations B.2, B.3 and B.4 gives

$$\epsilon_\nu = \Delta_\nu - \gamma_\nu$$

which, with equation B.5 gives the approximation

$$\Delta_\nu \simeq \epsilon_\nu. \quad (\text{B.6})$$

Given the measurable error, ϵ_ν , we wish to estimate the model error, Δ_0 , in the trial initial condition. If this is known, then we may perturb the initial condition to

give a new trial initial condition, $\mathbf{x}'_0 = \mathbf{x}_0 + \Delta_0$, which will satisfy

$$\tilde{\mathbf{f}}^\nu(\mathbf{x}'_0) \simeq \mathbf{y}_\nu$$

(see figure B.2 C and D), thereby increasing the ν -shadowing time, provided the trajectory still ν -shadows the earlier observations (*i.e.* in $t \in [0, \nu - 1]$). Following the new trial trajectory until it fails, the procedure is repeated and the trial initial condition will tend towards the optimal initial condition for that model; if the model is in fact the system used to generate the data, then the exact (system) initial condition will be an optimal initial condition.

To approximate Δ_0 , we will express ϵ_ν in terms of Δ_0 , and then invert the equation. In order to find an expression for ϵ_ν we appeal to stability analysis, as described in section 2.1.3, and must consider models of maps and flows separately [24].

B.3.1 Maps

Given a model $\tilde{\mathbf{f}}$ describing a map and taken to be a good approximation to \mathbf{f} , $\mathbf{x}_{t+1} = \tilde{\mathbf{f}}(\mathbf{x}_t)$ and $\mathbf{s}_{t+1} = \mathbf{f}(\mathbf{s}_t)$. Substituting into equation B.2,

$$\begin{aligned} \Delta_{t+1} &= \mathbf{s}_{t+1} - \mathbf{x}_{t+1} \\ &= \mathbf{f}(\mathbf{s}_t) - \tilde{\mathbf{f}}(\mathbf{x}_t) \\ &= \mathbf{f}(\mathbf{x}_t + \Delta_t) - \tilde{\mathbf{f}}(\mathbf{x}_t) \\ &= \mathbf{f}(\mathbf{x}_t) + \Delta_t \mathcal{J}(\mathbf{x}_t) + \frac{1}{2} \Delta_t^T \mathcal{H}(\mathbf{x}_t) \Delta_t + \mathcal{O}(|\Delta_t|^3) - \tilde{\mathbf{f}}(\mathbf{x}_t) \\ &\simeq \Delta_t \mathcal{J}(\mathbf{x}_t) + \frac{1}{2} \Delta_t^T \mathcal{H}(\mathbf{x}_t) \Delta_t + \mathcal{O}(|\Delta_t|^3). \end{aligned} \tag{B.7}$$

If the system equations are known, the Jacobian \mathcal{J} and Hessian \mathcal{H} can be calculated analytically, otherwise they can be approximated from the observational data. In the following we will assume the Jacobian is given.

Taking the perturbation, Δ_0 , in the initial condition to be sufficiently small (so that its evolution is approximately linear), we truncate equation B.7 to first order; applying repeatedly gives

$$\Delta_t = \left\{ \prod_{i=0}^{t-1} \mathcal{J}(\mathbf{x}_i) \right\} \Delta_0.$$

From equation 2.8 it follows that

$$\mathcal{M}(\mathbf{x}_0, t) = \left\{ \prod_{i=0}^{t-1} \mathcal{J}(\mathbf{x}_i) \right\}$$

and hence

$$\boldsymbol{\epsilon}_\nu \simeq \boldsymbol{\Delta}_\nu \equiv \mathcal{M}(\mathbf{x}_0, \nu) \boldsymbol{\Delta}_0,$$

giving an expression for $\boldsymbol{\epsilon}_\nu$ in terms of $\boldsymbol{\Delta}_0$ as required.

B.3.2 Flows

Implementing similar analysis for flows, we have $\frac{d\mathbf{x}}{dt} = \tilde{\mathbf{f}}(\mathbf{x})$, $\frac{d\mathbf{s}}{dt} = \mathbf{f}(\mathbf{s})$, so

$$\begin{aligned} \mathbf{s}_{t+\delta t} &= \mathbf{s}_t + \mathbf{f}(\mathbf{s}_t)\delta t + \mathcal{O}(\delta t^2) \\ \mathbf{x}_{t+\delta t} &= \mathbf{x}_t + \tilde{\mathbf{f}}(\mathbf{x}_t)\delta t + \mathcal{O}(\delta t^2) \end{aligned}$$

where $\delta t = \tau_s$ for discrete observational data. Substituting into equation B.2,

$$\begin{aligned} \boldsymbol{\Delta}_{t+\delta t} - \boldsymbol{\Delta}_t &= \mathbf{s}_{t+\delta t} - \mathbf{s}_t - \mathbf{x}_{t+\delta t} + \mathbf{x}_t \\ &= \mathbf{f}(\mathbf{s}_t)\delta t - \tilde{\mathbf{f}}(\mathbf{x}_t)\delta t + \mathcal{O}(\delta t^2). \end{aligned}$$

Again from equation B.2

$$\begin{aligned} \mathbf{f}(\mathbf{s}_t) &= \mathbf{f}(\mathbf{x}_t + \boldsymbol{\Delta}_t) \\ &\simeq \tilde{\mathbf{f}}(\mathbf{x}_t) + \boldsymbol{\Delta}_t \mathcal{J}(\mathbf{x}_t) + \frac{1}{2} \boldsymbol{\Delta}_t^2 \mathcal{H}(\mathbf{x}_t) + \mathcal{O}(\boldsymbol{\Delta}_t^3). \end{aligned}$$

Substituting into equation B.2 gives

$$\begin{aligned} \boldsymbol{\Delta}_{t+\delta t} - \boldsymbol{\Delta}_t &\simeq \boldsymbol{\Delta}_t \mathcal{J}(\mathbf{x}_t)\delta t + \frac{1}{2} \boldsymbol{\Delta}_t^2 \mathcal{H}(\mathbf{x}_t)\delta t + \mathcal{O}(\boldsymbol{\Delta}_t^3, \delta t^2) \\ \Rightarrow \boldsymbol{\Delta}_{t+\delta t} &\simeq \left[1 + \mathcal{J}(\mathbf{x}_t)\delta t + \frac{1}{2} \boldsymbol{\Delta}_t \mathcal{H}(\mathbf{x}_t)\delta t \right] \boldsymbol{\Delta}_t + \mathcal{O}(\boldsymbol{\Delta}_t^3, \delta t^2). \end{aligned} \quad (\text{B.8})$$

Again, taking the perturbation, $\boldsymbol{\Delta}_0$, in the initial condition to be sufficiently small, so that its evolution is approximately linear, and with *the additional assumption* that the time step is sufficiently small, we truncate equation B.8 to first order and apply repeatedly to obtain

$$\boldsymbol{\epsilon}_\nu \simeq \boldsymbol{\Delta}_\nu = \mathcal{M}(\mathbf{x}_0, \nu) \boldsymbol{\Delta}_0.$$

where now

$$\mathcal{M}(\mathbf{x}_0, t) = \left\{ \prod_{i=0}^{t-1} [1 + \mathcal{J}(\mathbf{x}_i)\delta t] \right\}.$$

B.3.3 Inversion

Having obtained an expression for ϵ_ν in terms of Δ_0 , we now need to invert the equation

$$\epsilon_\nu \simeq \mathcal{M}(\mathbf{x}_0, \nu) \Delta_0 \quad (\text{B.9})$$

to find Δ_0 . Since \mathcal{M} may be an almost singular matrix we use singular value decomposition to express \mathcal{M}_ν in the form

$$\mathcal{M}_\nu = U_\nu \Sigma_\nu V_\nu^T \quad (\text{B.10})$$

as discussed in Appendix C. The singular vectors, given by the rows of V_ν^T , will be the directions in which higher order terms first become important over time ν ; perturbations in these directions will have most effect on the evolution of the trajectory at time ν . Again, this assumes that the evolution of the particular initial condition is well approximated by the linearised equations, and nonlinear effects can be neglected. The singular values and vectors are calculated for a given time of evolution and will vary in evolution; they need to be calculated afresh for each value of ν .

Taking $E_\nu = U_\nu^T \epsilon_\nu$, we then ‘screen’ the singular values, applying a threshold below which they are set to zero (as done when forming a Moore-Penrose pseudo-inverse as described in section C.1 and [31]). Denoting this by

$$(\mathcal{K}_\nu)_l = \begin{cases} \frac{1}{\sigma_l} (E_\nu)_l & \sigma_l > \sigma_{\min} \\ 0 & \sigma_l \leq \sigma_{\min} \end{cases}, l = 1, \dots, m,$$

where σ_{\min} is the threshold value.

If we take $\Delta_0 = V_\nu \mathcal{K}_\nu$, a weighted combination of eigenvectors, then

$$\begin{aligned} \mathcal{M}_\nu \Delta_0 &= U_\nu \Sigma_\nu V_\nu^T V_\nu \mathcal{K}_\nu \\ &\simeq U \Sigma V^T V \Sigma^{-1} E_\nu \\ &\simeq U \Sigma V^T V \Sigma^{-1} U^T \epsilon_\nu \\ &\simeq \epsilon_\nu \end{aligned}$$

as required, the approximation arising from the use of a pseudo-inverse. If the l^{th} eigenvalue is below the threshold, *i.e.* $\sigma_l < \sigma_{\min}$, then $(\mathcal{M}_\nu \Delta_0)_l = 0$, and no perturbation occurs in the direction of the l^{th} eigenvector.

Given this approximation of Δ_0 we perturb the initial condition to give a new trial initial condition

$$\mathbf{x}'_0 = \mathbf{x}_0 + \Delta_0.$$

Providing that the the perturbation Δ_0 lies within the linear region, the new initial condition \mathbf{x}'_0 will yield a trajectory whose value, at time ν , will be contained in the corresponding neighbourhood of observation. Even in this case the trajectory of \mathbf{x}'_0 may fail to ι -shadow for some $t \in [0, \nu - 1]$; if such a failure occurs the change is discarded and we retain \mathbf{x}_0 as the trial initial condition. It should be noted that perturbations in the ‘null’ directions, for which $\sigma_l \leq \sigma_{\min}$, may be able to be used to prevent losing ι -shadowability at such earlier times. In practice the method of localised linear feedback is implemented in conjunction with the random search method, which is used when the suggested perturbation has evolution which fails for some reason (*e.g.* the perturbation lies outside the extent of the ‘linear’ region).

Appendix C

Singular Value Decomposition and its applications

Singular value decomposition, SVD, of a matrix or operator is a useful tool, introduced to nonlinear dynamical systems analysis by *Broomhead & King* [10]. It is a method of matrix decomposition whose power lies in the properties of the matrices formed, as illustrated by its application in singular systems analysis described here. Further applications (to the linear propagator) are used in Chapter 5 and are described therein.

C.1 SVD: a matrix decomposition

Considering *singular value decomposition* as a method of matrix decomposition, it is helpful in solving the matrix equation

$$A\mathbf{x} = \mathbf{b}$$

when A is singular or ill-conditioned. Any $m \times n$ matrix A may be decomposed into

$$A = U\Sigma V^T$$

where U is $m \times n$ orthogonal matrix, V is a $n \times n$ orthogonal matrix and Σ is an $n \times n$ diagonal matrix with non-negative elements σ_i [87]. The decomposition is unique up to: (i) permutation of columns of U and V , and elements of Σ , and (ii) linear combinations of any columns of U and V whose corresponding elements of Σ are exactly equal; elements of Σ are usually ordered so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$.

The columns of U (V) are known as the left (right) *singular vectors*, with corresponding *singular values* σ_i . The left singular vectors corresponding to non-zero singular values form an orthonormal basis for the range of A , while the right singular

vectors corresponding to zero singular values form an orthonormal basis for the null space of A .

When $\sigma_i > 0 \forall i$, \mathbf{x} can be determined using SVD since

$$\begin{aligned}\mathbf{x} &= V\Sigma^{-1}U^T\mathbf{b} \\ &= V[\text{diag}(1/\sigma_i)]U^T\mathbf{b}.\end{aligned}$$

If A is singular then there is no unique solution \mathbf{x} of equation C.1; singularity implies that there is at least one zero singular value. Even in this case, we may formulate the *Moore-Penrose pseudo-inverse* of Σ [31] by

$$\text{replacing } 1/\sigma_i \text{ by zero if } \sigma_i = 0. \quad (\text{C.1})$$

Using this pseudo-inverse a unique solution is obtained which minimises $|\mathbf{x}|^2$ over the possible set of solutions, providing the equation has a solution. If \mathbf{b} is not in the range of A , then the \mathbf{x} which minimises $|A\mathbf{x} - \mathbf{b}|$, *i.e.* the \mathbf{x} which is mapped closest to \mathbf{b} in the least squares sense, is found.

Implementation of SVD numerically requires adjusting the substitution in C.1 to

$$\text{replacing } 1/\sigma_i \text{ by zero if } \sigma_i < \sigma_{\min}, \quad (\text{C.2})$$

since in practice some σ_i may be very small ¹, and it is often preferable to discard the linear combinations corresponding to these singular values as they may be irretrievably contaminated by observational and/or numerical noise.

C.2 Singular systems analysis

Broomhead & King [10] introduced the method of singular value decomposition in order to calculate the optimal basis onto which data from a time series could be projected in order to reconstruct the attractor of the system. This method of *singular systems analysis* (SSA) was developed to deal with noise, unavoidable in observational data. The technique involves forming a $M \times n$ trajectory matrix A , whose rows are n -dimensional delay vectors, formed using the method of delays, see section 2.3.1, thus

$$A = \frac{1}{\sqrt{M}} \begin{bmatrix} \hat{\mathbf{Y}}_1^T \\ \hat{\mathbf{Y}}_2^T \\ \vdots \\ \hat{\mathbf{Y}}_n^T \end{bmatrix}.$$

¹Due to computational error.

SVD may be applied directly to A , though often it is more convenient to apply SVD to the $n \times n$ covariance matrix $C = A^T A$, yielding a set of right singular vectors \mathbf{v}_i and values σ_i which solve the eigenvalue problem

$$C\mathbf{v}_i = \sigma_i^2 \mathbf{v}_i;$$

the right singular vectors are normalised eigenvectors and σ_i^2 the eigenvalues. Considering the embedded time series as a cloud, or mass, of points, the singular vectors are the principle moments of the mass, known as *Empirical Orthogonal Functions* or EOFs, while the singular values give weightings. Hence SSA gives the optimal projection of the data onto a phase portrait, in the least squares sense of variance.

Any observational noise will be present in all singular values, and the complete basis will be needed to fully describe the distribution of data. However, it is possible to project the data onto a truncated basis, corresponding to the most significant singular values, in the hope that the most interesting or important dynamics of the system are restricted to this low-dimensional phase space. This will concentrate the variance of the data, *but not necessarily the information contained in the data*. Choice of the truncated basis is often achieved by taking the threshold, σ_{\min} from equation C.2, as the value of the noise floor, the limiting value to which the singular values decay. For a review see *Allen & Smith* [2] and references thereof. (SSA is not central to this thesis as an analysis technique, but is relevant since it provides an alternative basis for the reconstruction of dynamics. The techniques introduced in this thesis and illustrated for delay reconstruction models could equally be applied to models built upon other bases including the EOFs of SSA.)

Bibliography

- [1] H. AKAIKE. A new look at statistical model identification. *IEEE Trans. Automat. Contr. AC*, **19**:716–722, 1974.
- [2] M. R. ALLEN AND L. A. SMITH. Monte Carlo SSA: Detecting irregular oscillations in the presence of coloured noise. *J. Climate*, **9**:3373–3404, 1996.
- [3] K.T ALLIGOOD, T.D. SAUER, AND J.A. YORKE. *Chaos - an introduction to dynamical systems*. Springer-Verlag, New York, 1997.
- [4] D.V. ANOSOV. Geodesic flows and closed Riemannian manifolds with negative curvature. *Proc. Steklov Inst. Math.*, **90**, 1967.
- [5] V.I. ARNOLD AND A. AVEZ. *Problèmes Ergodiques de la Mécanique Classique*. Gauthier–Villars, Paris, 1967.
- [6] J. BARKMEIJER. Constructing fast-growing perturbations for the nonlinear regime. *J. Atmos. Sci.*, **53(19)**:2838–2851, 1996.
- [7] R. BOWEN. ω -limit sets for axiom A diffeomorphisms. *J. Diff. Eqns.*, **18**:333–339, 1975.
- [8] L. BRILLOUIN. *Scientific Uncertainty, and Information*. Academic Press, New York, 1964.
- [9] D.S. BROOMHEAD, R. INDIK, A.C. NEWELL, AND D.A. RAND. Local adaptive Galerkin bases for large-dimensional dynamic-systems. *Nonlinearity*, **4(2)**:159–197, 1991.
- [10] D.S. BROOMHEAD AND G.P. KING. Extracting qualitative dynamics from experimental data. *Physica D*, **20**:217–236, 1986.
- [11] D.S. BROOMHEAD AND D. LOWE. Multivariable functional interpolation and adaptive networks. *Compl. Syst.*, **2**:321–355, 1988.

- [12] R. BUIZZA. Optimal perturbation time evolution and sensitivity of ensemble prediction to perturbation amplitude. *Q. J. Roy. Meteor. Soc.*, **121**:1705–1738, 1995.
- [13] R. BUIZZA AND T.N. PALMER. The singular-vector structure of the atmospheric general circulation. *J. Atmos. Sci.*, **52(9)**:1434–1456, 1995.
- [14] R. BUIZZA, T. PETROLIAGIS, T.N. PALMER, J. BARKMEIJER, M. HAMRUD, A. HOLLINGSWORTH, A. SIMMONS, AND N. WEDI. Impact of model resolution and ensemble size on the performance of an Ensemble Prediction System. *Q. J. Roy. Meteor. Soc.*, **124(550B)**:1935–1960, 1998.
- [15] M. CASDAGLI. Nonlinear prediction of chaotic time series. *Physica D*, **35**:335–356, 1989.
- [16] E.M. COVEN, I. KAN, AND J.A. YORKE. Pseudo-orbit shadowing in the family of tent maps. *Trans. Amer. Math. Soc.*, **308(1)**:227–241, 1988.
- [17] A.G. DARBYSHIRE AND D.S. BROOMHEAD. Robust estimation of tangent maps and Liapunov spectra. *Physica D*, **89**:287–305, 1996.
- [18] D. DRAPER. Assessment and propagation of model uncertainty (with discussion). *J. Roy. Stat. Soc. B*, **57**:45–97, 1995.
- [19] P.G. DRAZIN. *Nonlinear systems*. Cambridge University Press, Cambridge, 1992.
- [20] J.A. DUTTON. *Dynamics of Atmospheric Motion*. Dover, New York, 1995.
- [21] J.P. ECKMANN AND D. RUELLE. Ergodic theory of chaos. *Rev. Mod. Phys.*, **45**:617–656, 1985.
- [22] *Predictability(1997)*, ECMWF, 1998. Shinfield Park, Reading, Berks., U.K. Proceedings of the ECMWF workshop on Predictability, 20–22 October, 1997.
- [23] R.M. ERRICO, T. VUKIĆEVIĆ, AND K. RAEDER. Examination of the accuracy of a tangent linear model. *Tellus A*, **45**:462–477, 1993.
- [24] S. EUBANK AND J.D. FARMER. An introduction to chaos and randomness. In E. JEN, editor, *Lectures in Complex Systems*, **Lecture II** of *SFI Studies in the Sciences of Complexity*. Addison-Wesley, 1990.

- [25] J.D. FARMER, E. OTT, AND J.A. YORKE. The dimension of chaotic attractors. *Physica D*, **7**:153–180, 1983.
- [26] J.D. FARMER AND J.J. SIDOROWICH. Exploiting chaos to predict the future and reduce noise. In Y.C. LEE, editor, *Evolution, Learning and Cognition*. World Scientific Press, 1988.
- [27] J.D. FARMER AND J.J. SIDOROWICH. Optimal shadowing and noise reduction. *Physica D*, **47**:373–392, 1991.
- [28] A.M. FRASER AND H.L. SWINNEY. Independent coordinates for strange attractors from mutual information. *Phys. Rev. A*, **33**:1134–1140, 1986.
- [29] W.-G. FRÜH AND P.L. READ. Wave interactions and the transition to chaos of baroclinic waves in a thermally driven rotating annulus. *Phil. Trans. Roy. Soc. Lond. A*, **355(1722)**:101–153, 1997.
- [30] J. GLEICK. *Chaos:making a new science*. Viking, 1987.
- [31] G.H. GOLUB AND C.F. VANLOAN. *Matrix Computations*. John Hopkins University Press, Baltimore, 1989.
- [32] C. GREBOGI, S.M. HAMMEL, J.A. YORKE, AND T. SAUER. Shadowing of physical trajectories in chaotic dynamics: containment and refinement. *Phys. Rev. Letts.*, **65(13)**:1527–1530, 1990.
- [33] G.J. HALTINER AND R.T. WILLIAMS. *Numerical prediction and dynamic meteorology*. John Wiley & Sons, 1990.
- [34] S. HAMMEL, C.K.R.T. JONES, AND J. MALONEY. Global dynamical behaviour of the optical field in a ring cavity. *J. Opt. Soc. Am. B*, **2**:552–564, 1985.
- [35] M. HÉNON. A two-dimensional mapping with a strange attractor. *Commun. Math. Phys.*, **50**:69–77, 1976.
- [36] R. HIDE AND P.J. MASON. Sloping convection in a rotating fluid. *Adv. Phys.*, **24**:47–100, 1975.
- [37] P.L. HOUTEKAMER AND J. DEROME. Prediction experiments with two-member ensembles. *Monthly Weather Review*, **122(9)**:2179–2191, 1994.

- [38] P.L. HOUTEKAMER, L. LEFAIVRE, J. DEROME, H. RITCHIE, AND H. MITCHELL. A system simulation approach to ensemble prediction. *Mon. Wea. Rev.*, **124(6)**:1225–1242, 1996.
- [39] K. IKEDA. Multiple valued stationary state and its instability of the transmitted light by a ring cavity system. *Opt. Commun.*, **30**:257–261, 1979.
- [40] I.N. JAMES. *Introduction to circulating atmospheres*. CUP, Cambridge, 1994.
- [41] H. KANTZ AND T. SCHREIBER. *Nonlinear time series analysis*. CUP, Cambridge, U.K., 1997.
- [42] M. KUCHTA. Shadowing property of continuous maps with zero topological entropy. *Proc. Am. Math. Soc.*, **119**:641–648, 1993.
- [43] J.F. LACARRA AND O. TALAGRAND. Short-range evolution of small perturbations in a barotropic model. *Tellus A*, **40**:81–95, 1988.
- [44] T.Y. LI AND J.A. YORKE. Period three implies chaos. *American Mathematical Monthly*, **82**:985–992, 1975.
- [45] E.N. LORENZ. Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**:130–141, 1963.
- [46] E.N. LORENZ. A study of the predictability of a 28-variable atmospheric model. *Tellus*, **17(3)**:321–333, 1965.
- [47] E.N. LORENZ. Atmospheric predictability experiments with a large numerical model. *Tellus*, **34**:505–513, 1982.
- [48] C.J. MARZEC AND E.A. SPIEGEL. Ordinary differential equations with strange attractors. *SIAM J. Appl. Math.*, **38(3)**:403–421, 1980.
- [49] R. MAY. Simple mathematical models with very complicated dynamics. *Nature*, **261**:459–467, 1976.
- [50] F. MOLteni, R. BUZZA, T.N. PALMER, AND T. PETROLIAGIS. The ECMWF ensemble prediction system: methodology and validation. *Q. J. R. Meteorol. Soc.*, **122**:73–120, 1996.
- [51] P.A.P. MORAN. Some remarks on animal population dynamics. *Biometrics*, **6**:250–258, 1950.

- [52] T. MULLIN. *The Nature of Chaos*. OUP, Oxford, 1993.
- [53] R. MUREAU, F. MOLTENI, AND T.N. PALMER. Ensemble prediction using dynamically conditioned perturbations. *Q. J. R. Meteorol. Soc.*, **119**:299–323, 1993.
- [54] J.D. MURRAY. *Mathematical Biology*. Springer-Verlag, 1993.
- [55] V.I. OSELEDEC. A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems. *Trans. Mosc. Math. Soc.*, **19**:197–231, 1968.
- [56] E. OTT. *Chaos in dynamical systems*. CUP, Cambridge, 1993.
- [57] E. OTT, T. SAUER, AND J.A. YORKE, editors. *Coping with Chaos*. Wiley, 1994.
- [58] N.H. PACKARD, J.P. CRUTCHFIELD, J.D. FARMER, AND R.S. SHAW. Geometry from a time series. *Phys. Rev. Letts.*, **45**:712–716, 1990.
- [59] T.N. PALMER, R. BUIZZA, AND F. LALAURETTE. Performance of the ECMWF prediction scheme. In ECMWF [22]. 12 pages.
- [60] T.N. PALMER, R. BUIZZA, F. MOLTENI, Y.Q. CHEN, AND S. CORTI. Singular vectors and the predictability of weather and climate. *Phil. Trans. Roy. Soc. Lond. A*, **348(1688)**:459–475, 1994.
- [61] T.N. PALMER, R. GELARO, J. BARKMEIJER, AND R. BUIZZA. Singular vectors, metrics, and adaptive observations. *J. Atmos. Sci.*, **55(4)**:633–653, 1998.
- [62] F. PAPARELLA, A. PROVENZALE, L.A. SMITH, C. TARICCO, AND R. VIO. Local random analogue prediction of nonlinear processes. *Phys. Letts. A*, **235(3)**:233–240, 1997.
- [63] S.K. PARK AND K.K. DROEGEMEIER. Validity of the tangent linear approximation in a moist convective cloud model. *Mon. Wea. Rev.*, **125**:3320–3340, 1997.
- [64] T.S. PARKER AND L.O. CHUA. *Practical Numerical Algorithms for Chaotic Systems*. Springer-Verlag, New York, 1989.

- [65] W.H. PRESS, B.P. FLANNERY, S.A. TEUKOLSKY, AND W.T. VETTERLING. *Numerical Recipes*. CUP, Cambridge, 1986.
- [66] G.D. QUINLAN AND S. TREMAINE. On the reliability of gravitational N-body integrations. *Mon. Not. R. Astron. Soc.*, **259**:505–518, 1992.
- [67] F. RABIER AND P. COURTIER. Four-dimensional assimilation in the presence of baroclinic instability. *Q. J. Roy. Meteor. Soc.*, **118**:649–672, 1992.
- [68] P. READ, M. J. BELL, D. W. JOHNSON, AND R. M. SMALL. Quasi-periodic and chaotic flow regimes in a thermally driven, rotating fluid annulus. *J. Fluid Mech.*, **238**:599–632, 1992.
- [69] W.E. RICKER. *J. Fisheries Res. Board Can.*, **11**:559–623, 1954.
- [70] J. RISSANEN. Consistent order estimates of autoregressive processes by shortest description of data. In O. JACOBS ET AL., editor, *Analysis and optimisation of stochastic systems*, New York, 1980. Academic Press.
- [71] F.J. ROMEIRAS, C. GREBOGI, E. OTT, AND W.P. DAYAWANSA. Controlling chaotic dynamical systems. *Physica D*, **58**:165–192, 1992.
- [72] C. RUHLA. *The Physics of Chance*. OUP, Oxford, 1992.
- [73] B. SALTZMAN. Finite amplitude free convection as an initial value problem - I. *J. Atmos. Sci.*, **19**:329–341, 1962.
- [74] A.N. SARKOVSKII. Coexistence of cycles of a continuous map of a line onto itself. *Ukrainian Math. J.*, **16**:61–71, 1964.
- [75] T. SAUER, J.A. YORKE, AND M. CASDAGLI. Embedology. *J. Stat. Phys.*, **65**:579–616, 1991.
- [76] J. SINAI. Gibbs measure in ergodic theory. *Russ. Math. Surveys*, **4**:21–64, 1972.
- [77] L. A. SMITH. Lacunarity and chaos in nature. Columbia University, New York, NY, 1987. Ph. D. Thesis, Appendix 1.
- [78] L. A. SMITH. Identification and prediction of low dimensional dynamics. *Physica D*, **58**:50–76, 1992.

- [79] L. A. SMITH. Does a meeting in Santa Fe imply chaos? In A.S WEIGEND AND N.A. GERSHENFIELD, editors, *Time series prediction: Forecasting the future and understanding the past*, **XV** of *SFI studies in the sciences of complexity*, pages 323–343. Addison-Wesley, 1993.
- [80] L. A. SMITH. Local optimal prediction: exploiting strangeness and the variation of sensitivity to initial condition. *Phil. Trans. R. Soc. Lond. A*, **348**:371–381, 1994.
- [81] L. A. SMITH, I. GILMOUR, AND P.L. READ. A brief overview of the seminar on atmospheric predictability. OCIAM Technical Report, OCIAM, 24–29 St. Giles’, Oxford, OX1 3LB, U.K. 1996.
- [82] L.A. SMITH. Accountability and error in ensemble prediction of baroclinic flows. In *Predictability(1995)*, **1**, pages 351–368, ECMWF, 1995. Shinfield Park, Reading, Berks., U.K. Proceedings of a seminar held at ECMWF, 4–8 September, 1995.
- [83] L.A. SMITH, 1999. Private communication.
- [84] L.A. SMITH AND I. GILMOUR. Accountability and internal consistency in ensemble formation. In ECMWF [22]. 15 pages.
- [85] L.A. SMITH, C. ZIEHMANN, AND K. FRAEDRICH. The limits of predictability. *Q. J. Roy. Meteor. Soc.*, 1998. 31 pages. Submitted (April 1997).
- [86] C. SPARROW. *The Lorenz equations: bifurcations, chaos, and strange attractors*. Springer–Verlag, New York, 1982.
- [87] G.W. STEWART. *Introduction to Matrix Computations*. Academic Press, New York, 1973.
- [88] G. STRANG. *Introduction to applied mathematics*. Wellesley-Cambridge press, 1986.
- [89] R. STROE AND J.F. ROYER. Comparison of different error growth formulas and predicability estimation in numerical extended-range forecasts. *Ann. Geophys.*, **11**:296–316, 1993.
- [90] S.H. STROGATZ. *Nonlinear dynamics and chaos*. Addison-Wesley, 1994.

- [91] F. TAKENS. Detecting strange attractors in fluid turbulence. In D. RAND AND L.-S. YOUNG, editors, *Dynamical Systems and Turbulence*, pages 366–381, Berlin, 1981. Springer-Verlag.
- [92] O. TALAGRAND AND P. COURTIER. Variational assimilation of meteorological observations with the adjoint vorticity equation - part I: theory. *Q.J.R. Meteorol. Soc.*, **113**:1311–1328, 1987.
- [93] Z. TOTH, 1997. Private communication.
- [94] Z. TOTH AND E. KALNAY. Ensemble forecasting at NMC: the generation of perturbations. *Bull. Am. Meteorol. Soc.*, **74(12)**:2317–2330, 1993.
- [95] Z. TOTH, E. KALNAY, S.M. TRACTON, R. WOBUS, AND J. IRWIN. A synoptic evaluation of the NCEP ensemble. *Wea. Forecasting*, **12(1)**:140–153, 1997.
- [96] H.M. VAN DEN DOOL. Searching for analogues, how long must we wait? *Tellus A*, **46**:314–324, 1994.
- [97] E.S. VAN VLECK. Numerical shadowing near hyperbolic trajectories. *SIAM*, **16**:1177–1189, 1995.
- [98] S. VANNITSEM AND C. NICOLIS. Lyapunov vectors and error growth patterns in a T21L3 quasigeostrophic model. *J. Atmos. Sci.*, **54**:347–361, 1997.
- [99] T. VUKIĆEVIĆ. Nonlinear and linear evolution of initial forecast errors. *Mon. Wea. Rev.*, **119**:1602–1611, 1991.
- [100] D.S. WILKS. *Statistical methods in the atmospheric sciences*. Academic Press, San Diego, California, 1995.
- [101] D.E. WILLMES. Computer assisted shadowing of nonhyperbolic maps. *Annals of the NY Academy of Sciences*, **751**:213–221, 1995.
- [102] A. WOLF, J.B. SWIFT, H.L. SWINNEY, AND J.A. VASTANO. Determining Lyapunov exponents from a time series. *Physica D*, **16**:285–317, 1985.