

Arnaud Doucet
Nando de Freitas
Neil Gordon
Editors

Sequential Monte Carlo Methods in Practice

Foreword by Adrian Smith

With 168 Illustrations



Contents

Foreword	v
Acknowledgments	vii
Contributors	xxi
I Introduction	1
1 An Introduction to Sequential Monte Carlo Methods	3
Arnaud Doucet, Nando de Freitas, and Neil Gordon	
1.1 Motivation	3
1.2 Problem statement	5
1.3 Monte Carlo methods	6
1.3.1 Perfect Monte Carlo sampling	7
1.3.2 Importance sampling	8
1.3.3 The Bootstrap filter	10
1.4 Discussion	13
II Theoretical Issues	15
2 Particle Filters - A Theoretical Perspective	17
Dan Crisan	
2.1 Introduction	17
2.2 Notation and terminology	17
2.2.1 Markov chains and transition kernels	18
2.2.2 The filtering problem	19
2.2.3 Convergence of measure-valued random variables .	20
2.3 Convergence theorems	21
2.3.1 The fixed observation case	21
2.3.2 The random observation case	24
2.4 Examples of particle filters	25
2.4.1 Description of the particle filters	25

2.4.2	Branching mechanisms	28
2.4.3	Convergence of the algorithm	31
2.5	Discussion	33
2.6	Appendix	35
2.6.1	Conditional probabilities and conditional expectations	35
2.6.2	The recurrence formula for the conditional distribution of the signal	38
3	Interacting Particle Filtering With Discrete Observations	43
Pierre Del Moral and Jean Jacod		
3.1	Introduction	43
3.2	Nonlinear filtering: general facts	46
3.3	An interacting particle system under Case A	48
3.3.1	Subcase A1	48
3.3.2	Subcase A2	55
3.4	An interacting particle system under Case B	60
3.4.1	Subcase B1	60
3.4.2	Subcase B2	67
3.5	Discretely observed stochastic differential equations	71
3.5.1	Case A	72
3.5.2	Case B	73
III	Strategies for Improving Sequential Monte Carlo Methods	77
4	Sequential Monte Carlo Methods for Optimal Filtering	79
Christophe Andrieu, Arnaud Doucet, and Elena Punskaya		
4.1	Introduction	79
4.2	Bayesian filtering and sequential estimation	79
4.2.1	Dynamic modelling and Bayesian filtering	79
4.2.2	Alternative dynamic models	80
4.3	Sequential Monte Carlo Methods	82
4.3.1	Methodology	82
4.3.2	A generic algorithm	85
4.3.3	Convergence results	86
4.4	Application to digital communications	88
4.4.1	Model specification and estimation objectives	89
4.4.2	SMC applied to demodulation	91
4.4.3	Simulations	93

5 Deterministic and Stochastic Particle Filters in State-Space Models	97
Erik Bølviken and Geir Storvik	
5.1 Introduction	97
5.2 General issues	98
5.2.1 Model and exact filter	98
5.2.2 Particle filters	99
5.2.3 Gaussian quadrature	100
5.2.4 Quadrature filters	101
5.2.5 Numerical error	102
5.2.6 A small illustrative example	104
5.3 Case studies from ecology	104
5.3.1 Problem area and models	104
5.3.2 Quadrature filters in practice	107
5.3.3 Numerical experiments	110
5.4 Concluding remarks	112
5.5 Appendix: Derivation of numerical errors	114
6 RESAMPLE-MOVE Filtering with Cross-Model Jumps	117
Carlo Berzuini and Walter Gilks	
6.1 Introduction	117
6.2 Problem statement	118
6.3 The RESAMPLE-MOVE algorithm	119
6.4 Comments	124
6.5 Central limit theorem	125
6.6 Dealing with model uncertainty	126
6.7 Illustrative application	129
6.7.1 Applying RESAMPLE-MOVE	131
6.7.2 Simulation experiment	134
6.7.3 Uncertainty about the type of target	135
6.8 Conclusions	138
7 Improvement Strategies for Monte Carlo Particle Filters	139
Simon Godsill and Tim Clapp	
7.1 Introduction	139
7.2 General sequential importance sampling	140
7.3 Markov chain moves	143
7.3.1 The use of bridging densities with MCMC moves .	144
7.4 Simulation example: TVAR model in noise	145
7.4.1 Particle filter algorithms for TVAR models	146
7.4.2 Bootstrap (SIR) filter	148
7.4.3 Auxiliary particle filter (APF)	149
7.4.4 MCMC resampling	150
7.4.5 Simulation results	152
7.5 Summary	157

7.6	Acknowledgements	158
8	Approximating and Maximising the Likelihood for a General State-Space Model	159
	Markus Hürzeler and Hans R. Künsch	
8.1	Introduction	159
8.2	Bayesian methods	159
8.3	Pointwise Monte Carlo approximation of the likelihood . .	161
8.3.1	Examples	161
8.4	Approximation of the likelihood function based on filter samples	164
8.5	Approximations based on smoother samples	166
8.5.1	Approximation of the likelihood function	167
8.5.2	Stochastic EM-algorithm	167
8.6	Comparison of the methods	168
8.6.1	AR(1) process	168
8.6.2	Nonlinear example, 3 parameters	171
8.6.3	Nonlinear model, 5 parameters	173
8.6.4	Discussion	173
8.7	Recursive estimation	173
9	Monte Carlo Smoothing and Self-Organising State-Space Model	177
	Genshiro Kitagawa and Seisho Sato	
9.1	Introduction	177
9.2	General state-space model and state estimation	178
9.2.1	The model and the state estimation problem	178
9.2.2	Non-Gaussian filter and smoother	179
9.3	Monte Carlo filter and smoother	180
9.3.1	Approximation of non-Gaussian distributions	180
9.3.2	Monte Carlo filtering	181
9.3.3	Derivation of the Monte Carlo filter	182
9.3.4	Monte Carlo smoothing	183
9.3.5	Non-Gaussian smoothing for the stochastic volatility model	186
9.3.6	Nonlinear Smoothing	188
9.4	Self-organising state-space models	189
9.4.1	Likelihood of the model and parameter estimation .	189
9.4.2	Self-organising state-space model	191
9.5	Examples	192
9.5.1	Self-organising smoothing for the stochastic volatility model	192
9.5.2	Time series with trend and stochastic volatility . .	194
9.6	Conclusion	195

10 Combined Parameter and State Estimation in Simulation-Based Filtering	197
Jane Liu and Mike West	
10.1 Introduction and historical perspective	197
10.2 General framework	199
10.2.1 Dynamic model and analysis perspective	199
10.2.2 Filtering for states	200
10.2.3 Filtering for states and parameters	202
10.3 The treatment of model parameters	202
10.3.1 Artificial evolution of parameters	202
10.3.2 Kernel smoothing of parameters	203
10.3.3 Reinterpreting artificial parameter evolutions	204
10.4 A general algorithm	206
10.5 Factor stochastic volatility modelling	208
10.6 Discussion and future directions	217
11 A Theoretical Framework for Sequential Importance Sampling with Resampling	225
Jun S. Liu, Rong Chen, and Tanya Logvinenko	
11.1 Introduction	225
11.2 Sequential importance sampling principle	227
11.2.1 Properly weighted sample	227
11.2.2 Sequential build-up	228
11.3 Operations for enhancing SIS	229
11.3.1 Reweighting, resampling and reallocation	230
11.3.2 Rejection control and partial rejection control	231
11.3.3 Marginalisation	234
11.4 Monte Carlo filter for state-space models	234
11.4.1 The general state-space model	235
11.4.2 Conditional dynamic linear model and the mixture Kalman filter	236
11.5 Some examples	237
11.5.1 A simple illustration	237
11.5.2 Target tracking with MKF	239
11.6 Discussion	241
11.7 Acknowledgements	242
12 Improving Regularised Particle Filters	247
Christian Musso, Nadia Oudjane, and Francois Le Gland	
12.1 Introduction	247
12.2 Particle filters	249
12.2.1 The (classical) interacting particle filter (IPF)	250
12.2.2 Regularised particle filters (RPF)	251
12.3 Progressive correction	255
12.3.1 Focus on the correction step	256

12.3.2 Principle of progressive correction	257
12.3.3 Adaptive choice of the decomposition	258
12.4 The local rejection regularised particle filter (L2RPF)	260
12.4.1 Description of the filter	260
12.4.2 Computing the coefficient $c_t^{(i)}(\alpha_t)$	263
12.5 Applications to tracking problems	264
12.5.1 Range and bearing	265
12.5.2 Bearings-only	266
12.5.3 Multiple model particle filter (MMPF)	269
13 Auxiliary Variable Based Particle Filters	273
Michael K. Pitt and Neil Shephard	
13.1 Introduction	273
13.2 Particle filters	274
13.2.1 The definition of particle filters	274
13.2.2 Sampling the empirical prediction density	274
13.2.3 Weaknesses of particle filters	276
13.3 Auxiliary variable	277
13.3.1 The basics	277
13.3.2 A generic SIR based auxiliary proposal	278
13.3.3 Examples of adaption	283
13.4 Fixed-lag filtering	288
13.5 Reduced random sampling	289
13.5.1 Basic ideas	289
13.5.2 Simple outlier example	290
13.6 Conclusion	292
13.7 Acknowledgements	293
14 Improved Particle Filters and Smoothing	295
Photis Stavropoulos and D.M. Titterington	
14.1 Introduction	295
14.2 The methods	296
14.2.1 The smooth bootstrap	296
14.2.2 Adaptive importance sampling	300
14.2.3 The kernel sampler of Hürzeler and Künsch	302
14.2.4 Partially smooth bootstrap	303
14.2.5 Roughening and sample augmentation	305
14.2.6 Application of the methods in particle filtering and smoothing	306
14.3 Application of smooth bootstrap procedures to a simple control problem	308
14.3.1 Description of the problem	308
14.3.2 An approach to the continuous-time version of the problem	309
14.3.3 An adaptation of Titterington's method	310

14.3.4 Probabilistic criterion 1	310
14.3.5 Probabilistic criterion 2: working directly with the cost	311
14.3.6 Unknown variances	311
14.3.7 Resampling implementation	312
14.3.8 Simulation results	314
14.3.9 Further work on this problem	317
IV Applications	319
15 Posterior Cramér-Rao Bounds for Sequential Estimation	321
Niclas Bergman	
15.1 Introduction	321
15.2 Review of the posterior Cramér-Rao bound	322
15.3 Bounds for sequential estimation	323
15.3.1 Estimation model	324
15.3.2 Posterior Cramér-Rao bound	325
15.3.3 Relative Monte Carlo evaluation	327
15.4 Example – terrain navigation	329
15.5 Conclusions	338
16 Statistical Models of Visual Shape and Motion	339
Andrew Blake, Michael Isard, and John MacCormick	
16.1 Introduction	339
16.2 Statistical modelling of shape	341
16.3 Statistical modelling of image observations	343
16.4 Sampling methods	345
16.5 Modelling dynamics	346
16.6 Learning dynamics	349
16.7 Particle filtering	352
16.8 Dynamics with discrete states	354
16.9 Conclusions	355
17 Sequential Monte Carlo Methods for Neural Networks	359
N de Freitas, C Andrieu, P Højen-Sørensen, M Niranjan, and A Gee	
17.1 Introduction	359
17.2 Model specification	360
17.2.1 MLP models for regression and classification	360
17.2.2 Variable dimension RBF models	362
17.3 Estimation objectives	365
17.4 General SMC algorithm	366
17.4.1 Importance sampling step	367
17.4.2 Selection step	368

17.4.3	MCMC Step	369
17.4.4	Exact step	371
17.5	On-line classification	371
17.5.1	Simple classification example	372
17.5.2	An application to fault detection in marine diesel engines	373
17.6	An application to financial time series	375
17.7	Conclusions	379
18	Sequential Estimation of Signals under Model Uncertainty	381
Petar M. Djurić		
18.1	Introduction	381
18.2	The problem of parameter estimation under uncertainty	383
18.3	Sequential updating of the solution	384
18.4	Sequential algorithm for computing the solution	389
18.4.1	A Sequential-Importance-Resampling scheme	390
18.4.2	Sequential sampling scheme based on mixtures	395
18.5	Example	397
18.6	Conclusions	400
18.7	Acknowledgment	400
19	Particle Filters for Mobile Robot Localization	401
Dieter Fox, Sebastian Thrun, Wolfram Burgard, and Frank Dellaert		
19.1	Introduction	401
19.2	Monte Carlo localization	403
19.2.1	Bayes filtering	403
19.2.2	Models of robot motion and perception	404
19.2.3	Implementation as particle filters	405
19.2.4	Robot results	408
19.2.5	Comparison to grid-based localization	410
19.3	MCL with mixture proposal distributions	414
19.3.1	The need for better sampling	414
19.3.2	An alternative proposal distribution	416
19.3.3	The mixture proposal distribution	419
19.3.4	Robot results	420
19.4	Multi-robot MCL	423
19.4.1	Basic considerations	423
19.4.2	Robot results	425
19.5	Conclusion	426
20	Self-Organizing Time Series Model	429
Tomoyuki Higuchi		
20.1	Introduction	429

20.1.1	Generalised state-space model	429
20.1.2	Monte Carlo filter	430
20.2	Self-organizing time series model	432
20.2.1	Genetic algorithm filter	432
20.2.2	Self-organizing state-space model	434
20.3	Resampling scheme for filtering	435
20.3.1	Selection scheme	435
20.3.2	Comparison of performance: simulation study	436
20.4	Application	438
20.4.1	Time-varying frequency wave in small count data	438
20.4.2	Self-organizing state-space model for time-varying frequency wave	439
20.4.3	Results	440
20.5	Conclusions	444
21	Sampling in Factored Dynamic Systems	445
Daphne Koller and Uri Lerner		
21.1	Introduction	445
21.2	Structured probabilistic models	448
21.2.1	Bayesian networks	448
21.2.2	Hybrid networks	449
21.2.3	Dynamic Bayesian networks	451
21.3	Particle filtering for DBNs	454
21.4	Experimental results	457
21.5	Conclusions	464
22	In-Situ Ellipsometry Solutions Using Sequential Monte Carlo	465
Alan D. Marrs		
22.1	Introduction	465
22.2	Application background	465
22.3	State-space model	467
22.3.1	Ellipsometry measurement model	468
22.3.2	System evolution model	471
22.3.3	Particle filter	472
22.4	Results	474
22.5	Conclusion	475
22.6	Acknowledgments	477
23	Manoeuvring Target Tracking Using a Multiple-Model Bootstrap Filter	479
Shaun McGinnity and George W. Irwin		
23.1	Introduction	479
23.2	Optimal multiple-model solution	481
23.3	The IMM algorithm	483

23.4	Multiple model bootstrap filter	484
23.4.1	Example	486
23.5	Target tracking examples	488
23.5.1	Target scenarios	488
23.5.2	Linear, Gaussian tests	488
23.5.3	Polar simulation results	492
23.6	Conclusions	495
23.7	Acknowledgments	496
24	Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks	499
Kevin Murphy and Stuart Russell		
24.1	Introduction	499
24.2	RBPF in general	500
24.2.1	How do we choose which nodes to sample?	503
24.3	The RBPF algorithm in detail	506
24.4	Application: concurrent localisation and map learning for a mobile robot	508
24.4.1	Results on a one-dimensional grid	511
24.4.2	Results on a two-dimensional grid	514
24.5	Conclusions and future work	515
25	Particles and Mixtures for Tracking and Guidance	517
David Salmond and Neil Gordon		
25.1	Introduction	517
25.1.1	Guidance as a stochastic control problem	518
25.1.2	Information state	519
25.1.3	Dynamic programming and the dual effect	520
25.1.4	Separability and certainty equivalence	521
25.1.5	Sub-optimal strategies	522
25.2	Derivation of control laws from particles	523
25.2.1	Certainty equivalence control	523
25.2.2	A scheme based on open-loop feedback control	524
25.3	Guidance in the presence of intermittent spurious objects and clutter	525
25.3.1	Introduction	525
25.3.2	Problem formulation	525
25.3.3	Simulation example	526
25.3.4	Guidance results	528
26	Monte Carlo Techniques for Automated Target Recognition	533
Anuj Srivastava, Aaron D. Lanterman, Ulf Grenander, Marc Loizeaux, and Michael I. Miller		
26.1	Introduction	533

26.1.1	The Bayesian posterior	535
26.1.2	Inference engines	536
26.2	Jump-diffusion sampling	539
26.2.1	Diffusion Processes	540
26.2.2	Jump processes	541
26.2.3	Jump-diffusion algorithm	544
26.3	Sensor models	545
26.4	Experiments	547
26.5	Acknowledgments	552
Bibliography		553
Index		577