
Contents

How to use this book

xi

Part 1: From Bayesian inference to Bayesian workflow

1

1 Bayesian theory and Bayesian practice

3

1.1 Why Bayes?

3

1.2 Varieties of Bayesian theory

7

1.3 There is no safe haven

9

1.4 Convincing evidence

13

1.5 Exercises

14

2 Statistical modeling and workflow

17

2.1 From Bayesian inference to Bayesian data analysis to Bayesian workflow

17

2.2 “Workflow” and its relation to statistical theory and practice

19

2.3 Four modeling scenarios

22

2.4 Psychological struggles

23

2.5 Exercises

24

3 Computational tools

25

3.1 Statistical programming environments

25

3.2 Software for Bayesian inference (probabilistic programming)

26

3.3 Additional workflow tools

26

3.4 Notation in equations and code

27

3.5 A simple example of probabilistic programming

28

3.6 Exercises

34

4 Introduction to workflow: Modeling performance on a multiple choice exam

37

4.1 Fitting a simple model to data

37

4.2 Expanding the model

39

4.3 Starting over with a fully generative model

49

4.4 Performing a simulation experiment to understand the model

54

4.5 Breaking the model

57

4.6 General lessons from the multiple-choice testing example

59

4.7 Exercises

60

Part 2: Statistical workflow

61

5 Building statistical models

63

5.1 Choosing an initial model

64

5.2 Relating a statistical model to subject-matter assumptions

66

5.3 Expressing a Bayesian model using probability distributions

69

5.4 A data model is not just a “likelihood”

70

5.5 Generative and partially generative models

72

5.6	Prior distribution	75
5.7	Specifying the data model and specifying the prior	86
5.8	Modeled and unmodeled data	89
5.9	Prior predictive checking	89
5.10	Tail behavior and how the prior and likelihood combine	93
5.11	Exercises	97
6	Using simulations to capture uncertainty and experiment with models	103
6.1	Simulation to express modeling and inference	103
6.2	Point estimates and uncertainties	106
6.3	Designing and performing simulated-data experiments	108
6.4	Simulating an underlying process, data collection, and inference	112
6.5	Exercises	114
7	Prediction, generalization, and causal inference	119
7.1	Poststratification: Inferences about old averages and new scenarios	119
7.2	Causal inference as generalization	125
7.3	From inference to decision	130
7.4	Different perspectives on statistical modeling and prediction	134
7.5	Exercises	134
8	Visualizing and checking fitted models	137
8.1	Displaying high-dimensional inference	137
8.2	Posterior predictive checking	142
8.3	Cross validation checking	147
8.4	Assessing the influence of individual data points and subsets of the data	148
8.5	Influence of likelihood and prior information	150
8.6	Exercises	154
9	Comparing and improving models	157
9.1	Big data need big models	157
9.2	A topology of models	159
9.3	Visualizing models in relation to each other	160
9.4	Model selection using predictive checking and predictive performance	162
9.5	Model selection and overfitting	167
9.6	Stacking and predictive model averaging	169
9.7	Model expansion: Predictive consistency and coherence	170
9.8	Exercises	171
10	Statistical inference and scientific inference	175
10.1	Science as iteration between modeling, data collection, and data analysis	176
10.2	A virtuous tangle	179
10.3	Software assisted workflow	181
10.4	The replication crisis and its connection to multiple levels of variation	181
10.5	Simulated-data experimentation as virtual replication	183
10.6	Exercises	189
Part 3: Computational workflow		191

11 Fitting statistical models	193
11.1 The typical set and the distribution of the log posterior density	193
11.2 Initial values	195
11.3 Adaptation and warmup	196
11.4 How many parallel chains and how many iterations	197
11.5 Effective sample size and Monte Carlo standard error	200
11.6 How many digits to report based on posterior uncertainty	201
11.7 Replicability of stochastic computation	205
11.8 Exercises	205
12 Diagnosing and fixing problems with fitting	209
12.1 Fit fast, fail fast	209
12.2 Discovering computational challenges by experimenting	210
12.3 Failure modes and steps forward	212
12.4 Using modeling ideas to address computing problems	226
12.5 What to do about convergence problems	234
12.6 Exercises	235
13 Approximate algorithms and approximate models	239
13.1 Approximations based on joint and conditional posterior modes	240
13.2 Variational inference and Pathfinder	241
13.3 Simulation-based and amortized inference	242
13.4 Divide-and-conquer algorithms	243
13.5 Fitting simpler models for computational purposes	244
13.6 Application of workflow ideas to approximate computation	245
13.7 Exercises	246
14 Simulation-based calibration checking	249
14.1 The process of simulation-based calibration checking	249
14.2 Visual and numerical calibration diagnostics	251
14.3 Incorporating calibration checking into workflow	252
14.4 Exercises	254
15 Statistical modeling as software development	255
15.1 Writing code for single or multiple uses	255
15.2 Version control smooths collaborations with others and with your past self	257
15.3 Testing as you go	257
15.4 Making it essentially reproducible	258
15.5 Making it readable and maintainable	259
15.6 Exercises	260
Part 4: Case studies	261
16 Coding a series of models: Simulated data of movie ratings	263
16.1 Model for two movies	263
16.2 Extending the model to J movies	265
16.3 Item-response model with parameters for raters and for movies	267
16.4 General lessons from the movie ratings example	272
16.5 Exercises	272

17 Prior specification for regression models: Reanalysis of a sleep study	275
17.1 Priors for linear models	276
17.2 Priors for linear multilevel models	280
17.3 Priors for generalized linear multilevel models	285
17.4 Model checking	288
17.5 General lessons from the prior-specification example	291
17.6 Exercises	292
18 Predictive model checking and comparison: Clinical trial	293
18.1 Initial models	295
18.2 Model extension and sensitivity analysis	295
18.3 Inference for the treatment effect	298
18.4 Sensitivity to prior and data model	300
18.5 Comparing predictive performance	302
18.6 General lessons from the clinical trial example	306
18.7 Exercises	308
19 Building up to a hierarchical model: Coronavirus testing	305
19.1 Background	305
19.2 Modeling a test with uncertain sensitivity and specificity	306
19.3 Hierarchical model for varying testing conditions	309
19.4 Prior sensitivity analysis	311
19.5 Extensions of the model	313
19.6 General lessons from the coronavirus testing example	315
19.7 Exercises	316
20 Using a fitted model for decision analysis: Classification competition	317
20.1 A decision problem	317
20.2 Data exploration	317
20.3 Fitting a mixture model and computing predicted probabilities	319
20.4 Solving the decision problem	320
20.5 General lessons from the time series competition example	321
20.6 Exercises	322
21 Posterior predictive checking: Stochastic learning in dogs	323
21.1 An early example of simulation-based model checking	323
21.2 Fitting a series of models	325
21.3 Further model checking	331
21.4 Summarizing the inferences from a fitted hierarchical model	334
21.5 Simulated-data experimentation	334
21.6 General lessons from the dogs example	338
21.7 Exercises	338
22 Incremental development and testing: Black cat adoptions	341
22.1 The sample and the inferential goal	341
22.2 Synthetic cat adoptions	342
22.3 Modeling censored data	343
22.4 Discrete outcomes instead of durations	349
22.5 General lessons from the cat adoptions example	350
22.6 Exercises	351

23 Debugging a model: World Cup football	353
23.1 Constructing a model for score differentials	353
23.2 Model checking and exploration	355
23.3 Finding the bug and checking the corrected model	358
23.4 Assessing the role of the prior team rankings	359
23.5 Expanding to a discrete-data model	361
23.6 Evaluating models based on predictive accuracy	364
23.7 General lessons from the World Cup example	366
23.8 Exercises	366
24 Leave-one-out cross validation model checking and comparison: Roaches	369
24.1 Poisson model	369
24.2 Negative binomial model	372
24.3 Poisson model with varying intercepts	376
24.4 Poisson model with varying intercepts and integrated LOO	381
24.5 Zero-inflated negative binomial model	384
24.6 General lessons from the cross validation example	387
24.7 Exercises	388
25 Model building and expansion: Golf putting	389
25.1 First model: Logistic regression	389
25.2 Modeling from first principles	389
25.3 Testing the fitted model on new data	391
25.4 A new model accounting for how hard the ball is hit	392
25.5 Expanding the model by including a fudge factor	394
25.6 General lessons from the golf example	398
25.7 Exercises	399
26 Model building with latent variables: Markov models for animal movement	401
26.1 Time series of white shark movements	401
26.2 Developing a hidden Markov model	402
26.3 State decoding	408
26.4 Model checking, expansion, and conclusions	410
26.5 General lessons from the animal movement example	415
26.6 Exercises	416
27 Model building: Time-series decomposition for birthdays	417
27.1 From 30 days to 366	417
27.2 Additive decompositions	418
27.3 Gaussian process models	420
27.4 Incremental model building	422
27.5 General lessons from the birthdays example	433
27.6 Exercises	434
28 Models for regression coefficients and variable selection: Student grades	437
28.1 Variable selection in general	437
28.2 Models for regression coefficients	438
28.3 Marginal posteriors	441
28.4 Model checking	442
28.5 Projection predictive variable selection	443
28.6 General lessons from the variable selection example	448
28.7 Exercises	448

29 Sampling problems with latent variables: No vehicles in the park	449
29.1 Fitting a model	449
29.2 Understanding the fitted model	450
29.3 Near-collinearity in the posterior distribution	452
29.4 Fitting a Bayesian model	454
29.5 How to choose parameterization and other tricks	460
29.6 General lessons from the vehicles in the park example	460
29.7 Exercises	461
30 Challenge of multimodality: Differential equation for planetary motion	463
30.1 Mechanistic model of motion	463
30.2 Analyzing a simplified model	464
30.3 Improving the inference	466
30.4 Analyzing the full model	468
30.5 General lessons from the planetary motion example	468
30.6 Exercises	469
31 Simulation-based calibration checking in model development workflow	471
31.1 Target model	471
31.2 Developing the mixture submodel	472
31.3 Developing the logistic regression submodel	476
31.4 Putting it together	479
31.5 General lessons from the calibration example	481
31.6 Exercises	482
Appendices	483
A Statistical and computational workflow for Bayesians and non-Bayesians	485
A.1 Use of prior information	485
A.2 Combining information from multiple sources	485
A.3 Regularization	486
A.4 Using simulations to capture uncertainty	486
A.5 Prediction, generalization, and causal inference	487
A.6 Visualizing model checking and model fit	487
A.7 Fitting a sequence of models rather than focusing on just one	488
A.8 Initial values and tuning parameters	488
A.9 Simulated-data experimentation	489
A.10 Understanding methods by applying them to multiple problems	489
A.11 Awareness of goals	490
B How to get the most out of <i>Bayesian Data Analysis</i>	491
B.1 BDA3 part I: Fundamentals of Bayesian inference	491
B.2 BDA3 part II: Fundamentals of Bayesian data analysis	494
B.3 BDA3 part III: Advanced computation	495
B.4 BDA3 parts IV and V: Regression models and nonlinear and nonparametric models	497
B.5 BDA3 Appendixes	499
References	500
Author Index	523
Subject Index	533