


Bayesian Analysis in Natural Language Processing

Second Edition



Synthesis Lectures on Human Language Technologies

Editor

Graeme Hirst, *University of Toronto*

Synthesis Lectures on Human Language Technologies is edited by Graeme Hirst of the University of Toronto. The series consists of 50- to 150-page monographs on topics relating to natural language processing, computational linguistics, information retrieval, and spoken language understanding. Emphasis is on important new techniques, on new applications, and on topics that combine two or more HLT subfields.

Bayesian Analysis in Natural Language Processing, Second Edition

Shay Cohen
2019

Argumentation Mining

Manfred Stede and Jodi Schneider
2018

Quality Estimation for Machine Translation

Lucia Specia, Carolina Scarton, and Gustavo Henrique Paetzold
2018

Natural Language Processing for Social Media, Second Edition

Atefeh Farzindar and Diana Inkpen
2017

Automatic Text Simplification

Horacio Saggion
2017

Neural Network Methods for Natural Language Processing

Yoav Goldberg
2017

Syntax-based Statistical Machine Translation

Philip Williams, Rico Sennrich, Matt Post, and Philipp Koehn
2016

Domain-Sensitive Temporal Tagging

Jannik Strötgen and Michael Gertz
2016

Linked Lexical Knowledge Bases: Foundations and Applications

Iryna Gurevych, Judith Eckle-Kohler, and Michael Matuschek
2016

Bayesian Analysis in Natural Language Processing

Shay Cohen
2016

Metaphor: A Computational Perspective

Tony Veale, Ekaterina Shutova, and Beata Beigman Klebanov
2016

Grammatical Inference for Computational Linguistics

Jeffrey Heinz, Colin de la Higuera, and Menno van Zaanen
2015

Automatic Detection of Verbal Deception

Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari
2015

Natural Language Processing for Social Media

Atefeh Farzindar and Diana Inkpen
2015

Semantic Similarity from Natural Language and Ontology Analysis

Sébastien Harispe, Sylvie Ranwez, Stefan Janaqi, and Jacky Montmain
2015

Learning to Rank for Information Retrieval and Natural Language Processing, Second Edition

Hang Li
2014

Ontology-Based Interpretation of Natural Language

Philipp Cimiano, Christina Unger, and John McCrae
2014

Automated Grammatical Error Detection for Language Learners, Second Edition

Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault
2014

Web Corpus Construction

Roland Schäfer and Felix Bildhauer
2013

Recognizing Textual Entailment: Models and Applications

Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto
2013

Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Morphology and Syntax

Emily M. Bender
2013

Semi-Supervised Learning and Domain Adaptation in Natural Language Processing

Anders Søgaard
2013

Semantic Relations Between Nominals

Vivi Nastase, Preslav Nakov, Diarmuid Ó Séaghdha, and Stan Szpakowicz
2013

Computational Modeling of Narrative

Inderjeet Mani
2012

Natural Language Processing for Historical Texts

Michael Piotrowski
2012

Sentiment Analysis and Opinion Mining

Bing Liu
2012

Discourse Processing

Manfred Stede
2011

Bitext Alignment

Jörg Tiedemann
2011

Linguistic Structure Prediction

Noah A. Smith
2011

Learning to Rank for Information Retrieval and Natural Language Processing

Hang Li
2011

Computational Modeling of Human Language Acquisition

Afra Alishahi

2010

Introduction to Arabic Natural Language Processing

Nizar Y. Habash

2010

Cross-Language Information Retrieval

Jian-Yun Nie

2010

Automated Grammatical Error Detection for Language Learners

Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault

2010

Data-Intensive Text Processing with MapReduce

Jimmy Lin and Chris Dyer

2010

Semantic Role Labeling

Martha Palmer, Daniel Gildea, and Nianwen Xue

2010

Spoken Dialogue Systems

Kristiina Jokinen and Michael McTear

2009

Introduction to Chinese Natural Language Processing

Kam-Fai Wong, Wenjie Li, Ruifeng Xu, and Zheng-sheng Zhang

2009

Introduction to Linguistic Annotation and Text Analytics

Graham Wilcock

2009

Dependency Parsing

Sandra Kübler, Ryan McDonald, and Joakim Nivre

2009

Statistical Language Models for Information Retrieval

ChengXiang Zhai

2008

© Springer Nature Switzerland AG 2022
Reprint of original edition ©Morgan & Claypool 2019

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Bayesian Analysis in Natural Language Processing, Second Edition
Shay Cohen

ISBN: 978-3-031-01042-2 paperback
ISBN: 978-3-031-02170-1 ebook
ISBN: 978-3-031-00181-9 hardcover

DOI 10.1007/978-3-031-02170-1

A Publication in the Springer series
SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES

Lecture #41
Series Editor: Graeme Hirst, *University of Toronto*
Series ISSN
Print 1947-4040 Electronic 1947-4059

Bayesian Analysis in Natural Language Processing

Second Edition

Shay Cohen
University of Edinburgh

SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES #41

ABSTRACT

Natural language processing (NLP) went through a profound transformation in the mid-1980s when it shifted to make heavy use of corpora and data-driven techniques to analyze language. Since then, the use of statistical techniques in NLP has evolved in several ways. One such example of evolution took place in the late 1990s or early 2000s, when full-fledged Bayesian machinery was introduced to NLP. This Bayesian approach to NLP has come to accommodate various shortcomings in the frequentist approach and to enrich it, especially in the unsupervised setting, where statistical learning is done without target prediction examples.

In this book, we cover the methods and algorithms that are needed to fluently read Bayesian learning papers in NLP and to do research in the area. These methods and algorithms are partially borrowed from both machine learning and statistics and are partially developed “in-house” in NLP. We cover inference techniques such as Markov chain Monte Carlo sampling and variational inference, Bayesian estimation, and nonparametric modeling. In response to rapid changes in the field, this second edition of the book includes a new chapter on representation learning and neural networks in the Bayesian context. We also cover fundamental concepts in Bayesian statistics such as prior distributions, conjugacy, and generative modeling. Finally, we review some of the fundamental modeling techniques in NLP, such as grammar modeling, neural networks and representation learning, and their use with Bayesian analysis.

KEYWORDS

natural language processing, computational linguistics, Bayesian statistics, Bayesian NLP, statistical learning, inference in NLP, grammar modeling in NLP, neural networks, representation learning

Dedicated to Mia

Contents

	List of Figures	xix
	List of Algorithms	xxi
	List of Generative Stories	xxiii
	Preface (First Edition)	xxv
	Acknowledgments (First Edition)	xxix
	Preface (Second Edition)	xxxi
1	Preliminaries	1
1.1	Probability Measures	1
1.2	Random Variables	2
1.2.1	Continuous and Discrete Random Variables	3
1.2.2	Joint Distribution over Multiple Random Variables	4
1.3	Conditional Distributions	5
1.3.1	Bayes' Rule	6
1.3.2	Independent and Conditionally Independent Random Variables	7
1.3.3	Exchangeable Random Variables	8
1.4	Expectations of Random Variables	9
1.5	Models	11
1.5.1	Parametric vs. Nonparametric Models	11
1.5.2	Inference with Models	12
1.5.3	Generative Models	14
1.5.4	Independence Assumptions in Models	16
1.5.5	Directed Graphical Models	17
1.6	Learning from Data Scenarios	19
1.7	Bayesian and Frequentist Philosophy (Tip of the Iceberg)	22
1.8	Summary	23
1.9	Exercises	24

2	Introduction	25
2.1	Overview: Where Bayesian Statistics and NLP Meet	26
2.2	First Example: The Latent Dirichlet Allocation Model	30
2.2.1	The Dirichlet Distribution	34
2.2.2	Inference	36
2.2.3	Summary	39
2.3	Second Example: Bayesian Text Regression	40
2.4	Conclusion and Summary	41
2.5	Exercises	42
3	Priors	43
3.1	Conjugate Priors	44
3.1.1	Conjugate Priors and Normalization Constants	47
3.1.2	The Use of Conjugate Priors with Latent Variable Models	48
3.1.3	Mixture of Conjugate Priors	50
3.1.4	Renormalized Conjugate Distributions	51
3.1.5	Discussion: To Be or not to Be Conjugate?	52
3.1.6	Summary	53
3.2	Priors Over Multinomial and Categorical Distributions	53
3.2.1	The Dirichlet Distribution Re-Visited	55
3.2.2	The Logistic Normal Distribution	59
3.2.3	Discussion	65
3.2.4	Summary	65
3.3	Non-Informative Priors	66
3.3.1	Uniform and Improper Priors	67
3.3.2	Jeffreys Prior	68
3.3.3	Discussion	69
3.4	Conjugacy and Exponential Models	69
3.5	Multiple Parameter Draws in Models	70
3.6	Structural Priors	73
3.7	Conclusion and Summary	74
3.8	Exercises	75
4	Bayesian Estimation	77
4.1	Learning with Latent Variables: Two Views	78
4.2	Bayesian Point Estimation	79

4.2.1	Maximum a Posteriori Estimation	79
4.2.2	Posterior Approximations Based on the MAP Solution	87
4.2.3	Decision-Theoretic Point Estimation	89
4.2.4	Discussion and Summary	90
4.3	Empirical Bayes	90
4.4	Asymptotic Behavior of the Posterior	93
4.5	Summary	93
4.6	Exercises	95
5	Sampling Methods	97
5.1	MCMC Algorithms: Overview	98
5.2	NLP Model Structure for MCMC Inference	99
5.2.1	Partitioning the Latent Variables	100
5.3	Gibbs Sampling	101
5.3.1	Collapsed Gibbs Sampling	105
5.3.2	Operator View	109
5.3.3	Parallelizing the Gibbs Sampler	111
5.3.4	Summary	112
5.4	The Metropolis–Hastings Algorithm	113
5.4.1	Variants of Metropolis–Hastings	114
5.5	Slice Sampling	115
5.5.1	Auxiliary Variable Sampling	117
5.5.2	The Use of Slice Sampling and Auxiliary Variable Sampling in NLP	117
5.6	Simulated Annealing	118
5.7	Convergence of MCMC Algorithms	119
5.8	Markov Chain: Basic Theory	121
5.9	Sampling Algorithms Not in the MCMC Realm	123
5.10	Monte Carlo Integration	126
5.11	Discussion	127
5.11.1	Computability of Distribution vs. Sampling	127
5.11.2	Nested MCMC Sampling	128
5.11.3	Runtime of MCMC Samplers	128
5.11.4	Particle Filtering	128
5.12	Conclusion and Summary	130
5.13	Exercises	132

6	Variational Inference	135
6.1	Variational Bound on Marginal Log-Likelihood	135
6.2	Mean-Field Approximation	138
6.3	Mean-Field Variational Inference Algorithm	139
6.3.1	Dirichlet-Multinomial Variational Inference	141
6.3.2	Connection to the Expectation-Maximization Algorithm	145
6.4	Empirical Bayes with Variational Inference	147
6.5	Discussion	148
6.5.1	Initialization of the Inference Algorithms	148
6.5.2	Convergence Diagnosis	149
6.5.3	The Use of Variational Inference for Decoding	150
6.5.4	Variational Inference as KL Divergence Minimization	151
6.5.5	Online Variational Inference	151
6.6	Summary	152
6.7	Exercises	153
7	Nonparametric Priors	155
7.1	The Dirichlet Process: Three Views	156
7.1.1	The Stick-Breaking Process	157
7.1.2	The Chinese Restaurant Process	159
7.2	Dirichlet Process Mixtures	161
7.2.1	Inference with Dirichlet Process Mixtures	161
7.2.2	Dirichlet Process Mixture as a Limit of Mixture Models	165
7.3	The Hierarchical Dirichlet Process	165
7.4	The Pitman–Yor Process	168
7.4.1	Pitman–Yor Process for Language Modeling	169
7.4.2	Power-Law Behavior of the Pitman–Yor Process	170
7.5	Discussion	171
7.5.1	Gaussian Processes	172
7.5.2	The Indian Buffet Process	172
7.5.3	Nested Chinese Restaurant Process	173
7.5.4	Distance-Dependent Chinese Restaurant Process	173
7.5.5	Sequence Memoizers	175
7.6	Summary	175
7.7	Exercises	176

8	Bayesian Grammar Models	177
8.1	Bayesian Hidden Markov Models	178
8.1.1	Hidden Markov Models with an Infinite State Space	179
8.2	Probabilistic Context-Free Grammars	181
8.2.1	PCFGs as a Collection of Multinomials	184
8.2.2	Basic Inference Algorithms for PCFGs	185
8.2.3	Hidden Markov Models as PCFGs	189
8.3	Bayesian Probabilistic Context-Free Grammars	189
8.3.1	Priors on PCFGs	189
8.3.2	Monte Carlo Inference with Bayesian PCFGs	190
8.3.3	Variational Inference with Bayesian PCFGs	192
8.4	Adaptor Grammars	193
8.4.1	Pitman–Yor Adaptor Grammars	194
8.4.2	Stick-Breaking View of PYAG	196
8.4.3	Inference with PYAG	198
8.5	Hierarchical Dirichlet Process PCFGs (HDP-PCFGs)	200
8.5.1	Extensions to the HDP-PCFG Model	201
8.6	Dependency Grammars	202
8.6.1	State-Split Nonparametric Dependency Models	203
8.7	Synchronous Grammars	205
8.8	Multilingual Learning	206
8.8.1	Part-of-Speech Tagging	206
8.8.2	Grammar Induction	208
8.9	Further Reading	209
8.10	Summary	211
8.11	Exercises	212
9	Representation Learning and Neural Networks	213
9.1	Neural Networks and Representation Learning: Why Now?	214
9.2	Word Embeddings	217
9.2.1	Skip-Gram Models for Word Embeddings	218
9.2.2	Bayesian Skip-Gram Word Embeddings	220
9.2.3	Discussion	221
9.3	Neural Networks	222
9.3.1	Frequentist Estimation and the Backpropagation Algorithm	225
9.3.2	Priors on Neural Network Weights	228

9.4	Modern Use of Neural Networks in NLP	230
9.4.1	Recurrent and Recursive Neural Networks	230
9.4.2	Vanishing and Exploding Gradient Problem	232
9.4.3	Neural Encoder-Decoder Models	235
9.4.4	Convolutional Neural Networks	239
9.5	Tuning Neural Networks	242
9.5.1	Regularization	242
9.5.2	Hyperparameter Tuning	243
9.6	Generative Modeling with Neural Networks	245
9.6.1	Variational Autoencoders	246
9.6.2	Generative Adversarial Networks	252
9.7	Conclusion	254
9.8	Exercises	256
	Closing Remarks	257
A	Basic Concepts	259
A.1	Basic Concepts in Information Theory	259
A.1.1	Entropy and Cross Entropy	259
A.1.2	Kullback–Leibler Divergence	260
A.2	Other Basic Concepts	260
A.2.1	Jensen’s Inequality	260
A.2.2	The Chain Rule for Differentiation	261
A.2.3	Transformation of Continuous Random Variables	261
A.2.4	The Expectation-Maximization Algorithm	262
A.3	Basic Concepts in Optimization	263
A.3.1	Stochastic Gradient Descent	264
A.3.2	Constrained Optimization	265
B	Distribution Catalog	267
B.1	The Multinomial Distribution	267
B.2	The Dirichlet Distribution	268
B.3	The Poisson Distribution	269
B.4	The Gamma Distribution	269
B.5	The Multivariate Normal Distribution	270
B.6	The Laplace Distribution	271
B.7	The Logistic Normal Distribution	271

B.8 The Inverse Wishart Distribution 272

B.9 The Gumbel Distribution 273

Bibliography 275

Author’s Biography 305

Index 307

List of Figures

1.1	Graphical model for the latent Dirichlet allocation model	18
2.1	Fully Bayesian latent Dirichlet allocation	34
2.2	The Beta distribution	37
2.3	A plot of sampled data from the Dirichlet distribution	38
3.1	A plot of sampled data from the logistic normal distribution	60
3.2	Depiction of prior levels	71
5.1	Demonstration of slice sampling	116
5.2	Demonstration of rejection sampling	125
6.1	The exp-digamma function	146
7.1	A graphical depiction of the stick breaking process	157
7.2	A graphical depiction of the Chinese restaurant process posterior	160
7.3	The behavior of the Pitman-Yor process	171
8.1	Graphical model for hidden Markov models	179
8.2	Phrase-structure tree example	181
8.3	Example of dependency tree with latent states	203
9.1	A logistic regression neural network	223
9.2	An example of a feed-forward neural network	223
9.3	A diagram of a recurrent neural network	231
9.4	A diagram depicting an LSTM cell	234
9.5	A diagram of an encoder-decoder model	236
9.6	A diagram of a neural encoder-decoder with an attention mechanism	238
9.7	Diagrams for a bidirectional encoder	239

9.8	An example of a convolutional neural network	241
-----	--	-----

List of Algorithms

5.1	The Gibbs sampling algorithm	102
5.2	The collapsed Gibbs sampling algorithm	105
5.3	The operator Gibbs sampling algorithm	110
5.4	The parallel Gibbs sampling algorithm	112
5.5	The Metropolis-Hastings algorithm	113
5.6	The rejection sampling algorithm	124
6.1	The mean-field variational inference algorithm	140
6.2	The mean-field variational expectation-maximization algorithm	148
8.1	Derivation sampling algorithm for PCFGs	185
8.2	Posterior sampling for PCFGs	192
9.1	The Bayesian optimization algorithm	245

List of Generative Stories

1.1	Gaussian mixture model	15
2.1	The latent Dirichlet allocation model	32
2.2	General Bayesian model	36
3.1	The logistic normal distribution	63
7.1	The Dirichlet process	158
7.2	The Dirichlet process with Chinese restaurant representation	160
7.3	The Dirichlet process mixture model	162
7.4	The Dirichlet process mixture model using the stick-breaking representation	164
7.5	Approximation of the Dirichlet process mixture model	165
7.6	LDA with the Dirichlet process	167
8.1	The infinite HMM	180
8.2	Pitman-Yor adaptor grammars	195
8.3	Stick-breaking representation of adaptor grammars	197
8.4	The HDP-PCFG model	199
8.5	Bilingual model for part-of-speech tagging	208

Preface (First Edition)

When writing about a topic that intersects two areas (in this case, Bayesian Statistics and Natural Language Processing), the focus and the perspective need to be considered. I took a rather practical one in writing this book, aiming to write it for those in the same position as myself during my graduate studies. At that time, I already had a reasonable grasp of the problems in natural language processing and knowledge of the basic principles in machine learning. I wanted to learn more about Bayesian statistics—in a rather abstract way—particularly the parts that are most relevant to NLP. Thus, this book is written from that perspective, providing abstract information about the key techniques, terminology and models that a computational linguist would need in order to apply the Bayesian approach to his or her work.

Most chapters in this book, therefore, are rather general and have relevance to other uses of Bayesian statistics. The last chapter in this book, though, presents some specific NLP applications for grammar models that are mostly (but not exclusively) used in NLP.

Ideally, this book should be read by a person who already has some idea about statistical modeling in NLP, and wants to gain more depth about the specific application of Bayesian techniques to NLP. The motivation for this decision to focus more on the mathematical aspects of Bayesian NLP is simple. Most computational linguists get exposed quite early in their graduate career or otherwise to the basic core terminology in NLP, the linguistic structures it predicts and perhaps some of the linguistic motivation behind them. Ideas from Bayesian statistics and other statistical tools are often “picked up” on the way. As such, there are sometimes misconceptions and a missing global picture. This book tries to provide some of these missing details to the reader.

There are several approaches to doing statistics, two of which are the frequentist approach and the Bayesian approach. The frequentist approach is also sometimes referred to as “classic statistics.” One of the things that motivated me to learn more about Bayesian statistics is the rich and colorful history it has. To this day, the famous frequentist-Bayesian divide still exists. This kind of divide regarding the philosophy that statistical analysis should follow is even more persistently and more ardently argued about than theories of grammar were in the famous “linguistics war” between generative semanticists and generative grammarians. It does not end there—even within the Bayesian camp there are those who support a subjective interpretation of probability and those who support an objective one.

Although I was captivated by the mathematical elegance of Bayesian statistics when I was first exposed to the core ideas (in principle, Bayesian statistics relies on one basic principle of applying Bayes’ rule to invert the relationship between data and parameters), I take a pragmatic approach and do not try to present Bayesian statistics as the ultimate theory for doing statistical

NLP. I also do not provide the philosophical arguments that support Bayesian statistics in this monograph. Instead, I provide the technical mechanisms behind Bayesian statistics, and advise the reader to determine whether the techniques work well for him or her in the problems they work on. Here and there, I also describe some connections that Bayesian statistics have to the frequentist approach, and other points of confluence. If the reader is interested in learning more about the philosophy behind Bayesian statistics, I suggest reading [Jaynes \(2003\)](#) and also looking at [Barnett \(1999\)](#). To better understand the history and the people behind Bayesian statistics, I suggest reading the book by [McGrayne \(2011\)](#). This book consists of eight chapters as following:

Chapter 1 is a refresher about probability and statistics as they relate to Bayesian NLP. We cover basic concepts such as random variables, independence between random variables, conditional independence, and random variable expectations; we also briefly discuss Bayesian statistics and how it differs from frequentist statistics. Most of this chapter can probably be skipped if one already has some basic background in computer science or statistics.

Chapter 2 introduces Bayesian analysis in NLP with two examples (the latent Dirichlet allocation model and Bayesian text regression), and also provides a high-level overview of the topic.

Chapter 3 covers an important component in Bayesian statistical modeling—the prior. We discuss the priors that are most commonly used in Bayesian NLP, such as the Dirichlet distribution, non-informative priors, the normal distribution and others.

Chapter 4 covers ideas that bring together frequentist statistics and Bayesian statistics through the summarization of the posterior distribution. It details approaches to calculate a point estimate for a set of parameters while maintaining a Bayesian mindset.

Chapter 5 covers one of the main inference approaches in Bayesian statistics: Markov chain Monte Carlo. It details the most common sampling algorithms used in Bayesian NLP, such as Gibbs sampling and Metropolis-Hastings sampling.

Chapter 6 covers another important inference approach in Bayesian NLP, variational inference. It describes mean-field variational inference and the variational expectation-maximization algorithm.

Chapter 7 covers an important modeling technique in Bayesian NLP, nonparametric modeling. We discuss nonparametric models such as the Dirichlet process and the Pitman-Yor process.

Chapter 8 covers basic grammar models in NLP (such as probabilistic context-free grammars and synchronous grammars), and the way to frame them in a Bayesian context (using models such as adaptor grammars, hierarchical Dirichlet process PCFGs and others).

In addition, the book includes two appendices that provide background information that offers additional context for reading this book. Each chapter is accompanied by at least five exercises. This book (perhaps with the exercises) could be used as teaching material. Specifically, it could be used to teach a number of lectures about Bayesian analysis in NLP. If a significant amount of time is devoted to Bayesian NLP in class (say, four lectures), I would suggest devoting one lecture to chapter 3, one lecture to chapter 4, one lecture to chapters 5 and 6 together, and one lecture to chapter 7. Topics from chapter 8 (such as adaptor grammars or Bayesian PCFGs) can be injected into the various lectures as examples.

Acknowledgments (First Edition)

I am indebted to all the people who helped me to write this book. First, I would like to especially thank Lea Frermann, Trevor Cohn, and Jacob Eisenstein, who carefully read a draft of this book, and left detailed feedback. I would also like to thank all other people who gave feedback in one form or another: Omri Abend, Apoorv Agarwal, Anahita Bhiwandiwalla, Jordan Boyd-Graber, Daniel Gildea, Sharon Goldwater, Mark Johnson, Mirella Lapata, Shalom Lappin, Adam Lopez, Brendan O'Connor, Mohammad Sadegh Rasooli, Siva Reddy, Stefan Riezler, Giorgio Satta, Stuart Shieber, Mark Steedman, Karl Stratos, Swabha Swayamdipta, Bonnie Webber and Dani Yogatama. Thanks also to Sharon Rosenfeld, who proofread this book to help make it more readable. Thanks also to Samantha Draper, Graeme Hirst, Michael Morgan and CL Tondo for helping with the publication of this book.

I would also like to thank all of the great students who attended my class at the Department of Computer Science in Columbia University in Spring 2013 (“Bayesian analysis in NLP”) and indirectly helped me better understand what is the level of coverage needed from this book for young researchers making their first steps in the area of Bayesian NLP (students such as Jessica Forde, Daniel Perlmutter, and others whom I have already mentioned). Thanks also go to my collaborators on projects in the area of Bayesian NLP, who helped me shape my understanding of it: David Blei, Jordan Boyd-Graber, Kevin Gimpel, and Ke Zhai.

I want to also thank my mentors throughout all the years, and especially my advisor, Noah Smith, with whom I first studied Bayesian NLP, Michael Collins, my post-doc advisor, who supported me in spending time writing this book during my post-doctoral fellowship and in teaching the Bayesian NLP class at Columbia, and Mark Johnson, whose work, as well as our conversations and email exchanges, influenced me in writing this book.

Also, thanks to Sylvia Cohen, my spouse, who has always been there for me during the writing of this book. Similar thanks to Sylvia’s family, who always made me feel at home in Pittsburgh, while I was studying topics such as Bayesian analysis. Finally, I would like to thank my parents and siblings—whose prior beliefs in me never changed, no matter what the evidence was.

Shay Cohen
Edinburgh
May 2016

Preface (Second Edition)

I did not expect to release a second edition for this book so quickly, but the last few years of fast-paced and exciting developments in the world of Natural Language Processing (NLP) have called for various updates, leading to this second edition.

The main addition to this book is Chapter 9, which focuses on representation learning and neural networks in NLP, particularly in a Bayesian context. This chapter was written based on the observation that in the past five years or so, NLP literature has been dominated by the use of neural networks; and as such, I believe the fundamentals needed to be addressed in this book. Adapting the content to the Bayesian “mission” of this book (coupled with the NLP context) was not always easy, and I will let the reader be the judge of whether I have accomplished my mission.

In addition to introducing this new chapter in this edition, several typographical errors have been fixed, and some additional content was integrated into various chapters.

There are several people who have helped with this edition. I want to thank Trevor Cohn, Marco Damonte, Jacob Eisenstein, Lea Frermann, Annie Louis, Chunchuan Lyu, Nikos Papanastasiou, Shashi Narayan, Mark Steedman, Rico Sennrich, and Ivan Titov for their help and comments. I also want to thank my students and postdocs who have taught me more than I have taught them in some areas of the new material in this book.

Shay Cohen
Edinburgh
February 2019