
AN INTRODUCTION TO SEARCH ENGINES AND WEB NAVIGATION

MARK LEVENE

*Department of Computer Science and Information Systems
Birkbeck University of London, UK*



A JOHN WILEY & SONS, INC., PUBLICATION

*AN INTRODUCTION TO
SEARCH ENGINES AND
WEB NAVIGATION*

AN INTRODUCTION TO SEARCH ENGINES AND WEB NAVIGATION

MARK LEVENE

*Department of Computer Science and Information Systems
Birkbeck University of London, UK*

 **WILEY**

A JOHN WILEY & SONS, INC., PUBLICATION

Copyright © 2010 by John Wiley & Sons, Inc. All rights reserved

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Levene, M. (Mark), 1957-

An introduction to search engines and web navigation / Mark Levene.

p. cm.

ISBN 978-0-470-52684-2 (pbk.)

1. Internet searching. 2. Web search engines. I. Title.

ZA4230.L48 2010

025.0425-dc22

2010008435

Printed in Singapore

10 9 8 7 6 5 4 3 2 1

*To my wife Sara and three children
Tamara, Joseph and Oren*

CONTENTS

<i>PREFACE</i>	xiv
<i>LIST OF FIGURES</i>	xvii
CHAPTER 1 <i>INTRODUCTION</i>	1
1.1 Brief Summary of Chapters	2
1.2 Brief History of Hypertext and the Web	3
1.3 Brief History of Search Engines	6
CHAPTER 2 <i>THE WEB AND THE PROBLEM OF SEARCH</i>	9
2.1 Some Statistics	10
2.1.1 Web Size Statistics	10
2.1.2 Web Usage Statistics	15
2.2 Tabular Data Versus Web Data	18
2.3 Structure of the Web	20
2.3.1 Bow-Tie Structure of the Web	21
2.3.2 Small-World Structure of the Web	23
2.4 Information Seeking on the Web	24
2.4.1 Direct Navigation	24
2.4.2 Navigation within a Directory	25
2.4.3 Navigation using a Search Engine	26
2.4.4 Problems with Web Information Seeking	27
2.5 Informational, Navigational, and Transactional Queries	28
2.6 Comparing Web Search to Traditional Information Retrieval	29
2.6.1 Recall and Precision	30
2.7 Local Site Search Versus Global Web Search	32
2.8 Difference Between Search and Navigation	34
CHAPTER 3 <i>THE PROBLEM OF WEB NAVIGATION</i>	38
3.1 Getting Lost in Hyperspace and the Navigation Problem	39
3.2 How Can the Machine Assist in User Search and Navigation	42
3.2.1 The Potential Use of Machine Learning Algorithms	42
3.2.2 The Naive Bayes Classifier for Categorizing Web Pages	43
3.3 Trails Should be First Class Objects	46
3.4 Enter Markov Chains and Two Interpretations of Its Probabilities	49
3.4.1 Markov Chains and the Markov Property	49
3.4.2 Markov Chains and the Probabilities of Following Links	50
3.4.3 Markov Chains and the Relevance of Links	52

- 3.5 Conflict Between Web Site Owner and Visitor 54
- 3.6 Conflict Between Semantics of Web Site and the Business Model 57

CHAPTER 4 *SEARCHING THE WEB***60**

- 4.1 Mechanics of a Typical Search 61
- 4.2 Search Engines as Information Gatekeepers of the Web 64
- 4.3 Search Engine Wars, is the Dust Settling? 68
 - 4.3.1 Competitor Number One: Google 69
 - 4.3.2 Competitor Number Two: Yahoo 70
 - 4.3.3 Competitor Number Three: Bing 70
 - 4.3.4 Other Competitors 72
- 4.4 Statistics from Studies of Search Engine Query Logs 73
 - 4.4.1 Search Engine Query Logs 73
 - 4.4.2 Search Engine Query Syntax 75
 - 4.4.3 The Most Popular Search Keywords 77
- 4.5 Architecture of a Search Engine 78
 - 4.5.1 The Search Index 79
 - 4.5.2 The Query Engine 80
 - 4.5.3 The Search Interface 81
- 4.6 Crawling the Web 81
 - 4.6.1 Crawling Algorithms 82
 - 4.6.2 Refreshing Web Pages 84
 - 4.6.3 The Robots Exclusion Protocol 84
 - 4.6.4 Spider Traps 85
- 4.7 What Does it Take to Deliver a Global Search Service? 85

CHAPTER 5 *HOW DOES A SEARCH ENGINE WORK***91**

- 5.1 Content Relevance 94
 - 5.1.1 Processing Web Pages 94
 - 5.1.2 Interpreting the Query 96
 - 5.1.3 Term Frequency 96
 - 5.1.4 Inverse Document Frequency 99
 - 5.1.5 Computing Keyword TF-IDF Values 100
 - 5.1.6 Caching Queries 102
 - 5.1.7 Phrase Matching 102
 - 5.1.8 Synonyms 102
 - 5.1.9 Link Text 103
 - 5.1.10 URL Analysis 104
 - 5.1.11 Date Last Updated 104
 - 5.1.12 HTML Structure Weighting 104
 - 5.1.13 Spell Checking 105
 - 5.1.14 Non-English Queries 106
 - 5.1.15 Home Page Detection 107
 - 5.1.16 Related Searches and Query Suggestions 107
- 5.2 Link-Based Metrics 108
 - 5.2.1 Referential and Informational Links 109
 - 5.2.2 Combining Link Analysis with Content Relevance 110
 - 5.2.3 Are Links the Currency of the Web? 110

5.2.4	PageRank Explained	112
5.2.5	Online Computation of PageRank	116
5.2.6	Monte Carlo Methods in PageRank Computation	116
5.2.7	Hyperlink-Induced Topic Search	117
5.2.8	Stochastic Approach for Link-Structure Analysis	120
5.2.9	Counting Incoming Links	122
5.2.10	The Bias of PageRank against New Pages	123
5.2.11	PageRank within a Community	123
5.2.12	Influence of Weblogs on PageRank	124
5.2.13	Link Spam	125
5.2.14	Citation Analysis	127
5.2.15	The Wide Ranging Interest in PageRank	129
5.3	Popularity-Based Metrics	130
5.3.1	Direct Hit's Popularity Metric	130
5.3.2	Document Space Modification	132
5.3.3	Using Query Log Data to Improve Search	132
5.3.4	Learning to Rank	133
5.3.5	BrowseRank	134
5.4	Evaluating Search Engines	136
5.4.1	Search Engine Awards	136
5.4.2	Evaluation Metrics	136
5.4.3	Performance Measures	138
5.4.4	Eye Tracking Studies	139
5.4.5	Test Collections	141
5.4.6	Inferring Ranking Algorithms	142

CHAPTER 6 *DIFFERENT TYPES OF SEARCH ENGINES*

148

6.1	Directories and Categorization of Web Content	150
6.2	Search Engine Advertising	152
6.2.1	Paid Inclusion	152
6.2.2	Banner Ads	153
6.2.3	Sponsored Search and Paid Placement	153
6.2.4	Behavioral Targeting	157
6.2.5	User Behavior	158
6.2.6	The Trade-Off between Bias and Demand	160
6.2.7	Sponsored Search Auctions	161
6.2.8	Pay per Action	165
6.2.9	Click Fraud and Other Forms of Advertising Fraud	166
6.3	Metasearch	168
6.3.1	Fusion Algorithms	169
6.3.2	Operational Metasearch Engines	170
6.3.3	Clustering Search Results	173
6.3.4	Classifying Search Results	175
6.4	Personalization	178
6.4.1	Personalization versus Customization	180
6.4.2	Personalized Results Tool	180
6.4.3	Privacy and Scalability	182
6.4.4	Relevance Feedback	182
6.4.5	Personalized PageRank	184

- 6.4.6 Outride's Personalized Search 186
- 6.5 Question Answering (Q&A) on the Web 187
 - 6.5.1 Natural Language Annotations 188
 - 6.5.2 Factual Queries 190
 - 6.5.3 Open Domain Question Answering 191
 - 6.5.4 Semantic Headers 193
- 6.6 Image Search 194
 - 6.6.1 Text-Based Image Search 195
 - 6.6.2 Content-Based Image Search 196
 - 6.6.3 VisualRank 198
 - 6.6.4 CAPTCHA and reCAPTCHA 200
 - 6.6.5 Image Search for Finding Location-Based Information 200
- 6.7 Special Purpose Search Engines 201

CHAPTER 7 NAVIGATING THE WEB

- 7.1 Frustration in Web Browsing and Navigation 211
 - 7.1.1 HTML and Web Site Design 211
 - 7.1.2 Hyperlinks and Surfing 211
 - 7.1.3 Web Site Design and Usability 212
- 7.2 Navigation Tools 213
 - 7.2.1 The Basic Browser Tools 213
 - 7.2.2 The Back and Forward Buttons 214
 - 7.2.3 Search Engine Toolbars 215
 - 7.2.4 The Bookmarks Tool 216
 - 7.2.5 The History List 219
 - 7.2.6 Identifying Web Pages 219
 - 7.2.7 Breadcrumb Navigation 221
 - 7.2.8 Quicklinks 222
 - 7.2.9 Hypertext Orientation Tools 223
 - 7.2.10 Hypercard Programming Environment 224
- 7.3 Navigational Metrics 225
 - 7.3.1 The Potential Gain 226
 - 7.3.2 Structural Analysis of a Web Site 228
 - 7.3.3 Measuring the Usability of Web Sites 229
- 7.4 Web Data Mining 230
 - 7.4.1 Three Perspectives on Data Mining 230
 - 7.4.2 Measuring the Success of a Web Site 231
 - 7.4.3 Web Analytics 233
 - 7.4.4 E-Metrics 233
 - 7.4.5 Web Analytics Tools 234
 - 7.4.6 Weblog File Analyzers 235
 - 7.4.7 Identifying the Surfer 236
 - 7.4.8 Sessionizing 237
 - 7.4.9 Supplementary Analyses 237
 - 7.4.10 Markov Chain Model of Web Site Navigation 238
 - 7.4.11 Applications of Web Usage Mining 242
 - 7.4.12 Information Extraction 244
- 7.5 The Best Trail Algorithm 245
 - 7.5.1 Effective View Navigation 245

- 7.5.2 Web Usage Mining for Personalization 246
- 7.5.3 Developing a Trail Engine 246
- 7.6 Visualization that Aids Navigation 252
 - 7.6.1 How to Visualize Navigation Patterns 252
 - 7.6.2 Overview Diagrams and Web Site Maps 253
 - 7.6.3 Fisheye Views 255
 - 7.6.4 Visualizing Trails within a Web Site 257
 - 7.6.5 Visual Search Engines 258
 - 7.6.6 Social Data Analysis 259
 - 7.6.7 Mapping Cyberspace 262
- 7.7 Navigation in Virtual and Physical Spaces 262
 - 7.7.1 Real-World Web Usage Mining 262
 - 7.7.2 The Museum Experience Recorder 264
 - 7.7.3 Navigating in the Real World 265

CHAPTER 8 THE MOBILE WEB

272

- 8.1 The Paradigm of Mobile Computing 273
 - 8.1.1 Wireless Markup Language 274
 - 8.1.2 The i-mode Service 275
- 8.2 Mobile Web Services 277
 - 8.2.1 M-commerce 277
 - 8.2.2 Delivery of Personalized News 278
 - 8.2.3 Delivery of Learning Resources 281
- 8.3 Mobile Device Interfaces 282
 - 8.3.1 Mobile Web Browsers 282
 - 8.3.2 Information Seeking on Mobile Devices 284
 - 8.3.3 Text Entry on Mobile Devices 284
 - 8.3.4 Voice Recognition for Mobile Devices 286
 - 8.3.5 Presenting Information on a Mobile Device 287
- 8.4 The Navigation Problem in Mobile Portals 291
 - 8.4.1 Click-Distance 291
 - 8.4.2 Adaptive Mobile Portals 292
 - 8.4.3 Adaptive Web Navigation 294
- 8.5 Mobile Search 295
 - 8.5.1 Mobile Search Interfaces 296
 - 8.5.2 Search Engine Support for Mobile Devices 298
 - 8.5.3 Focused Mobile Search 299
 - 8.5.4 Laid Back Mobile Search 300
 - 8.5.5 Mobile Query Log Analysis 301
 - 8.5.6 Personalization of Mobile Search 302
 - 8.5.7 Location-Aware Mobile Search 303

CHAPTER 9 SOCIAL NETWORKS

309

- 9.1 What is a Social Network? 311
 - 9.1.1 Milgram's Small-World Experiment 312
 - 9.1.2 Collaboration Graphs 313
 - 9.1.3 Instant Messaging Social Network 314

xii CONTENTS

- 9.1.4 The Social Web 314
- 9.1.5 Social Network Start-Ups 316
- 9.2 Social Network Analysis 320
 - 9.2.1 Social Network Terminology 320
 - 9.2.2 The Strength of Weak Ties 322
 - 9.2.3 Centrality 322
 - 9.2.4 Web Communities 324
 - 9.2.5 Pajek: Large Network Analysis Software 326
- 9.3 Peer-to-Peer Networks 326
 - 9.3.1 Centralized P2P Networks 327
 - 9.3.2 Decentralized P2P Networks 328
 - 9.3.3 Hybrid P2P Networks 330
 - 9.3.4 Distributed Hash Tables 331
 - 9.3.5 BitTorrent File Distribution 331
 - 9.3.6 JXTA P2P Search 332
 - 9.3.7 Incentives in P2P Systems 332
- 9.4 Collaborative Filtering 333
 - 9.4.1 Amazon.com 333
 - 9.4.2 Collaborative Filtering Explained 334
 - 9.4.3 User-Based Collaborative Filtering 335
 - 9.4.4 Item-Based Collaborative Filtering 337
 - 9.4.5 Model-Based Collaborative Filtering 338
 - 9.4.6 Content-Based Recommendation Systems 339
 - 9.4.7 Evaluation of Collaborative Filtering Systems 340
 - 9.4.8 Scalability of Collaborative Filtering Systems 341
 - 9.4.9 A Case Study of Amazon.co.uk 341
 - 9.4.10 The Netflix Prize 342
 - 9.4.11 Some Other Collaborative Filtering Systems 346
- 9.5 Weblogs (Blogs) 347
 - 9.5.1 Blogrolling 348
 - 9.5.2 Blogspace 348
 - 9.5.3 Blogs for Testing Machine Learning Algorithms 349
 - 9.5.4 Spreading Ideas via Blogs 349
 - 9.5.5 The Real-Time Web and Microblogging 350
- 9.6 Power-Law Distributions in the Web 352
 - 9.6.1 Detecting Power-Law Distributions 353
 - 9.6.2 Power-Law Distributions in the Internet 355
 - 9.6.3 A Law of Surfing and a Law of Participation 355
 - 9.6.4 The Evolution of the Web via Preferential Attachment 357
 - 9.6.5 The Evolution of the Web as a Multiplicative Process 359
 - 9.6.6 The Evolution of the Web via HOT 360
 - 9.6.7 Small-World Networks 361
 - 9.6.8 The Robustness and Vulnerability of a Scale-Free Network 366
- 9.7 Searching in Social Networks 369
 - 9.7.1 Social Navigation 369
 - 9.7.2 Social Search Engines 370
 - 9.7.3 Navigation Within Social Networks 373
 - 9.7.4 Navigation Within Small-World Networks 375
 - 9.7.5 Testing Navigation Strategies in Social Networks 379
- 9.8 Social Tagging and Bookmarking 379

9.8.1	Flickr—Sharing Your Photos	380
9.8.2	YouTube—Broadcast Yourself	380
9.8.3	Delicious for Social Bookmarking	382
9.8.4	Communities Within Content Sharing Sites	382
9.8.5	Sharing Scholarly References	383
9.8.6	Folksonomy	383
9.8.7	Tag Clouds	384
9.8.8	Tag Search and Browsing	385
9.8.9	The Efficiency of Tagging	388
9.8.10	Clustering and Classifying Tags	389
9.9	Opinion Mining	390
9.9.1	Feature-Based Opinion Mining	391
9.9.2	Sentiment Classification	392
9.9.3	Comparative Sentence and Relation Extraction	393
9.10	Web 2.0 and Collective Intelligence	393
9.10.1	Ajax	394
9.10.2	Syndication	395
9.10.3	Open APIs, Mashups, and Widgets	396
9.10.4	Software as a Service	398
9.10.5	Collective Intelligence	399
9.10.6	Algorithms for Collective Intelligence	401
9.10.7	Wikipedia—The World’s Largest Encyclopedia	402
9.10.8	eBay—The World’s Largest Online Trading Community	407

CHAPTER 10	<i>THE FUTURE OF WEB SEARCH AND NAVIGATION</i>	419
-------------------	---	------------

<i>BIBLIOGRAPHY</i>	424
---------------------	------------

<i>INDEX</i>	463
--------------	------------

PREFACE

MOTIVATION

Searching and navigating the web have become part of our daily online lives. Web browsers and the standard navigation tools embedded in them provide a showcase of successful software technology with a global user-base, that has changed the way in which we search for and interact with information. Search engine technology has become ubiquitous, providing a standard interface to the endless amount of information that the web contains. Since the inception of the web, search engines have delivered a continuous stream of innovations, satisfying their users with increasingly accurate results through the implementation of advanced retrieval algorithms and scalable distributed architectures. Search and navigation technologies are central to the smooth operation of the web and it is hard to imagine finding information without them. Understanding the computational basis of these technologies and the models underlying them is of paramount importance both for IT students and practitioners.

There are several technical books on web search and navigation but the ones I have seen are either very academic in nature, that is, targeted at the post-graduate student or advanced researcher, and therefore have a limited audience, or they concentrate on the user interface and web site usability issues, ignoring the technicalities of what is happening behind the scenes. These books do not explain at an introductory level how the underlying computational tools work. This book answers the need for an introductory, yet technical, text on the topic.

My research into web search and navigation technologies started during the beginning of the 1990s just before the internet boom, when, together with my colleagues, we began looking at hypertext as a model for unstructured (or semistructured) data connected via a network of links, much in the same way web pages are connected. Of particular interest to us was the infamous “navigation problem” when we lose our way navigating (or what has become known as “surfing”) through the myriad of information pages in the network. Tackling this problem has provided continued impetus for my research.

In a wider context, the activity of information seeking, that is, the process we go through when searching and locating information in order to augment our state of knowledge, has been of major concern to all involved in the development of technologies that facilitate web interaction.

I have been using browser navigation tools and search engines since their early days, and have been fascinated by the flow of new ideas and the improvements that each new tool has delivered. One of my aims in this text is to demystify the technology underlying the tools that we use in our day-to-day interaction with

the web, and another is to inform readers about upcoming technologies, some of which are still in the research and development stage.

I hope that this book will instill in you some of my enthusiasm for the possibilities that these technologies have and are creating to extend our capabilities of finding and sharing information.

AUDIENCE AND PREREQUISITES

The book is intended as an undergraduate introductory text on search and navigation technologies, but could also be used to teach an option on the subject. It is also intended as a reference book for IT professionals wishing to know how these technologies work and to learn about the bigger picture in this area.

The course has no formal prerequisites, all that is required is for the learner to be a user of the web and to be curious to know how these technologies work. All the concepts that are introduced are explained in words, and simple examples from my own experience are given to illustrate various points. Occasionally, to add clarity to an important concept, a formula is given and explained. Each chapter starts with a list of learning objectives and ends with a brief bullet-pointed summary. There are several exercises at the end of each chapter. Some of these aim to get the student to explore further issues, possibly with a reference which can be followed up, some get the student to discuss an aspect of the technology, and others are mini-projects (which may involve programming) to add to the student's understanding through a hands-on approach. The book ends with a set of notes containing web addresses to items mentioned in the book, and an extensive bibliography of the articles and books cited in the book.

Readers should be encouraged to follow the links in the text and to discover new and related links that will help them understand how search and navigation tools work, and to widen their knowledge with related information.

TIMELINESS

I believe that due to the importance of the topic it is about time that such a book should appear. Search and navigation technologies are moving at a very fast pace due to the continued growth of the web and its user base, and improvements in computer networking and hardware. There is also strong competition between different service providers to lock-in users to their products. This is good news for web users, but as a result some of the numerics in the text may be out of date. I have qualified the statistics I have given with dates and links, which can be found in the notes, so the reader can follow these to get an up-to-date picture and follow the trends. I do not expect the core technologies I have covered to radically change in the near future and I would go so far as to claim that in essence they are fundamental to the web's working, but innovation and new ideas will continue to flourish and mold the web's landscape.

If you find any errors or omissions please let me know so that I can list them on the book's web site. I will also be grateful to receive any constructive comments and suggestions, which can be used to improve the text.

ACKNOWLEDGMENTS

First I would like to thank my wife and family who have been extremely supportive throughout this project, encouraging me to put in the extra hours needed to complete such a task. I would also like to thank my colleagues at the Department of Computer Science and Information Systems at Birkbeck, who have read and commented on parts of the book. Special thanks to my editors at Wiley, Lucy Hitz and George Telecki, who have patiently guided me through the publication process. Finally, I would like to thank the reviewers for their constructive comments.

The people who have built the innovative technologies that drive today's web are the real heroes of the revolution that the World Wide Web has brought upon us. Without them, this book could not have been written. Not only in terms of the content of the book, but also in terms of the tools I have been using daily to augment my knowledge on how search and navigation technologies work in practice.

Mark Levene
London, June 2010

LIST OF FIGURES

1.1	Bush's memex	3
1.2	Nelson's Xanadu	5
1.3	Engelbart's NLS	5
1.4	Mosaic browser initially released in 1993	6
2.1	A database table versus an example web site	19
2.2	Partial map of the Internet (2005)	21
2.3	A view of web pages related to www.dcs.bbk.ac.uk	22
2.4	Bow-tie shape of the Web	23
2.5	Information seeking	26
2.6	Recall versus precision	31
3.1	The navigation problem	40
3.2	Being lost in hyperspace	41
3.3	Example web site	47
3.4	Four trails within the web site	47
3.5	Search engine results for the query "mark research" submitted to Google	48
3.6	Relevant trail for the query "mark research"	49
3.7	Markov chain for example web site	51
3.8	Two trails in the Markov chain	51
3.9	Probabilities of the four trails	52
3.10	Scoring web pages	53
3.11	Constructing a Markov chain from scores	54
4.1	Query "computer chess" submitted to Google	62
4.2	Results for "computer chess" from Google	62
4.3	Relevant category from the directory for "computer chess" from Google	64
4.4	Results for phrase "computer chess" from Google	65
4.5	Simplified search engine architecture	78
4.6	Basic crawler algorithm	82
5.1	Results from "~ food ~ facts" Google	103
5.2	Web pages both A and B have a link to C	109
5.3	Example web site to illustrate PageRank	113
5.4	Example of a rank sink	113
5.5	Page A is an authority	118
5.6	Page A is a hub	119
5.7	The base set for the HITS technique	120
5.8	Web pages related to Yahoo	129
5.9	Heatmap of Google's golden triangle	140

6.1	Query “chess” submitted to Overture	155
6.2	Query “netbook” submitted to Google	155
6.3	Query “netbook” submitted to Yahoo	156
6.4	Query “netbook” submitted to Bing	156
6.5	Query “chess” submitted to Clusty	173
6.6	Query “chess” submitted to CatS	176
6.7	The right-hand window, generated by PResTo! when the query “salsa” is typed into Google	181
6.8	Question “who is the prime minister of the uk?” submitted to Ask Jeeves	190
6.9	Query “who is the prime minister of the uk?” submitted to Wolfram Alpha	191
6.10	Similarity graph generated from the top 1000 search results of “Mona-Lisa”	199
7.1	Navigation tools built-in to the browser	214
7.2	GroupLab Internet Explorer History System	220
7.3	Google sitelinks for query “Birkbeck”	222
7.4	Nielsen’s hypertext implemented in Apple’s Hypercard	225
7.5	Example web site	226
7.6	Pie chart showing the page request to my site	236
7.7	Pie chart showing the keywords that led to referrals to my site	236
7.8	Markov chain example	239
7.9	Frequent trails	240
7.10	Markov chain constructed from search engine scores	247
7.11	A navigation tree	249
7.12	Trail search for the query “knowledge technologies”	250
7.13	Nav-search for the query “knowledge technologies”	250
7.14	Visual search for the query “knowledge technologies”	251
7.15	The WebBrain user interface	254
7.16	Open Directory home page	254
7.17	Hierarchical site map	255
7.18	A web site for an online journal application	256
7.19	Query-specific site map	256
7.20	RSVP browser on a small screen	258
7.21	VISVIP visualization of user trails laid over the web site	259
7.22	Sea anemones in an aquarium	260
7.23	Anemone web site usage visualization	260
7.24	Carrot’s topic map for the query “beatles”	261
7.25	Kartoo’s topic map for the query “beatles”	261
7.26	The trail of a visitor to a museum	263
7.27	Navigation engine architecture	266
8.1	The Daily Learner adaptive news agent	279
8.2	A map in the Knowledge Sea system	282
8.3	Standard mobile phone 12-key keypad	285
8.4	A web page thumbnail overview (a) and a detailed view of a selected segment (b) as presented by SmartView	288
8.5	Accordion page summary on Power Browser	289
8.6	Example summary generated by BCL Technologies (www.bcltechnologies.com) summarizer	290

8.7	A typical navigation session in a mobile portal	293
8.8	The portal menu after personalization	294
8.9	Web search on a standard mobile phone	296
8.10	Screenshot of WebTwig prototype interface	297
9.1	Stanford University social web	315
9.2	Example of a collaboration graph	320
9.3	Example to illustrate centrality in social networks	323
9.4	Netflix prize final leaderboard	344
9.5	The user interface of MovieLens	346
9.6	Trend graph for keywords: “google” and “yahoo”	350
9.7	Web inlink data from May 1999	354
9.8	Scale-free network with 130 nodes	358
9.9	Example of network growth and preferential attachment	359
9.10	First 10,000 nodes of a HOT network with 100,000 nodes and $\alpha = 10$	362
9.11	Log–log plot for the cumulative degree distribution of the network shown in Fig. 9.10	363
9.12	The emergence of a small-world network: (a) regular, (b) small world, and (c) random	364
9.13	Average distance and clustering coefficient for the network construction shown in Fig. 9.12	364
9.14	User interface of the context-aware proxy-based system for the query “context-aware”	371
9.15	StumbleUpon toolbar	373
9.16	A two-dimensional lattice with random shortcuts	376
9.17	The local knowledge including the neighbors of a neighbor	377
9.18	Tag cloud of all time popular tags on Flickr; June 24, 2009	385
9.19	Tag cloud of popular tags on Delicious; June 24, 2009	385
9.20	Wordle tag cloud of Search Engine Land RSS feeds; June 24, 2009	386
9.21	MrTaggy’s search results for the query “social bookmarking”	387
9.22	The RSS logo	395
9.23	Live bookmarks in Firefox	396
9.24	Treemap visualization of news	397
9.25	The growth of English Wikipedia based on a logistic growth model	404

INTRODUCTION

“People keep asking me what I think of it now it’s done. Hence my protest:
The Web is not done!”

— Tim Berners-Lee, Inventor of the World Wide Web

THE LAST two decades have seen dramatic revolutions in information technology; not only in computing power, such as processor speed, memory size, and innovative interfaces, but also in the everyday use of computers. In the late 1970s and during the 1980s, we had the revolution of the personal computer (PC), which brought the computer into the home, the classroom, and the office. The PC then evolved into the desktop, the laptop, and the netbook as we know them today.

The 1990s was the decade of the World Wide Web (the Web), built over the physical infrastructure of the Internet, radically changing the availability of information and making possible the rapid dissemination of digital information across the globe. While the Internet is a physical network, connecting millions of computers together globally, the Web is a virtual global network linking together a massive amount of information. Search engines now index many billions of web pages and that number is just a fraction of the totality of information we can access on the Web, much of it residing in searchable databases not directly accessible to search engines.

Now, in the twenty-first century we are in the midst of a third wave of novel technologies, that of mobile and wearable computing devices, where computing devices have already become small enough so that we can carry them around with us at all times, and they also have the ability to interact with other computing devices, some of which are embedded in the environment. While the Web is mainly an informational and transactional tool, mobile devices add the dimension of being a location-aware ubiquitous social communication tool.

Coping with, organizing, visualizing, and acting upon the massive amount of information with which we are confronted when connected to the Web are amongst the main problems of *web interaction* [421]. Searching and navigating (or surfing) the Web are the methods we employ to help us find information

on the web, using search engines and navigation tools that are either built-in or plugged-in to the browser or are provided by web sites.

In this book, we explore search and navigation technologies to their full, present the State-of-the art tools, and explain how they work. We also look at ways of modeling different aspects of the Web that can help us understand how the Web is evolving and how it is being and can be used. The potential of many of the technologies we introduce has not yet been fully realized, and many new ideas to improve the ways in which we interact with the Web will inevitably appear in this dynamic and exciting space.

1.1 BRIEF SUMMARY OF CHAPTERS

This book is roughly divided into three parts. The first part (Chapters 1–3) introduces the problems of web interaction dealt with in the book, the second part (Chapters 4–6) deals with web search engines, and the third part (Chapters 7–9) looks at web navigation, the mobile web, and social network technologies in the context of search and navigation. Finally, in Chapter 10, we look ahead at the future prospects of search and navigation on the Web.

Chapters 1–3 introduce the reader to the problems of search and navigation and provide background material on the Web and its users. In particular, in the remaining part of Chapter 1, we give brief histories of hypertext and the Web, and of search engines. In Chapter 2, we look at some statistics regarding the Web, investigate its structure, and discuss the problems of information seeking and web search. In Chapter 3, we introduce the navigation problem, discuss the potential of machine learning to improve search and navigation tools, and propose Markov chains as a model for user navigation.

Chapters 4–6 cover the architectural and technical aspects of search engines. In particular, in Chapter 4, we discuss the search engine wars, look at some usage statistics of search engines, and introduce the architecture of a search engine, including the details of how the Web is crawled. In Chapter 5, we dissect a search engine's ranking algorithm, including content relevance, link- and popularity-based metrics, and different ways of evaluating search engines. In Chapter 6, we look at different types of search engines, namely, web directories, search engine advertising, metasearch engines, personalization of search, question answering engines, and image search and special purpose engines.

Chapters 7–9 concentrate on web navigation, and looks beyond at the mobile web and at how viewing the Web in social network terms is having a major impact on search and navigation technologies. In particular, in Chapter 7, we discuss a range of navigation tools and metrics, introduce web data mining and the Best Trail algorithm, discuss some visualization techniques to assist navigation, and look at the issues present in real-world navigation. In Chapter 8, we introduce the mobile web in the context of mobile computing, look at the delivery of mobile web services, discuss interfaces to mobile devices, and present the problems of search and navigation in a mobile context. In Chapter 9,

we introduce social networks in the context of the Web, look at social network analysis, introduce peer-to-peer networks, look at the technology of collaborative filtering, introduce weblogs as a medium for personal journalism on the Web, look at the ubiquity of power-law distributions on the Web, present effective searching strategies in social networks, introduce opinion mining as a way of obtaining knowledge about users opinions and sentiments, and look at Web 2.0 and collective intelligence that have generated a lot of hype and inspired many start-ups in recent years.

1.2 BRIEF HISTORY OF HYPERTEXT AND THE WEB

The history of the Web dates back to 1945 when Vannevar Bush, then an advisor to President Truman, wrote his visionary article “As We May Think,” and described his imaginary desktop machine called *memex*, which provides personal access to all the information we may need [119]. An artist’s impression of memex is shown in Fig. 1.1.

The memex is a “sort of mechanized private file and library,” which supports “associative indexing” and allows navigation whereby “any item may be caused at will to select immediately and automatically another.” Bush emphasizes that “the process of tying two items together is an important thing.” By repeating this process of creating links, we can form a *trail* which can be traversed by the user; in Bush’s words, “when numerous items have been thus joined together to form a trail they can be reviewed in turn.” The motivation for the memex’s support of trails as first-class objects was that the human mind “operates by association” and “in accordance to some intricate web of trails carried out by the cells of the brain.”




Figure 1.1 Bush’s memex. (Source: Life Magazine 1945;9(11):123.)

Bush also envisaged the “new profession of trailblazers” who create trails for other memex users, thus enabling sharing and exchange of knowledge. The memex was designed as a personal desktop machine, where information is stored locally on the machine. Trigg [647] emphasizes that Bush views the activities of creating a new trail and following a trail as being connected. Trails can be authored by trailblazers based on their experience and can also be created by memex, which records all user navigation sessions. In his later writings on the memex, published in Ref. 509, Bush revisited and extended the memex concept. In particular, he envisaged that memex could “learn from its own experience” and “refine its trails.” By this, Bush means that memex collects statistics on the trails that the user follows and “notifies” the ones that are most frequently followed. Oren [516] calls this extended version *adaptive memex*, stressing that adaptation means that trails can be constructed dynamically and given semantic justification; for example, by giving these new trails meaningful names.

The term *hypertext* [503] was coined by Ted Nelson in 1965 [495], who considers “a literature” (such as the scientific literature) to be a *system of interconnected writings*. The process of referring to other connected writings, when reading an article or a document, is that of *following links*. Nelson’s vision is that of creating a repository of all the documents that have ever been written thus achieving a universal hypertext. Nelson views his hypertext system, which he calls *Xanadu*, as a network of distributed documents that should be allowed to grow without any size limit, such that users, each corresponding to a node in the network, may link their documents to any other documents in the network. Xanadu can be viewed as a generalized memex system, which is both for private and public use. As with memex, Xanadu remained a vision that was not fully implemented; a mockup of Xanadu’s linking mechanism is shown in Fig. 1.2. Nelson’s pioneering work in hypertext is materialized to a large degree in the Web, since he also views his system as a means of publishing material by making it universally available to a wide network of interconnected users.

Douglas Engelbart’s *on-line system (NLS)* [205] was the first working hypertext system, where documents could be linked to other documents and thus groups of people could work collaboratively. The video clips of Engelbart’s historic demonstration of NLS from December 1968 are archived on the Web,¹ and a recollection of the demo can be found in Ref. 204; a picture of Engelbart during the demo is shown in Fig. 1.3.

About 30 years later in 1990, Tim Berners-Lee—then working for Cern, the world’s largest particle physics laboratory—turned the vision of hypertext into reality by creating the World Wide Web as we know it today [77].²

The Web works using three conventions: (i) the URL (unified resource locator) to identify web pages, (ii) HTTP (hypertext transfer protocol) to exchange messages between a browser and web server, and (iii) HTML (hypertext markup language) [501] to display web pages. More recently, Tim Berners-Lee has been

¹Video clips from Engelbart’s demo can be found at <http://sloan.stanford.edu/mousesite/1968Demo.html>.

²A little history of the World Wide Web from 1945 to 1995. www.w3.org/History.html.




Figure 1.2 Nelson's Xanadu. (Source: Figure 1.3, Xanalogical structure, needed now more than ever: Parallel documents, deep links to content, deep versioning, and deep re-use, by Nelson TH. www.cs.brown.edu/memex/ACM_HypertextTestbed/papers/60.html.)



Figure 1.3 Engelbart's NLS. (Source: Home video of the birth of the hyperlink. www.ratchetup.com/eyes/2004/01/wired_recently_.html.)

promoting the semantic web [78] together with XML (extensible markup language) [259], and RDF (resource description framework) [544], as a means of creating machine understandable information that can better support end user web applications. Details on the first web browser implemented by Tim Berners-Lee in 1990 can be found at www.w3.org/People/Berners-Lee/WorldWideWeb.

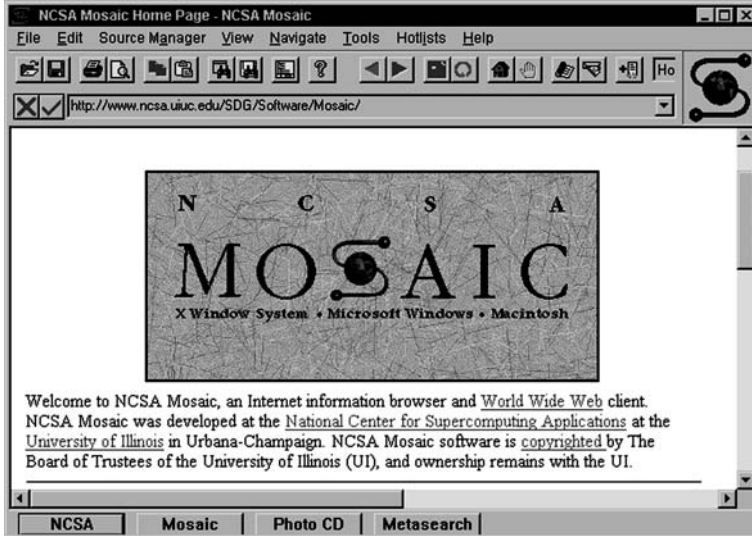


Figure 1.4 Mosaic browser initially released in 1993. (Source: <http://gladiator.ncsa.illinois.edu/Images/press-images/mosaic.gif>.)

The creation of the Mosaic browser by Marc Andreessen in 1993 followed by the creation of Netscape early in 1994 were the historic events that marked the beginning of the internet boom that lasted throughout the rest of the 1990s, and led to the mass uptake in web usage that continues to increase to this day. A screenshot of an early version of Mosaic is shown in Fig. 1.4.

1.3 BRIEF HISTORY OF SEARCH ENGINES

The roots of web search engine technology are in information retrieval (IR) systems, which can be traced back to the work of Luhn at IBM during the late 1950s [444]. IR has been an active field within information science since then, and has been given a big boost since the 1990s with the new requirements that the Web has brought.

Many of the methods used by current search engines can be traced back to the developments in IR during the 1970s and 1980s. Especially influential is the SMART (system for the mechanical analysis and retrieval of text) retrieval system, initially developed by Gerard Salton and his collaborators at Cornell University during the early 1970s [583]. An important treatment of the traditional approaches to IR was given by Keith van Rijsbergen [655], while more modern treatments with reference to the Web can be found in Refs 45, 68, 453, and 164. More recent developments, which concentrate on web technologies, are the probabilistic perspective on modeling the Web as in Ref. 46 and the data mining perspective on managing web information, which can be found in Refs 128 and 435.

Owing to the massive amount of information on the Web, right from the early days of the Web, search engines have become an indispensable tool for web users. A history of search engines detailing some of the early search services can be found in Ref. 659.³

Here, we will be very selective and mention only a few of the early and current search engines; see <http://searchenginewatch.com/links> and http://en.wikipedia.org/wiki/List_of_search_engines for up-to-date listings of the major search engines. More details on many of the current search engines are spread throughout the book.

- Yahoo (www.yahoo.com), which started up in February 1994, was one of the earliest search services.⁴ Initially, Yahoo was only providing a browsable directory, organizing web pages into categories which were classified by human editors. Yahoo continues to maintain a strong brand and has evolved into a full-fledged search engine by acquiring existing search engine technology in mid-2003. (You can get some insight on the latest innovations in Yahoo's search engine from its weblog at www.ysearchblog.com.)
- InfoSeek, which started up in July 1994, was the first search engine that I was using on a regular basis, and as with many of the innovative web tools, users voted with their clicks and its reputation spread by word of mouth. In July 1998, Infoseek merged with Walt Disney's Buena Vista Internet Group to form Go.com, which was ultimately abandoned in January 2001.
- Inktomi, which started up in September 1995, provides search engine infrastructure rather than delivering the service from their web site. Until it was acquired by Yahoo in March 2003, it was providing search services to some of the major search engines.
- AltaVista (www.altavista.com), which started up in December 1995, was the second search engine that I was using on a regular basis. It was initially a research project in Digital Equipment Corporation, and was eventually acquired by Overture in April 2003.
- AlltheWeb (www.alltheweb.com) was launched in May 1999 by Fast Search & Transfer, and in a very short time was able to build a very large and fresh index with fast and accurate search results. It was also acquired by Overture in April 2003.
- Ask Jeeves (www.ask.com) started up in April 1996. It went public in July 1999, and is one of the survivors in the search engine game. Its strong brand and distinctive question answering facility have evolved into a general search service through its acquisition of Teoma in September 2001, which has enabled it to manage a proprietary search service and develop its own search technology. It was acquired by e-commerce conglomerate IAC (InterActiveCorp) in July 2005.

³See also, A history of search engines, by W. Sonnenreich. www.wiley.com/legacy/compbooks/sonnenreich/history.html.

⁴The history of Yahoo!—How it all started. <http://docs.yahoo.com/info/misc/history.html>.

- Overture (www.overture.com) started up as Goto.com in September 1997, and pioneered pay-per-click search engine advertising. It was renamed as Overture in September 2001 and was acquired by Yahoo in July 2003. In April 2005, Overture was rebranded as Yahoo Search Marketing (<http://searchmarketing.yahoo.com>).
- Bing (www.bing.com) is Microsoft's search engine that went online in June 2009. It replaced Live search, released in September 2006, which replaced MSN search, originally launched in August 1995, coinciding with the release of Windows 95. Initially, MSN search partnered with major search engines to provide the search facility for their site. Realizing the strategic importance of search to Microsoft's core business, Microsoft announced, in 2003, that it would develop its own proprietary search technology. The beta version of the search engine was released by MSN in November 2004, and in February 2005 MSN search was officially delivering search results from its internally developed engine. (You can get some insight on the latest innovations in Bing's search engine from its weblog at www.bing.com/community/blogs/search.)
- Google (www.google.com) was started up in September 1998, by Larry Page and Sergey Brin, then PhD students at Stanford University.⁵ Google was the third search engine that I was using on a regular basis and am still using today, although I do consult other search services as well. It became a public company in August 2004, and, as of late 2004, has been the most popular search engine. You will find a wealth of information in this book on the innovative features that Google and other search engines provide. (You can get some insight on the latest innovations in Google's search engine from its weblog at <http://googleblog.blogspot.com>.)

⁵Google History. www.google.com/corporate/history.html.