



Semantic Web Programming

John Hebler
Matthew Fisher
Ryan Blace
Andrew Perez-Lopez

*Foreword by Mike Dean,
Principal Engineer, BBN Technologies*



Semantic Web Programming



Semantic Web Programming

John Hebler
Matthew Fisher
Ryan Blace
Andrew Perez-Lopez



WILEY

Wiley Publishing, Inc.

Semantic Web Programming

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2009 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-41801-7

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data is available from the publisher.

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 Sections 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, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. 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/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Web site may provide or recommendations it may make. Further, readers should be aware that Internet Web sites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc. is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

To my wife, Christi, who for twenty-five years continues to offer support, wisdom, and love while putting up with my innate geekiness.

And to my dad, John, who gave me the gift of curiosity. Thank you!

—John Hebler

To Brenna, Denny, Brody, Mallory, Grace, and Olivia: You had patience in a father whose playtime and energy slipped while I wrote this book, but your love never faltered—you are each a blessing. To Erin, my wife, who had the world at her feet and still chose to be with me. You make this world a better place to be. I am the luckiest.

—Matthew Fisher

To my parents, Jorge and Kathleen; to my siblings, Dan, Tom, Anya, and Tonya; and to Erika. Without your love, patience, and support I could never have written this book. Thank you!

—Andrew Perez-Lopez

To my beautiful and infinitely patient wife Luci, for allowing me to spend nights and weekends writing this book. And to my "kids," Daisy, Mini, Midas, India, and Lily, for providing plenty of mental health breaks.

—Ryan Blace

To my wife, Nancy, and my sons, Jason and Noah, for allowing me the time to review chapters.

—Mike Dean



About the Authors

John Hebler is an avid, aging, yet still excited explorer of new technologies for the development of large-scale, distributed systems. In the last five years, he has focused on the Semantic Web and emergent, distributed systems. He has published several papers, has co-written a P2P networking book, and presents at major technical conferences around the world. He is currently pursuing his PhD in Information Systems at the University of Maryland. He is a division scientist for BBN technologies.

Matthew Fisher has over fifteen years experience in the fields of software and systems development. He has worked in a wide range of engineering environments, ranging from small technology startups and research and development companies to large Fortune 50 firms. He regularly contributes to the Semantic Report and has been involved with conferences such as OWLED, ISWC, and the Semantic Technology Conference. Matthew is a principal systems engineer at Progeny Systems and holds a BS in Computer Science from Penn State University and a MS in Computer Science from George Mason University.

Andrew Perez-Lopez is a software developer who has worked at BBN Technologies since 2005 on large-scale information integration systems using the Semantic Web technologies discussed in this book. He holds an MS in Computer Science from Virginia Tech and an BA in Cognitive Science from the University of Virginia.

Ryan Blace has been a Semantic Web developer and BBN Technologies employee for five years. He works on multiple large-scale Semantic Web-based knowledge management systems for the government and commercial sectors. Ryan holds a BS in Computer Engineering from the University of Maryland and is pursuing his master's in Computer Science at Virginia Tech.

When not spending late nights on his computer hacking away, Ryan spends his time cycling, mountain biking, and instructing at car club track days.

About The Technical Editors

Mike Dean (reviewer/editor) is principal engineer at BBN Technologies, where he has worked since 1984. He started working with Semantic Web in 2000, as a principal investigator in the DARPA Agent Markup Language (DAML) program. He was co-editor of the W3C OWL Reference, a co-author of SWRL, and has developed various Semantic Web tools, data sets, and applications. He currently provides technical direction for a number of Semantic Web projects at BBN. He holds a BS in Computer Engineering from Stanford University.

Mike Smith (technical editor) is a senior engineer at Clark & Parsia LLC, a software development and consulting firm specializing in the development and application of artificial intelligence technologies. Mike is a member of the W3 OWL WG, participates actively in the OWL community, and publishes content at <http://clarkparsia.com/weblog/>. He is one of the primary developers of Pellet, the open-source OWL reasoner, and frequently contributes to the Protégé and OWL API projects. He holds BS and MS degrees in Systems and Information Engineering from the University of Virginia.



Executive Editor

Robert Elliott

Development Editor

Christopher J. Rivera

Technical Editor

Michael Smith

Production Editor

Melissa Lopez

Copy Editor

Linda Recktenwald

Editorial Manager

Mary Beth Wakefield

Production Manager

Tim Tate

**Vice President and Executive
Group Publisher**

Richard Swadley

**Vice President and Executive
Publisher**

Barry Pruett

Associate Publisher

Jim Minatel

Proofreader

Corina Copp and Jen Larsen,
Word One

Indexer

Ted Laux

Cover Image

Tony Sweet/ Digital Vision/
Getty Images



Acknowledgments

The idea for this book grew over two years with the support of many BBN folks but especially Pete Pflugrath, our Semantic Web visionary; Ted Benson, an all-round awesome dude who motivated us to take on this challenge; Dana Moore, whose ideas and enthusiasm are simply limitless; and Mike Dean, whose boundless knowledge and expertise in all things technological is simply an inspiration to us all.

Strong support went well beyond BBN to include Walt Kitonis, Mike MacKay and Fred Vignovich of Progeny Systems and Gary Sikora for his advocacy of Semantic Web solutions in industry. Also to Tom Dietz, vice president of iJet, a truly rare and special person whose confidence in our abilities never wavered even when ours did.

A special thank-you to Mike Smith for his detailed technical reviews that gave the book its high quality and for keeping us on the leading edge of the rapidly advancing Semantic Web. And thanks to all the folks at Wiley publishing, especially Bob Elliott (our executive editor), whose initial belief in the project made it all possible, and Christopher Rivera (our editor), whose patience and whip kept us in line and writing throughout the entire process.



Contents

Foreword	xxiii
Introduction	xxv
Part One Introducing Semantic Web Programming	1
Chapter 1 Preparing to Program a Semantic Web of Data	3
Defining the Semantic Web	4
Identifying the Major Programming Components	10
Determining Impacts on Programming	13
Establishing a Web Data–Centric Perspective	13
Expressing Semantic Data	14
Sharing Data	16
Making Data Dynamic and Flexible	18
Avoiding the Roadblocks, Myths, and Hype	19
Semantic Web Roadblocks	19
Semantic Web Myths	21
Semantic Web Hype	22
Understanding Semantic Web Origins	23
Exploring Semantic Web Examples	26
Semantic Wikis (<i>semantic-mediawiki.org</i>)	26
Twine (<i>www.twine.com</i>)	27
The FOAF Project (<i>www.foaf-project.org</i>)	28
RDFa and Microformats	30
Semantic Query Endpoint (<i>dbpedia.org/sparql</i>)	32
Semantic Search (<i>www.trueknowledge.com</i>)	32
Summary and Onward	34
Notes	34

Chapter 2	Hello Semantic Web World	35
	Setting Up Your Semantic Web Development Environment	36
	Programming the Hello Semantic Web World Application	38
	Summary	58
Part Two	Foundations of Semantic Web Programming	61
Chapter 3	Modeling Information	63
	Modeling Information in Software	64
	Sharing Information: Syntax and Semantics	65
	Serialized Objects	66
	Relational Databases	66
	Extensible Markup Language (XML)	66
	Metadata and Data in Information Sharing	67
	The Semantic Web Information Model: The Resource	
	Description Framework (RDF)	68
	Nodes: Resources and Literals	69
	Edges: Predicates	71
	Exchanging Information with RDF	72
	Statements as Points	73
	RDF Serializations	74
	RDF/XML	74
	Terse RDF Triple Language (Turtle)	78
	N-Triples	81
	Quick Hack	82
	More RDF	84
	Blank Nodes	84
	Reification	88
	RDF Organizational Constructs	88
	Summary	91
Chapter 4	Incorporating Semantics	93
	Semantics on the Web	94
	Motivating Factors	94
	Understanding the World Wide Web	95
	Knowledge Domain Integration	97
	Expressing Semantics in RDF	98
	Vocabularies, Taxonomies, and Ontologies	99
	A Vocabulary Language for RDF	100
	An Ontology Language for the Web	101
	Introduction to Ontologies	102
	Distributed Knowledge	102
	Open World Assumption	103
	No Unique Names Assumption	104
	Overview of Ontology Elements	104

Ontology Header	105
Classes and Individuals	105
Properties	106
Annotations	106
Data types	106
Elements of an Ontology	107
OWL 2 Typing	107
Ontology Header	108
Annotations	109
Basic Classification	110
Classes and Individuals	110
<i>rdfs:SubClassOf</i>	111
Instance versus Subclass	112
<i>owl:Thing</i> and <i>owl:Nothing</i>	113
Defining and Using Properties	113
Property Domain and Range	114
Describing Properties	115
<i>rdfs:subPropertyOf</i>	115
Top and Bottom Properties	116
Inverse Properties	116
Disjoint Properties	117
Property Chains	118
Symmetric, Reflexive, and Transitive Properties	119
Functional and Inverse Functional Properties	120
Keys	121
Datatypes	122
Data type Restrictions	122
Defining Datatypes in Terms of Other Datatypes	124
Negative Property Assertions	126
Property Restrictions	127
Value Restrictions	127
Cardinality Restrictions	130
Qualified Cardinality Restrictions	132
Advanced Class Description	134
Enumerating Class Membership	134
Set Operators	134
Disjoint Classes	136
Equivalence in OWL	137
Equivalence among Individuals	138
Equivalence among Classes and Properties	139
Summary	139
Chapter 5 Modeling Knowledge in the Real World	141
Exploring the Components of the Semantic Web	141
Semantic Web Frameworks	143

Storing and Retrieving RDF	144
RDF Store Implementations	144
Retrieving Information in a Knowledgebase	146
Realizing the Semantics of OWL	147
Understanding Forward Chaining Inference	148
Understanding Backward Chaining Inference	152
Choosing the Right Inference Method	153
Common Frameworks and Components	153
RDF Store Implementations	154
Retrieval Components	156
Reasoning Engines	156
Knowledgebase Performance	157
Exploring the Profiles of OWL	158
OWL Full and OWL DL	159
The Profiles of OWL	160
OWL EL	160
OWL QL	161
OWL RL	162
Demonstrating OWL Inference	163
The Ontology	163
The Example Application	165
The Results	169
Performing No Inference	169
Performing RDFS Inference	170
Performing OWL Inference	171
Working with Ontologies	172
Decoupling the Knowledge Model from the Application	173
Sharing across Domain and Application Boundaries	174
What Is a Foundational Ontology?	175
Common Foundational Ontologies	177
BFO	177
Cyc and OpenCyc	178
DOLCE	179
SUMO	179
Dublin Core Metadata Initiative	179
FOAF	180
GeoRSS and OWL-Time	180
Finding Ontologies to Reuse or Extend	181
Choosing the Right Foundational Ontologies	183
Summary	184
Chapter 6 Discovering Information	185
Navigating the Semantic Web	186
Searching the Semantic Web	190

Querying the Semantic Web	192
Quickstart with SPARQL	192
Four Foundational Query Forms	196
SELECT Essentials	197
RDF Datasets, FROM and FROM NAMED	202
Query Modifiers	208
DISTINCT	209
REDUCED	210
ORDER BY	210
Data Streaming with OFFSET and LIMIT	211
Flexible Querying with FILTER and OPTIONAL	213
FILTER	213
OPTIONAL	215
UNION	219
CONSTRUCT Essentials	222
DESCRIBE Essentials	224
ASK Essentials	225
SPARQL Entailment	226
Unsupported Functionality	228
Data Modification	228
Subqueries	228
Aggregation	228
Summary	229
Chapter 7 Adding Rules	231
What Are Rules?	232
Reasons for Rules	232
No Support for Property Composition	232
Use of Built-ins	232
Ontological Mediation	233
Limiting Assumptions	233
Rule Languages	233
SWRL Essentials	234
The Abstract Syntax	235
The XML Concrete Syntax	237
var	238
imp	239
_rlab	239
_body	239
_head	240
classAtom	240
datarangeAtom	241
individualPropertyAtom	241
datavaluedPropertyAtom	242
sameIndividualAtom	242

differentIndividualsAtom	243
builtinAtom	243
The RDF Concrete Syntax	243
Built-ins	244
Examples	244
DL-Safe Rules	245
Ontological Mediation	252
Mapping Friends without Upsetting Any of Them	253
The Power of Rules	256
Jena Rules	257
Rule Interchange Format	259
Delving into the Details	260
The Future of RIF	260
Summary	261
Part Three Building Semantic Web Applications	263
Chapter 8 Applying a Programming Framework	265
Framing the Semantic Web	266
The Jena Semantic Web Framework	269
Defining Jena Programming Concepts	269
Programming with Jena	273
Establishing the Jena Development Environment	276
Establishing the Knowledgebase: Setting Up the Model	276
Populating the Model with Semantic Web Data	279
Combining Semantic Web Data	282
Interrogating Semantic Web Data	283
Reasoning across Semantic Web Data	285
Exporting Semantic Web Data	290
Deallocating Semantic Web Data Resources	290
Managing Semantic Web Data	291
Getting Information Regarding Your Semantic Web Data	291
Generating Events Based on Semantic Web Data	292
Dealing with Concurrency and Your Semantic Web Data	293
Customizing the Jena Framework	295
Serializing Semantic Web Data	297
Common App Overview: FriendTracker	298
Summary	299
Chapter 9 Combining Information	301
Combining Information	303
Representing Information	303
Translating between Representations	304
Addressing the Challenges of Translation	305

Maintaining Fidelity	305
Tracking Provenance Information	305
Reversing the Process	306
Handling Varying Data	306
Managing Data Volume	306
Introducing the FriendTracker Data Sources	307
Facebook XML Web Service	307
Jabber Java Client	308
Upcoming.org XML Web Service	308
WordPress Relational Database	308
Exposing XML-Based Web Services as RDF	309
Introducing the Weather.gov XML Feed	310
Exposing XML Using XSL Transformations	311
Traversing XML Documents with XPath	313
Applying XSLT to a Simple Example	315
Processing XML and XSLT Programmatically	319
Applying XSLT to the Facebook Data Source	322
Weighing the Benefits and the Costs of XSLT	326
Exposing XML Using XML Bindings and Velocity	328
Generating Java Bindings for XML Data	329
Unmarshalling XML Data into Java	331
Introducing the Velocity Template Engine	333
Generating RDF with Velocity	336
Weighing the Benefits and the Costs	336
Exposing Relational Databases as RDF	337
Exposing a WordPress Blog Using D2RQ	338
Creating D2RQ Mappings for the WordPress Database	339
Wrapping the D2RQ Instance in a Jena Model	341
Querying the D2RQ Exposed WordPress Database	342
Weighing the Benefits and the Costs of D2RQ	345
Exposing Other Sources of Data	346
Exposing Jabber with a Custom Streaming RDF Writer	346
Exposing Java Objects Using Reflection	352
Applying the RDF Generator to the Weather.gov XML Feed	356
Applying the RDF Generator to the Upcoming.org XML Feed	358
Summary	359

Chapter 10 **Aligning Information** **361**

Data Source, Domain, and Application Ontologies	361
Aligning Ontologies	362
Ontology Constructs	363
Translation via Rules	365

Explicit Translation	365
Ad Hoc Approaches to Translation	366
FriendTracker	366
Aligning Ontologies with OWL and SWRL	372
Aligning Ontologies with XSLT	376
Aligning Ontologies with Code	381
Aligning Simple Ontologies with RDFS	382
Record Linkage	385
Summary	388
Chapter 11 Sharing Information	389
Microformats	390
eRDF	392
RDFa	395
Supported Attributes	396
xmlns	396
rel	396
rev	398
content	398
href	399
src	399
about	399
property	399
resource	400
datatype	401
typeof	401
Blank Nodes	402
Language Support	403
Tools and Frameworks	404
RDF Transformational Tools	404
SPARQL Endpoints	404
Joseki Installation and Operation	405
xOperator	408
Installation and Operation	409
Example Query	410
FriendTracker in RDFa	411
Summary	417
Part Four Expanding Semantic Web Programming	419
Chapter 12 Developing and Using Semantic Services	421
Background	422
Discovery	424
Invocation	424
Negotiation	425

Error Handling	425
Monitoring	425
Composition	425
Implementing Semantic Services	426
Semantic Markup for Web Services	427
ServiceProfile	428
ServiceModel	428
ServiceGrounding	428
Web Service Modeling Ontology	429
Semantic Annotations for WSDL	432
SAWSDL Example	433
SAWSDL Tools	434
Summary	435
Chapter 13 Managing Space and Time	437
Space and Time in Software	437
Spatial Information	438
Temporal Information	440
Representing Spatiotemporal Data on the Semantic Web	441
Spatial and Temporal Software with Jena	450
Working with Spatial Data	452
Example: Spatial Queries	453
Framing the Problem	453
Approach and Rationale	454
Components	454
Example: Transaction Time–Bounded Queries	461
Framing the Problem	461
Approach and Rationale	462
Components	462
Summary	465
Notes	465
Chapter 14 Semantic Web Patterns and Best Practices	467
Aggregating Disparate Data Sources	468
Exposing Data Sources as RDF	469
Bringing Data into the Domain Knowledge Model	470
Storing Information in the Knowledgebase	472
Initiating the Flow of Data	472
Annotating Unstructured Data	472
Annotation Management	474
Ontology Management	475
Unstructured Data Sources and the Client Application	476
Coordinating Semantic Services	478
Applying Semantic Web Best Practices	480

Creating URIs	480
Making URIs Unique	480
Making URIs Consistent	481
Making URIs Resolvable	483
Specifying Units of Measurement	483
Unit-Specific Properties and Datatypes	484
Statement Reification	485
Value Containers	485
Representing N-ary Relationships	486
Managing Bad Data	487
Summary	487
Chapter 15 Moving Forward	489
Advancing Ontologies	491
Ontology Repositories and Registries	491
Linked Data	492
Versioning	493
Ontology Metrics	495
Advancing Integration	497
Semantic Pipes	497
Distributed Queries	498
Alignment	500
Advancing Reasoning	506
Rule Interchange Format (RIF)	506
Probabilistic Reasoning	507
Trust: Proof Markup Language	510
LarKC: The Large Knowledge Collider	512
Advancing Visualization	513
Summary	516
Appendix A RDF	519
Appendix B The OWL Web Ontology Language	525
Appendix C SWRL	533
Appendix D SPARQL	555
Appendix E Jena Reference Guide	567
Appendix F Installation Reference Guide	577
Index	585



Foreword

Our group at BBN Technologies has been working at the forefront of the Semantic Web since 2000, first as part of the DARPA Agent Markup Language (DAML) program and then in developing a variety of tools, data sets, and applications for other government and commercial customers. The authors and technical editor of this book are current or former members of this group, which has grown to about 30 employees. *Semantic Web Programming* reflects our backgrounds as software developers, the experience we've gained over the past eight years, and a number of hard-won insights.

The Semantic Web is an international effort to represent data (including World Wide Web data currently designed for human users) in formats amenable to automated processing, integration, and reasoning. Data is king, and it provides even greater value when it's connected with other data sources to create a linked data web. Current applications include data integration from mash-ups to the enterprise, improved search, service composition, intelligent agents, desktop and mobile applications, and collaboration.

Catalyzed by U.S. and EU research programs, the growing community includes the W3C Semantic Web Activity, a host of large and small vendors, several Semantic Web and Semantic Technology conference series, and a large number of open-source developers and projects.

While Web 3.0 is in many ways an appropriate moniker for the Semantic Web, the Semantic Web has always emphasized Web 2.0 social networking and collaboration aspects through FOAF, RSS 1.0, various semantic wiki projects, and participatory collections such as MusicBrainz. Semantic Web ontologies provide more structure than Web 2.0 tags, microformats, and folksonomies, while retaining much of their flexibility.

Semantic Web standards including RDF, OWL, and SPARQL continue to evolve based on usage. A wide range of high-quality tools, many of them

open source, have been developed for different programming environments. The Linking Open Data initiative has addressed a critical need by providing foundational data for many applications and continues to grow. Many tools and applications are now highly scalable.

Developers often benefit from seeing other people's code. Throughout this book, we've taken a pragmatic approach, with lots of examples and an application that spans multiple chapters.

We hope that you'll also find that Semantic Web technologies provide an effective means of addressing current and upcoming computing challenges and that you'll enjoy working with them as much as we have.

Mike Dean

Ann Arbor, Michigan

November 2008



Introduction

Semantic Web Programming takes the Semantic Web directly and boldly into solving practical, real-world problems that flexibly deliver real value from our growing ability to access information and services from our laptop to the enterprise to the World Wide Web. The chapters form a solid, code-based path addressing information and service challenges. As the code examples build, we pragmatically explore the many technologies that form the Semantic Web, including the knowledge representations such as microformats, Resource Description Framework (RDF), RDF Schema (RDFS), the Web Ontology Language (OWL) including its latest release OWL 2 and Semantic Web Rule Language (SWRL), Semantic Web programming frameworks such as Jena, and useful Semantic Web tools. We explore these technologies, not as ends in themselves but rather for their role and merits in solving real problems. Thus, your learning is based on results—the results that each technology brings to address your application challenges.

Semantic Web Programming benefits from our many years of experience in developing large-scale Semantic Web solutions, building Semantic Web tools, and contributing to the Semantic Web standards. We know this stuff! This background provides you with not only an understanding of this new powerful technology but the ability to apply it directly to your real-world application and information challenges.

Overview of the Book and Technology

The Semantic Web offers a powerful, practical approach to gain mastery over the multitude of information and information services. Semantics offer the leverage to make *more* information *better* and not overwhelmingly worse. This

requires new data representations that improve our ability to capture and share knowledge and new programming constructs and tools to make this information work for your application.

This book explores it all through actual data formats, working code, and tools. We take a *developer* perspective aimed at application results. We focus the explanations and justifications on what you need to build and manage your Semantic Web applications. The multitude of working code examples throughout the book provides the credibility and insights that truly augment the background and explanatory text. In many cases, the code does the talking. We strongly recommend that you get *hands on* and adjust the examples to your needs. This will help you gain the understanding and perspective necessary to put the Semantic Web to work for you immediately.

How This Book Is Organized

The book has 15 chapters organized in four parts. Also included is an extensive set of references in the appendices for the key technologies.

Part 1: “Introducing Semantic Web Programming,” covers Chapters 1 and 2. This section quickly introduces you to Semantic Web programming. Chapter 1, “Preparing to Program a Semantic Web of Data,” covers the main Semantic Web concepts and their relationship with one another. This establishes your Semantic Web developer vocabulary. Chapter 1 also points out the advantages and programming impacts; it ends with some compelling examples of the Semantic Web in use today. Chapter 2, “Hello Semantic Web World,” dives right into working code with an exhaustive Hello Semantic World Web program. The example takes you from setting up your development environment to using reasoners. The explanations are brief because this chapter is merely an introduction to the rest of the book. This section is critical if you are new to the Semantic Web. Seasoned readers may choose to skim these two chapters.

Part 2, “Foundations of Semantic Web Programming,” covers Chapters 3 through 7. Two main areas drive a Semantic Web application: knowledge representation and application integration. This section focuses on the former—representing and manipulating knowledge. Chapter 3, “Modeling Information,” establishes the data model through RDF. Chapter 4, “Incorporating Semantics,” adds an ontology to create a knowledge model using RDFS and OWL 2. Chapter 5, “Modeling Knowledge in the Real World,” exercises the working ontology via application frameworks and reasoners. Chapter 6, “Discovering Information,” dives into the knowledge model to extract useful information through search, navigation, and formal queries via SPARQL. Chapter 7, “Adding Rules,” rounds out the knowledge representation through an exploration of the semantic rule languages, including the W3C standard SWRL.

Part 3, “Building Semantic Web Applications,” covers Chapters 8 through 11. This section deals with the second main area—integrating the knowledgebase with an application that acts upon it. This part provides a solid programming base for the Semantic Web. Chapter 8, “Applying a Programming Framework,” fully explores Semantic Web frameworks with extensive examples from the Jena Semantic Web Framework. The chapter ends with an outline of our FriendTracker Semantic Web application. This example spans the next three chapters as we explore methods to integrate, align, and output data and information in many formats and locations. Chapter 9, “Combining Information,” focuses on integrating the information into a knowledge model from sources such as relational databases, web services, and other formats. Chapter 10, “Aligning Information,” focuses on aligning the data along ontological concepts to unify the disparate information. Chapter 11, “Sharing Information,” outputs the information into many formats, including RDFa, microformats, SPARQL endpoints, and more. All along we add to the FriendTracker application to directly demonstrate the programming concepts.

Part 4, “Expanding Semantic Web Programming,” covers chapters 12 through 15. Here we build on your solid base of knowledge representation and Semantic Web application development to expand into powerful, useful areas, including semantic services, time and space, Semantic Web architectures and best practices, and unfolding Semantic Web tools that are almost here. Chapter 12, “Developing and Using Semantic Services,” adds semantics to services to allow them to participate in the Semantic Web. Chapter 13, “Managing Space and Time,” adds space and time considerations to your knowledge representations. Chapter 14, “Applying Patterns and Best Practices,” is a retrospective of sorts. It builds on everything we covered so far in the book by presenting a series of architecture patterns for constructing various Semantic Web applications. Chapter 15, “Moving Forward,” concludes the book by peering into the future. It focuses on four critical, evolving areas for the Semantic Web: ontology management, advanced integration and distribution, advanced reasoning, and visualization. This provides a solid view into what is on its way in the actively evolving Semantic Web.

Who Should Read This Book

The book provides a comprehensive, practical view for developing applications that use the Semantic Web. The Semantic Web takes advantage of the multitude of distributed information and services that exist in the World Wide Web, the business enterprise, and your personal resources. Therefore, many technical readers would benefit from this book whether you focus on the entire application or only the information.

Developers gain first-hand experience with the many code examples throughout the book. These include both applications developers and information developers who focus on data in its many forms, from database schemas to XML formats. This book provides all the tools, background, and rich examples to jump-start your applications.

Architects gain insights into the role of the Semantic Web within a larger application. The Semantic Web offers many benefits to any system that uses information—which is just about any system—and can quickly extend your system’s capabilities to better leverage available information and services. The overall applications serve the system architect, whereas the detailed information and data management areas benefit information architects responsible for data formats and data processing.

Technical management gains insight into the power, risks, and benefits of the Semantic Web. The Semantic Web is a strategic technology—one that truly provides a solution with a significant advantage. It offers a new approach to extremely tough but lucrative challenges that employ vast amounts of information and services. Awareness of the Semantic Web is required for any solution that depends on dynamic information and service resources. The code examples provide credibility to the technology and insights into its own challenges for better planning.

Tools You Will Need

We highly recommend that you reinforce your learning by downloading and customizing the numerous coding examples throughout the book. All the software tools are open source and readily available from the World Wide Web. We include all necessary links and instructions. Your computer is compatible with all of these tools as long as your operating system supports a Java 1.5 virtual machine. That’s it! As we cover each tool in the book, we provide download, installation, and configuration instructions. In addition, we summarize all the tools with instructions in Appendix F.

What’s on the Website

The book comes with an extensive website companion at <http://semwebprogramming.org>. Here you can access all related articles, complete code examples, and ontologies, as well as have an opportunity to get involved in the ongoing discussions and activities. The site also contains any book and code updates to reflect the continual expansion and evolution of the Semantic Web. We welcome comments on the book and examples.