

2
eBooks

Addison-Wesley Professional Ruby Series

RUBY ON RAILS TUTORIAL

LEARN RAILS™ BY EXAMPLE

MICHAEL HARTL
FOREWORDS BY DEREK SIVERS AND
DAVID HEINEMEIER HANSSON

Addison-Wesley Professional Ruby Series

THE RAILS™ 3 WAY

Forewords by David Heinemeier Hansson, creator of Ruby on Rails
and Yehuda Katz, Rails Core

OBIE FERNANDEZ
with CURRAN JORDAN, JON LARKOWSKI,
XAVIER NORIA, and TIM POPE

THE **RUBY ON RAILS 3**
TUTORIAL & REFERENCE
COLLECTION

The Ruby on Rails 3 Tutorial and Reference Collection

◆◆Addison-Wesley

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Cape Town • Sydney • Tokyo • Singapore • Mexico City

Note from the Publisher

The *Ruby on Rails 3 Tutorial and Reference Collection* consists of two bestselling Rails eBooks:

- *Ruby on Rails 3 Tutorial: Learn Rails by Example* by Michael Hartl
- *The Rails 3 Way* by Obie Fernandez

Ruby on Rails 3 Tutorial: Learn Rails by Example is a hands-on guide to the Rails 3 environment. Through detailed instruction, you develop your own complete sample application using the latest techniques in Rails Web development. *The Rails 3 Way* addresses real challenges development teams face, showing how to use Rails 3 to maximize your productivity. It is an essential reference for professional developers to deliver production-quality code using Rails 3.

To simplify access to each book, we've appended "A" to pages of *Ruby on Rails 3 Tutorial* and "B" to pages of *The Rails 3 Way*. This enabled us to produce a single, comprehensive table of contents and dedicated indexes.

We hope you find this collection useful!

—The editorial and production teams at Addison-Wesley Professional

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales
international@pearson.com

Visit us on the Web: informit.com/aw

Ruby on Rails 3 Tutorial copyright © 2011 Michael Hartl

The Rails 3 Way copyright © 2011 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission to use material from this work, please submit a written request to Pearson Education, Inc., Permissions Department, One Lake Street, Upper Saddle River, New Jersey 07458, or you may fax your request to (201) 236-3290.

ISBN-13: 978-0-13-292800-7

ISBN-10: 0-13-292800-0

Table of Contents

RUBY ON RAILS 3 TUTORIAL

1	From Zero to Deploy	.1A
1.1	Introduction	.3A
1.1.1	Comments for Various Readers	.4A
1.1.2	“Scaling” Rails	.7A
1.1.3	Conventions in This Book	.7A
1.2	Up and Running	.9A
1.2.1	Development Environments	.9A
1.2.2	Ruby, RubyGems, Rails, and Git	.11A
1.2.3	The First Application	.15A
1.2.4	Bundler	.16A
1.2.5	<code>rails server</code>	.20A
1.2.6	Model-View-Controller (MVC)	.22A
1.3	Version Control with Git	.24A
1.3.1	Installation and Setup	.24A
1.3.2	Adding and Committing	.26A
1.3.3	What Good Does Git Do You?	.28A
1.3.4	GitHub	.29A
1.3.5	Branch, Edit, Commit, Merge	.31A
1.4	Deploying	.35A
1.4.1	Heroku Setup	.36A
1.4.2	Heroku Deployment, Step One	.37A
1.4.3	Heroku Deployment, Step Two	.37A
1.4.4	Heroku Commands	.39A
1.5	Conclusion	.40A
2	A Demo App	.41A
2.1	Planning the Application	.41A
2.1.1	Modeling Users	.43A
2.1.2	Modeling Microposts	.44A
2.2	The Users Resource	.44A
2.2.1	A User Tour	.46A
2.2.2	MVC in Action	.49A
2.2.3	Weaknesses of This Users Resource	.58A

2.3	The Microposts Resource	.58A
2.3.1	A Micropost Microtour	.58A
2.3.2	Putting the <i>micro</i> in Microposts	.61A
2.3.3	A User <i>has_many</i> Microposts	.63A
2.3.4	Inheritance Hierarchies	.66A
2.3.5	Deploying the Demo App	.68A
2.4	Conclusion	.69A
3	Mostly Static Pages	.71A
3.1	Static Pages	.74A
3.1.1	Truly Static Pages	.75A
3.1.2	Static Pages with Rails	.78A
3.2	Our First Tests	.84A
3.2.1	Testing Tools	.84A
3.2.2	TDD: Red, Green, Refactor	.86A
3.3	Slightly Dynamic Pages	.103A
3.3.1	Testing a Title Change	.103A
3.3.2	Passing Title Tests	.106A
3.3.3	Instance Variables and Embedded Ruby	.108A
3.3.4	Eliminating Duplication with Layouts	.112A
3.4	Conclusion	.115A
3.5	Exercises	.116A
4	Rails-Flavored Ruby	.119A
4.1	Motivation	.119A
4.1.1	A <code>title</code> Helper	.119A
4.1.2	Cascading Style Sheets	.122A
4.2	Strings and Methods	.125A
4.2.1	Comments	.125A
4.2.2	Strings	.126A
4.2.3	Objects and Message Passing	.129A
4.2.4	Method Definitions	.132A
4.2.5	Back to the <code>title</code> Helper	.133A
4.3	Other Data Structures	.134A
4.3.1	Arrays and Ranges	.134A
4.3.2	Blocks	.137A

4.3.3 Hashes and Symbols	139A
4.3.4 CSS Revisited	142A
4.4 Ruby Classes	144A
4.4.1 Constructors	144A
4.4.2 Class Inheritance	145A
4.4.3 Modifying Built-In Classes	148A
4.4.4 A Controller Class	150A
4.4.5 A User Class	152A
4.5 Exercises	154A
5 Filling in the Layout	157A
5.1 Adding Some Structure	157A
5.1.1 Site Navigation	159A
5.1.2 Custom CSS	164A
5.1.3 Partials	171A
5.2 Layout Links	177A
5.2.1 Integration Tests	178A
5.2.2 Rails Routes	181A
5.2.3 Named Routes	183A
5.3 User Signup: A First Step	186A
5.3.1 Users Controller	186A
5.3.2 Signup URL	188A
5.4 Conclusion	191A
5.5 Exercises	191A
6 Modeling and Viewing Users, Part I	193A
6.1 User Model	194A
6.1.1 Database Migrations	196A
6.1.2 The Model File	201A
6.1.3 Creating User Objects	203A
6.1.4 Finding User Objects	207A
6.1.5 Updating User Objects	208A
6.2 User Validations	210A
6.2.1 Validating Presence	210A
6.2.2 Length Validation	217A
6.2.3 Format Validation	218A
6.2.4 Uniqueness Validation	222A

6.3 Viewing Users	227A
6.3.1 Debug and Rails Environments	227A
6.3.2 User Model, View, Controller	230A
6.3.3 A Users Resource	232A
6.4 Conclusion	236A
6.5 Exercises	237A
7 Modeling and Viewing Users, Part II	239A
7.1 Insecure Passwords	239A
7.1.1 Password Validations	240A
7.1.2 A Password Migration	244A
7.1.3 An Active Record Callback	247A
7.2 Secure Passwords	250A
7.2.1 A Secure Password Test	251A
7.2.2 Some Secure Password Theory	252A
7.2.3 Implementing <code>has_password?</code>	254A
7.2.4 An Authenticate Method	258A
7.3 Better User Views	262A
7.3.1 Testing the User Show Page (With Factories)	263A
7.3.2 A Name and A Gravatar	268A
7.3.3 A User Sidebar	276A
7.4 Conclusion	279A
7.4.1 Git Commit	279A
7.4.2 Heroku Deploy	280A
7.5 Exercises	280A
8 Sign Up	283A
8.1 Signup Form	283A
8.1.1 Using <code>form_for</code>	286A
8.1.2 The Form HTML	288A
8.2 Signup Failure	292A
8.2.1 Testing Failure	292A
8.2.2 A Working Form	295A
8.2.3 Signup Error Messages	299A
8.2.4 Filtering Parameter Logging	303A
8.3 Signup Success	305A
8.3.1 Testing Success	305A
8.3.2 The Finished Signup Form	308A

8.3.3	The Flash	.308A
8.3.4	The First Signup	.312A
8.4	RSpec Integration Tests	.313A
8.4.1	Integration Tests with Style	.315A
8.4.2	Users Signup Failure Should not Make a New User	.315A
8.4.3	Users Signup Success Should Make a New User	.319A
8.5	Conclusion	.321A
8.6	Exercises	.321A
9	Sign In, Sign Out	.325A
9.1	Sessions	.325A
9.1.1	Sessions Controller	.326A
9.1.2	Signin Form	.328A
9.2	Signin Failure	.332A
9.2.1	Reviewing form Submission	.333A
9.2.2	Failed Signin (Test and Code)	.335A
9.3	Signin Success	.338A
9.3.1	The Completed create Action	.338A
9.3.2	Remember Me	.340A
9.3.3	Current User	.345A
9.4	Signing Out	.354A
9.4.1	Destroying Sessions	.354A
9.4.2	Signin Upon Signup	.356A
9.4.3	Changing the Layout Links	.358A
9.4.4	Signin/Out Integration Tests	.362A
9.5	Conclusion	.363A
9.6	Exercises	.363A
10	Updating, Showing, and Deleting Users	.365A
10.1	Updating Users	.365A
10.1.1	Edit Form	.366A
10.1.2	Enabling Edits	.373A
10.2	Protecting Pages	.376A
10.2.1	Requiring Signed-In Users	.376A
10.2.2	Requiring the Right User	.379A
10.2.3	Friendly Forwarding	.382A

10.3 Showing Users	384A
10.3.1 User Index	385A
10.3.2 Sample Users	389A
10.3.3 Pagination	392A
10.3.4 Partial Refactoring	398A
10.4 Destroying Users	399A
10.4.1 Administrative Users	399A
10.4.2 The <code>destroy</code> Action	404A
10.5 Conclusion	408A
10.6 Exercises	409A
11 User Microposts	411A
11.1 A Micropost Model	411A
11.1.1 The Basic Model	412A
11.1.2 User/Micropost Associations	414A
11.1.3 Micropost Refinements	419A
11.1.4 Micropost Validations	423A
11.2 Showing Microposts	425A
11.2.1 Augmenting the User Show Page	426A
11.2.2 Sample Microposts	432A
11.3 Manipulating Microposts	434A
11.3.1 Access Control	436A
11.3.2 Creating Microposts	439A
11.3.3 A Proto-feed	444A
11.3.4 Destroying Microposts	452A
11.3.5 Testing the New Home Page	456A
11.4 Conclusion	457A
11.5 Exercises	458A
12 Following Users	461A
12.1 The Relationship Model	463A
12.1.1 A Problem with the Data Model (and a Solution)	464A
12.1.2 User/Relationship Associations	470A
12.1.3 Validations	473A
12.1.4 Following	474A
12.1.5 Followers	479A

12.2 A Web Interface for Following and Followers	482A
12.2.1 Sample Following Data	482A
12.2.2 Stats and a Follow Form	484A
12.2.3 Following and Followers Pages	494A
12.2.4 A Working Follow Button the Standard Way	498A
12.2.5 A Working Follow Button with Ajax	502A
12.3 The Status Feed	507A
12.3.1 Motivation and Strategy	508A
12.3.2 A First Feed Implementation	511A
12.3.3 Scopes, Subselects, and a Lambda	513A
12.3.4 The New Status Feed	518A
12.4 Conclusion	519A
12.4.1 Extensions to the Sample Application	520A
12.4.2 Guide to Further Resources	522A
12.5 Exercises	523A
Index	527A

THE RAILS 3 WAY

1 Rails Environments and Configuration	1B
1.1 Bundler	2B
1.1.1 Gemfile	3B
1.1.2 Installing Gems	5B
1.1.3 Gem Locking	7B
1.1.4 Packaging Gems	7B
1.2 Startup and Application Settings	8B
1.2.1 application.rb	8B
1.2.2 Initializers	11B
1.2.3 Additional Configuration	15B
1.3 Development Mode	15B
1.3.1 Automatic Class Reloading	16B
1.3.2 Whiny Nils	18B
1.3.3 Error Reports	18B
1.3.4 Caching	18B
1.3.5 Raise Delivery Errors	19B

1.4 Test Mode	19B
1.5 Production Mode	20B
1.5.1 Asset Hosts	22B
1.5.2 Threaded Mode	22B
1.6 Logging	23B
1.6.1 Rails Log Files	24B
1.6.2 Log File Analysis	26B
1.7 Conclusion	29B
2 Routing	31B
2.1 The Two Purposes of Routing	32B
2.2 The routes.rb File	33B
2.2.1 Regular Routes	34B
2.2.2 URL Patterns	35B
2.2.3 Segment Keys	36B
2.2.4 Spotlight on the :id Field	38B
2.2.5 Optional Segment Keys	38B
2.2.6 Constraining Request Methods	38B
2.2.7 Redirect Routes	39B
2.2.8 The Format Segment	40B
2.2.9 Routes as Rack Endpoints	41B
2.2.10 Accept Header	42B
2.2.11 Segment Key Constraints	43B
2.2.12 The Root Route	44B
2.3 Route Globbing	45B
2.4 Named Routes	46B
2.4.1 Creating a Named Route	46B
2.4.2 name path vs. name url	47B
2.4.3 What to Name Your Routes	48B
2.4.4 Argument Sugar	49B
2.4.5 A Little More Sugar with Your Sugar?	50B
2.5 Scoping Routing Rules	50B
2.5.1 Controller	51B
2.5.2 Path Prefix	51B
2.5.3 Name Prefix	52B
2.5.4 Namespaces	52B
2.5.5 Bundling Constraints	52B

2.6 Listing Routes53B
2.7 Conclusion54B

3 REST, Resources, and Rails55B

3.1 REST in a Rather Small Nutshell55B
3.2 Resources and Representations56B
3.3 REST in Rails57B
3.4 Routing and CRUD58B
 3.4.1 REST Resources and Rails59B
 3.4.2 From Named Routes to REST Support ..59B
 3.4.3 Reenter the HTTP Verb60B
3.5 The Standard RESTful Controller Actions61B
 3.5.1 Singular and Plural RESTful Routes62B
 3.5.2 The Special Pairs: new/create
 and edit/update63B
 3.5.3 The PUT and DELETE Cheat64B
 3.5.4 Limiting Routes Generated64B
3.6 Singular Resource Routes64B
3.7 Nested Resources65B
 3.7.1 RESTful Controller Mappings66B
 3.7.2 Considerations67B
 3.7.3 Deep Nesting?67B
 3.7.4 Shallow Routes68B
3.8 RESTful Route Customizations69B
 3.8.1 Extra Member Routes70B
 3.8.2 Extra Collection Routes72B
 3.8.3 Custom Action Names72B
 3.8.4 Mapping to a Different Controller72B
 3.8.5 Routes for New Resources73B
 3.8.6 Considerations for Extra Routes73B
3.9 Controller-Only Resources74B
3.10 Different Representations of Resources76B
 3.10.1 The respond to Method76B
 3.10.2 Formatted Named Routes77B
3.11 The RESTful Rails Action Set78B
 3.11.1 Index78B
 3.11.2 Show80B
 3.11.3 Destroy80B

3.11.4 New and Create	.81B
3.11.5 Edit and Update	.82B
3.12 Conclusion	.83B

4 Working with Controllers85B

4.1 Rack	.86B
4.1.1 Configuring Your Middleware Stack	.87B
4.2 Action Dispatch: Where It All Begins	.88B
4.2.1 Request Handling	.89B
4.2.2 Getting Intimate with the Dispatcher	.89B
4.3 Render unto View	.92B
4.3.1 When in Doubt, Render	.92B
4.3.2 Explicit Rendering	.93B
4.3.3 Rendering Another Action's Template	.93B
4.3.4 Rendering a Different Template Altogether	.94B
4.3.5 Rendering a Partial Template	.95B
4.3.6 Rendering Inline Template Code	.96B
4.3.7 Rendering Text	.96B
4.3.8 Rendering Other Types of Structured Data	.96B
4.3.9 Rendering Nothing	.97B
4.3.10 Rendering Options	.98B
4.4 Additional Layout Options	.101B
4.5 Redirecting	.101B
4.5.1 The <code>redirect</code> to Method	.102B
4.6 Controller/View Communication	.104B
4.7 Filters	.105B
4.7.1 Filter Inheritance	.106B
4.7.2 Filter Types	.107B
4.7.3 Filter Chain Ordering	.108B
4.7.4 Around Filters	.109B
4.7.5 Filter Chain Skipping	.110B
4.7.6 Filter Conditions	.110B
4.7.7 Filter Chain Halting	.111B
4.8 Verification	.111B
4.8.1 Example Usage	.111B
4.8.2 Options	.112B

4.9 Streaming	.112B
4.9.1 Via render :text => proc	.112B
4.9.2 send_data(data, options = {})	.113B
4.9.3 send_file(path, options = {})	.114B
4.10 Conclusion	.117B
5 Working with Active Record	.119B
5.1 The Basics	.120B
5.2 Macro-Style Methods	.121B
5.2.1 Relationship Declarations	.121B
5.2.2 Convention over Configuration	.122B
5.2.3 Setting Names Manually	.122B
5.2.4 Legacy Naming Schemes	.122B
5.3 Defining Attributes	.123B
5.3.1 Default Attribute Values	.123B
5.3.2 Serialized Attributes	.125B
5.4 CRUD: Creating, Reading, Updating, Deleting	.127B
5.4.1 Creating New Active Record Instances	.127B
5.4.2 Reading Active Record Objects	.128B
5.4.3 Reading and Writing Attributes	.128B
5.4.4 Accessing and Manipulating Attributes Before They Are Typecast	.131B
5.4.5 Reloading	.131B
5.4.6 Cloning	.131B
5.4.7 Dynamic Attribute-Based Finders	.132B
5.4.8 Dynamic Scopes	.133B
5.4.9 Custom SQL Queries	.133B
5.4.10 The Query Cache	.135B
5.4.11 Updating	.136B
5.4.12 Updating by Condition	.137B
5.4.13 Updating a Particular Instance	.138B
5.4.14 Updating Specific Attributes	.139B
5.4.15 Convenience Updaters	.139B
5.4.16 Touching Records	.139B
5.4.17 Controlling Access to Attributes	.140B

5.4.18	Readonly Attributes	.141B
5.4.19	Deleting and Destroying	.141B
5.5	Database Locking	.142B
5.5.1	Optimistic Locking	.143B
5.5.2	Pessimistic Locking	.145B
5.5.3	Considerations	.145B
5.6	Where Clauses	.146B
5.6.1	where(*conditions)	.146B
5.6.2	order(*clauses)	.148B
5.6.3	limit(number) and offset(number)	.149B
5.6.4	select(*clauses)	.149B
5.6.5	from(*tables)	.150B
5.6.6	group(*args)	.150B
5.6.7	having(*clauses)	.150B
5.6.8	includes(*associations)	.151B
5.6.9	joins	.151B
5.6.10	readonly	.152B
5.6.11	exists?	.152B
5.6.12	arel_table	.152B
5.7	Connections to Multiple Databases in Different Models	.153B
5.8	Using the Database Connection Directly	.154B
5.8.1	The DatabaseStatements Module	.154B
5.8.2	Other Connection Methods	.156B
5.9	Other Configuration Options	.158B
5.10	Conclusion	.159B
6	Active Record Migrations	.161B
6.1	Creating Migrations	.161B
6.1.1	Sequencing Migrations	.162B
6.1.2	Irreversible Migrations	.162B
6.1.3	create_table(name, options, & block)	.164B
6.1.4	change_table(table name, & block)	.165B
6.1.5	API Reference	.165B
6.1.6	Defining Columns	.167B
6.1.7	Command-line Column Declarations	.172B

6.2 Data Migration	.173B
6.2.1 Using SQL	.173B
6.2.2 Migration Models	.174B
6.3 schema.rb	.174B
6.4 Database Seeding	.175B
6.5 Database-Related Rake Tasks	.176B
6.6 Conclusion	.179B
7 Active Record Associations	.181B
7.1 The Association Hierarchy	.181B
7.2 One-to-Many Relationships	.183B
7.2.1 Adding Associated Objects to a Collection	.185B
7.2.2 Association Collection Methods	.186B
7.3 The belongs_to Association	.191B
7.3.1 Reloading the Association	.192B
7.3.2 Building and Creating Related Objects via the Association	.192B
7.3.3 belongs_to Options	.193B
7.4 The has_many Association	.200B
7.4.1 has_many Options	.200B
7.5 Many-to-Many Relationships	.209B
7.5.1 has_and_belongs_to_many	.209B
7.5.2 has_many :through	.215B
7.5.3 has_many :through Options	.220B
7.6 One-to-One Relationships	.223B
7.6.1 has_one	.223B
7.7 Working with Unsaved Objects and Associations	.226B
7.7.1 One-to-One Associations	.226B
7.7.2 Collections	.226B
7.7.3 Deletion	.227B
7.8 Association Extensions	.227B
7.9 The AssociationProxy Class	.229B
7.10 Conclusion	.230B

8 Validations	231B
8.1 Finding Errors	231B
8.2 The Simple Declarative Validations	232B
8.2.1 <code>validates_acceptance_of</code>	232B
8.2.2 <code>validates_associated</code>	233B
8.2.3 <code>validates_confirmation_of</code>	233B
8.2.4 <code>validates_each</code>	234B
8.2.5 <code>validates_format_of</code>	235B
8.2.6 <code>validates_inclusion_of</code> and <code>validates_exclusion_of</code>	236B
8.2.7 <code>validates_length_of</code>	236B
8.2.8 <code>validates_numericality_of</code>	237B
8.2.9 <code>validates_presence_of</code>	238B
8.2.10 <code>validates_uniqueness_of</code>	239B
8.2.11 <code>validates_with</code>	241B
8.2.12 <code>RecordInvalid</code>	242B
8.3 Common Validation Options	242B
8.3.1 <code>:allow_blank</code> and <code>:allow_nil</code>	242B
8.3.2 <code>:if</code> and <code>:unless</code>	242B
8.3.3 <code>:message</code>	242B
8.3.4 <code>:on</code>	243B
8.4 Conditional Validation	243B
8.4.1 Usage and Considerations	244B
8.4.2 Validation Contexts	245B
8.5 Short-form Validation	245B
8.6 Custom Validation Techniques	246B
8.6.1 Add Custom Validation Macros to Your Application	247B
8.6.2 Create a Custom Validator Class	248B
8.6.3 Add a <code>validate</code> Method to Your Model	248B
8.7 Skipping Validations	249B
8.8 Working with the Errors Hash	249B
8.8.1 Checking for Errors	250B
8.9 Testing Validations with Shoulda	250B
8.10 Conclusion	250B

- 9 Advanced Active Record251B**
 - 9.1 Scopes251B
 - 9.1.1 Scope Parameters252B
 - 9.1.2 Chaining Scopes252B
 - 9.1.3 Scopes and has many252B
 - 9.1.4 Scopes and Joins253B
 - 9.1.5 Scope Combinations253B
 - 9.1.6 Default Scopes254B
 - 9.1.7 Using Scopes for CRUD255B
 - 9.2 Callbacks256B
 - 9.2.1 Callback Registration256B
 - 9.2.2 One-Liners257B
 - 9.2.3 Protected or Private257B
 - 9.2.4 Matched before/after Callbacks258B
 - 9.2.5 Halting Execution259B
 - 9.2.6 Callback Usages259B
 - 9.2.7 Special Callbacks:
after_initialize and after_find262B
 - 9.2.8 Callback Classes263B
 - 9.3 Calculation Methods265B
 - 9.3.1 average(column_name,
*options)267B
 - 9.3.2 count(column_name,
*options)267B
 - 9.3.3 maximum(column_name,
*options)267B
 - 9.3.4 minimum(column_name,
*options)267B
 - 9.3.5 sum(column_name, *options)267B
 - 9.4 Observers268B
 - 9.4.1 Naming Conventions268B
 - 9.4.2 Registration of Observers269B
 - 9.4.3 Timing269B
 - 9.5 Single-Table Inheritance (STI)269B
 - 9.5.1 Mapping Inheritance to the
Database271B
 - 9.5.2 STI Considerations273B
 - 9.5.3 STI and Associations274B

9.6 Abstract Base Model Classes	276B
9.7 Polymorphic has many Relationships	277B
9.7.1 In the Case of Models with Comments	278B
9.8 Foreign-key Constraints	281B
9.9 Using Value Objects	281B
9.9.1 Immutability	283B
9.9.2 Custom Constructors and Converters	283B
9.9.3 Finding Records by a Value Object	284B
9.10 Modules for Reusing Common Behavior	285B
9.10.1 A Review of Class Scope and Contexts	287B
9.10.2 The <code>included</code> Callback	288B
9.11 Modifying Active Record Classes at Runtime	289B
9.11.1 Considerations	290B
9.11.2 Ruby and Domain-Specific Languages	291B
9.12 Conclusion	292B
10 Action View	293B
10.1 Layouts and Templates	294B
10.1.1 Template Filename Conventions	294B
10.1.2 Layouts	294B
10.1.3 Yielding Content	295B
10.1.4 Conditional Output	296B
10.1.5 Decent Exposure	297B
10.1.6 Standard Instance Variables	298B
10.1.7 Displaying <code>flash</code> Messages	300B
10.1.8 <code>flash.now</code>	301B
10.2 Partial	302B
10.2.1 Simple Use Cases	302B
10.2.2 Reuse of Partial	303B
10.2.3 Shared Partial	304B
10.2.4 Passing Variables to Partial	305B
10.2.5 Rendering Collections	306B
10.2.6 Logging	308B
10.3 Conclusion	308B

11 All About Helpers	309B
11.1 ActiveModelHelper	309B
11.1.1 Reporting Validation Errors	310B
11.1.2 Automatic Form Creation	313B
11.1.3 Customizing the Way Validation Errors Are Highlighted	315B
11.2 AssetTagHelper	316B
11.2.1 Head Helpers	316B
11.2.2 Asset Helpers	319B
11.2.3 Using Asset Hosts	321B
11.2.4 Using Asset Timestamps	323B
11.2.5 For Plugins Only	324B
11.3 AtomFeedHelper	324B
11.4 CacheHelper	326B
11.5 CaptureHelper	326B
11.6 DateHelper	328B
11.6.1 The Date and Time Selection Helpers	328B
11.6.2 The Individual Date and Time Select Helpers	329B
11.6.3 Common Options for Date Selection Helpers	332B
11.6.4 <code>distance_in_time</code> Methods with Complex Descriptive Names	332B
11.7 DebugHelper	333B
11.8 FormHelper	333B
11.8.1 Creating Forms for Models	334B
11.8.2 How Form Helpers Get Their Values ..	342B
11.8.3 Integrating Additional Objects in One Form	343B
11.8.4 Customized Form Builders	347B
11.8.5 Form Inputs	348B
11.9 FormOptionsHelper	350B
11.9.1 Select Helpers	350B
11.9.2 Option Helpers	351B
11.10 FormTagHelper	355B
11.11 JavaScriptHelper	358B
11.12 NumberHelper	359B
11.13 PrototypeHelper	361B

11.14	RawOutputHelper	361B
11.15	RecordIdentificationHelper	362B
11.16	RecordTagHelper	363B
11.17	SanitizeHelper	364B
11.18	TagHelper	366B
11.19	TextHelper	367B
11.20	TranslationHelper and the I18n API	372B
11.20.1	Localized Views	373B
11.20.2	TranslationHelper Methods	374B
11.20.3	I18n Setup	374B
11.20.4	Setting and Passing the Locale	375B
11.20.5	Setting Locale from Client Supplied Information	379B
11.20.6	Internationalizing Your Application	380B
11.20.7	Organization of Locale Files	382B
11.20.8	Looking up Translations	383B
11.20.9	How to Store Your Custom Translations	386B
11.20.10	Overview of Other Built-In Methods that Provide I18n Support	388B
11.20.11	Exception Handling	391B
11.21	UrlHelper	391B
11.22	Writing Your Own View Helpers	398B
11.22.1	Small Optimizations: The Title Helper	398B
11.22.2	Encapsulating View Logic: The photo for Helper	399B
11.22.3	Smart View: The breadcrumbs Helper	400B
11.23	Wrapping and Generalizing Partial	401B
11.23.1	A tiles Helper	401B
11.23.2	Generalizing Partial	404B
11.24	Conclusion	407B
12	Ajax on Rails	409B
12.0.1	Changes in Rails 3	410B
12.0.2	Firebug	410B
12.1	Unobtrusive JavaScript	411B
12.1.1	UJS Usage	411B

12.2 Writing JavaScript in Ruby with RJS	.412B
12.2.1 RJS Templates	.414B
12.2.2 <<(javascript)	.415B
12.2.3 [(id)	.415B
12.2.4 alert(message)	.416B
12.2.5 call(function, *arguments, & block)	.416B
12.2.6 delay(seconds = 1)416B
12.2.7 draggable(id, options = {})	.416B
12.2.8 drop receiving(id, options = {})	.417B
12.2.9 hide(*ids)	.417B
12.2.10 insert_html(position, id, *options_for_render)	.417B
12.2.11 literal(code)	.417B
12.2.12 redirect to(location)	.418B
12.2.13 remove(*ids)	.418B
12.2.14 replace(id, *options for render)	.418B
12.2.15 replace html(id, *options for render)	.418B
12.2.16 select(pattern)	.418B
12.2.17 show(*ids)	.418B
12.2.18 sortable(id, options = {})	.418B
12.2.19 toggle(*ids)	.419B
12.2.20 visual effect(name, id = nil, options = {})	.419B
12.3 Ajax and JSON	.419B
12.3.1 Ajax link to	.419B
12.4 Ajax and HTML	.421B
12.5 Ajax and JavaScript	.423B
12.6 Conclusion	.424B
13 Session Management	.425B
13.1 What to Store in the Session	.426B
13.1.1 The Current User	.426B
13.1.2 Session Use Guidelines	.426B
13.2 Session Options	.427B
13.3 Storage Mechanisms	.427B
13.3.1 Active Record Session Store	.427B
13.3.2 Memcache Session Storage	.428B

13.3.3	The Controversial CookieStore	429B
13.3.4	Cleaning Up Old Sessions	430B
13.4	Cookies	431B
13.4.1	Reading and Writing Cookies	431B
13.5	Conclusion	432B
14	Authentication	433B
14.1	Authlogic	434B
14.1.1	Getting Started	434B
14.1.2	Creating the Models	434B
14.1.3	Setting Up the Controllers	435B
14.1.4	Controller, Limiting Access to Actions	436B
14.1.5	Configuration	437B
14.1.6	Summary	439B
14.2	Devise	439B
14.2.1	Getting Started	439B
14.2.2	Modules	439B
14.2.3	Models	440B
14.2.4	Controllers	441B
14.2.5	Devise, Views	442B
14.2.6	Configuration	442B
14.2.7	Extensions	443B
14.2.8	Summary	443B
14.3	Conclusion	443B
15	XML and Active Resource	445B
15.1	The <code>to_xml</code> Method	445B
15.1.1	Customizing <code>to_xml</code> Output	446B
15.1.2	Associations and <code>to_xml</code>	448B
15.1.3	Advanced <code>to_xml</code> Usage	451B
15.1.4	Dynamic Runtime Attributes	452B
15.1.5	Overriding <code>to_xml</code>	453B
15.2	The XML Builder	454B
15.3	Parsing XML	456B
15.3.1	Turning XML into Hashes	456B
15.3.2	Typecasting	457B

15.4 Active Resource	.457B
15.4.1 List	.458B
15.4.2 Show	.459B
15.4.3 Create	.460B
15.4.4 Update	.462B
15.4.5 Delete	.462B
15.4.6 Headers	.462B
15.4.7 Customizing URLs	.463B
15.4.8 Hash Forms	.464B
15.5 Active Resource Authentication	.465B
15.5.1 HTTP Basic Authentication	.465B
15.5.2 HTTP Digest Authentication	.466B
15.5.3 Certificate Authentication	.466B
15.5.4 Proxy Server Authentication	.466B
15.5.5 Authentication in the Web Service Controller	.467B
15.6 Conclusion	.469B
16 Action Mailer	.471B
16.1 Setup	.471B
16.2 Mailer Models	.472B
16.2.1 Preparing Outbound Email Messages	.472B
16.2.2 HTML Email Messages	.474B
16.2.3 Multipart Messages	.475B
16.2.4 Attachments	.475B
16.2.5 Generating URLs	.476B
16.2.6 Mailer Layouts	.476B
16.2.7 Sending an Email	.477B
16.3 Receiving Emails	.477B
16.3.1 Handling Incoming Attachments	.478B
16.4 Server Configuration	.479B
16.5 Testing Email Content	.479B
16.6 Conclusion	.481B
17 Caching and Performance	.483B
17.1 View Caching	.483B
17.1.1 Caching in Development Mode?	.484B
17.1.2 Page Caching	.484B

17.1.3	Action Caching	.484B
17.1.4	Fragment Caching	.486B
17.1.5	Expiration of Cached Content	.488B
17.1.6	Automatic Cache Expiry with Sweepers	.490B
17.1.7	Cache Logging	.492B
17.1.8	Action Cache Plugin	.492B
17.1.9	Cache Storage	.493B
17.2	General Caching	.495B
17.2.1	Eliminating Extra Database Lookups	.495B
17.2.2	Initializing New Caches	.496B
17.2.3	fetch Options	.496B
17.3	Control Web Caching	.497B
17.3.1	expires_in(seconds, options = {})	.498B
17.3.2	expires_now	.498B
17.4	ETags	.498B
17.4.1	fresh_when(options)	.499B
17.4.2	stale?(options)	.499B
17.5	Conclusion	.500B
18	RSpec	.501B
18.1	Introduction	.501B
18.2	Basic Syntax and API	.504B
18.2.1	describe and context	.504B
18.2.2	let(:name) (expression)	.504B
18.2.3	let!(:name) (expression)	.506B
18.2.4	before and after	.506B
18.2.5	it	.507B
18.2.6	specify	.507B
18.2.7	expect	.508B
18.2.8	pending	.509B
18.2.9	should and should_not	.510B
18.2.10	Implicit Subject	.511B
18.2.11	Explicit Subject	.511B
18.2.12	its	.512B
18.3	Predicate Matchers	.513B

18.4 Custom Expectation Matchers514B
18.4.1 Custom Matcher DSL516B
18.4.2 <i>Fluent</i> Chaining516B
18.5 Shared Behaviors517B
18.6 RSpec's Mocks and Stubs517B
18.7 Running Specs520B
18.8 RSpec Rails Gem521B
18.8.1 Installation521B
18.8.2 Model Specs524B
18.8.3 Mocked and Stubbed Models526B
18.8.4 Controller Specs526B
18.8.5 View Specs529B
18.8.6 Helper Specs531B
18.9 RSpec Tools531B
18.9.1 RSpactor531B
18.9.2 watchr532B
18.9.3 Spork532B
18.9.4 Specjour532B
18.9.5 RCov532B
18.9.6 Heckle532B
18.10 Conclusion533B
19 Extending Rails with Plugins535B
19.1 The Plugin System536B
19.1.1 Plugins as RubyGems536B
19.1.2 The Plugin Script536B
19.2 Writing Your Own Plugins537B
19.2.1 The <code>init.rb</code> Hook538B
19.2.2 The <code>lib</code> Directory539B
19.2.3 Extending Rails Classes540B
19.2.4 The <code>README</code> and <code>MIT-LICENSE</code> File541B
19.2.5 The <code>install.rb</code> and <code>uninstall.rb</code> Files542B
19.2.6 Custom Rake Tasks543B
19.2.7 The Plugin's Rakefile544B
19.2.8 Including Assets With Your Plugin545B

19.2.9 Testing Plugins545B
19.2.10 Rallities546B
19.3 Conclusion547B
20 Background Processing549B
20.1 Delayed Job550B
20.1.1 Getting Started550B
20.1.2 Creating Jobs551B
20.1.3 Running552B
20.1.4 Summary552B
20.2 Resque553B
20.2.1 Getting Started553B
20.2.2 Creating Jobs554B
20.2.3 Hooks554B
20.2.4 Plugins555B
20.2.5 Running556B
20.2.6 Monitoring556B
20.2.7 Summary557B
20.3 Rails Runner557B
20.3.1 Getting Started558B
20.3.2 Usage Notes558B
20.3.3 Considerations559B
20.3.4 Summary559B
20.4 Conclusion559B
A Active Model API Reference561B
A.1 AttributeMethods561B
A.1.1 active_model/ attribute_methods.rb562B
A.2 Callbacks563B
A.2.1 active_model/callbacks.rb563B
A.3 Conversion563B
A.3.1 active_model/conversion.rb563B
A.4 Dirty564B
A.4.1 active_model/dirty.rb565B
A.5 Errors565B
A.5.1 active_model/errors.rb566B

- A.6 `Lint::Tests`567B
- A.7 `MassAssignmentSecurity`567B
 - A.7.1 `active_model/mass_assignment_security.rb`567B
- A.8 `Name`568B
 - A.8.1 `active_model/naming.rb`569B
- A.9 `Naming`569B
 - A.9.1 `active_model/naming.rb`569B
- A.10 `Observer`569B
 - A.10.1 `active_model/observing.rb`570B
- A.11 `Observing`570B
 - A.11.1 `active_model/observing.rb`571B
- A.12 `Serialization`571B
 - A.12.1 `active_model/serialization.rb`571B
- A.13 `Serializers::JSON`572B
 - A.13.1 `active_model/serializers/json.rb`572B
- A.14 `Serializers::Xml`572B
 - A.14.1 `active_model/serializers/xml.rb`573B
- A.15 `Translation`573B
 - A.15.1 `active_model/translation.rb`573B
- A.16 `Validations`574B
 - A.16.1 `active_model/validations.rb`574B
- A.17 `Validator`578B
 - A.17.1 `active_model/validator.rb`578B

B Active Support API Reference579B

- B.1 `Array`579B
 - B.1.1 `active_support/core_ext/array/access`579B
 - B.1.2 `active_support/core_ext/array/conversions`580B
 - B.1.3 `active_support/core_ext/array/extract_options`582B
 - B.1.4 `active_support/core_ext/array/grouping`583B

B.1.5 active_support/core_ext/ array/random_access	584B
B.1.6 active_support/core_ext/ array/uniq_by	584B
B.1.7 active_support/core_ext/ array/wrap	584B
B.1.8 active_support/core_ext/ object/blank	585B
B.1.9 active_support/core_ext/ object/to_param	585B
B.2 ActiveSupport::BacktraceCleaner	585B
B.2.1 active_support/ backtrace_cleaner	585B
B.3 ActiveSupport::Base64	586B
B.3.1 active_support/base64	586B
B.4 ActiveSupport::BasicObject	586B
B.4.1 active_support/basic_object	586B
B.5 ActiveSupport::Benchmarkable	587B
B.5.1 active_support/ benchmarkable	587B
B.6 BigDecimal	588B
B.6.1 active_support/core_ext/ big_decimal/conversions	588B
B.6.2 active_support/json/ encoding	588B
B.7 ActiveSupport::BufferedLogger	588B
B.7.1 active_support/ buffered_logger	589B
B.8 ActiveSupport::Cache::Store	590B
B.9 ActiveSupport::Callbacks	595B
B.9.1 active_support/callbacks	596B
B.10 Class	598B
B.10.1 active_support/core_ext/ class/attribute	598B
B.10.2 active_support/core_ext/ class/attribute_accessors	599B
B.10.3 active_support/core_ext/ class/attribute_accessors	600B
B.10.4 active_support/core_ext/ class/delegating_attributes	600B

B.10.5 active_support/core_ext/ class/inheritable_attributes600B
B.10.6 active_support/core_ext/ class/subclasses601B
B.11 ActiveSupport::Concern602B
B.11.1 active_support/concern602B
B.12 ActiveSupport::Configurable603B
B.12.1 active_support/ configurable603B
B.13 Date603B
B.13.1 active_support/core_ext/ date/acts_like603B
B.13.2 active_support/core_ext/ date/calculations603B
B.13.3 active_support/core_ext/ date/conversions607B
B.13.4 active_support/core_ext/ date/freeze608B
B.13.5 active_support/json/ encoding609B
B.14 DateTime609B
B.14.1 active_support/core_ext/ date_time/acts_like609B
B.14.2 active_support/core_ext/ date_time/calculations609B
B.14.3 active_support/core_ext/ date_time/conversions611B
B.14.4 active_support/core_ext/ date_time/zones612B
B.14.5 active_support/json/ encoding613B
B.15 ActiveSupport::Dependencies613B
B.15.1 active_support/ dependencies/autoload614B
B.16 ActiveSupport::Deprecation617B
B.17 ActiveSupport::Duration617B
B.17.1 active_support/duration617B
B.18 Enumerable619B
B.18.1 active_support/core_ext/ enumerable619B
B.18.2 active_support/json/ encoding620B

B.19 ERB::Util	620B
B.19.1 active_support/core_ext/ string/output_safety	620B
B.20 FalseClass	621B
B.20.1 active_support/core_ext/ object/blank	621B
B.20.2 active_support/json/ encoding	621B
B.21 File	621B
B.21.1 active_support/core_ext/ file/atomic	621B
B.21.2 active_support/core_ext/ file/path	622B
B.22 Float	622B
B.22.1 active_support/core_ext/ float/rounding	622B
B.23 Hash	622B
B.23.1 active_support/core_ext/ hash/conversions	622B
B.23.2 active_support/core_ext/ hash/deep_merge	623B
B.23.3 active_support/core_ext/ hash/diff	624B
B.23.4 active_support/core_ext/ hash/except	624B
B.23.5 active_support/core_ext/ hash/indifferent_access	624B
B.23.6 active_support/core_ext/ hash/keys	625B
B.23.7 active_support/core_ext/ hash/reverse_merge	626B
B.23.8 active_support/core_ext/ hash/slice	626B
B.23.9 active_support/core_ext/ object/to_param	627B
B.23.10 active_support/core_ext/ object/to_query	627B
B.23.11 active_support/json/ encoding	627B
B.23.12 active_support/core_ext/ object/blank	627B

- B.24 HashWithIndifferentAccess627B
 - B.24.1 active_support/
hash_with_indifferent_access627B
- B.25 ActiveSupport::Inflector::
Inflections628B
 - B.25.1 active_support/inflector/
inflections629B
 - B.25.2 active_support/inflector/
transliteration631B
- B.26 Integer632B
 - B.26.1 active_support/core_ext/
integer/inflections633B
 - B.26.2 active_support/core_ext/
integer/multiple633B
- B.27 ActiveSupport::JSON633B
 - B.27.1 active_support/json/
decoding633B
 - B.27.2 active_support/json/
encoding634B
- B.28 Kernel634B
 - B.28.1 active_support/core_ext/
kernel/agnostics634B
 - B.28.2 active_support/core_ext/
kernel/debugger634B
 - B.28.3 active_support/core_ext/
kernel/reporting634B
 - B.28.4 active_support/core_ext/
kernel/requires635B
 - B.28.5 active_support/core_ext/
kernel singleton_class635B
- B.29 Logger635B
 - B.29.1 active_support/core_ext/
logger636B
- B.30 ActiveSupport::
MessageEncryptor636B
 - B.30.1 active_support/
message_encryptor637B
- B.31 ActiveSupport::
MessageVerifier637B
 - B.31.1 active_support/
message_verifier637B

- B.32 Module638B
 - B.32.1 active_support/core_ext/
module/aliasing638B
 - B.32.2 active_support/core_ext/
module/anonymous639B
 - B.32.3 active_support/
core_ext/module/
attr_accessor_with_default640B
 - B.32.4 active_support/core_ext/
module/attr_internal640B
 - B.32.5 active_support/core_ext/
module/attribute_accessors640B
 - B.32.6 active_support/core_ext/
module/delegation641B
 - B.32.7 active_support/core_ext/
module/introspection643B
 - B.32.8 active_support/core_ext/
module/synchronization644B
 - B.32.9 active_support/
dependencies644B
- B.33 ActiveSupport::Multibyte::
Chars645B
 - B.33.1 active_support/
multibyte/chars645B
 - B.33.2 active_support/
multibyte/unicode646B
 - B.33.3 active_support/
multibyte/utils647B
- B.34 NilClass648B
 - B.34.1 active_support/core_ext/
object/blank648B
 - B.34.2 active_support/json/
encoding648B
 - B.34.3 active_support/whiny_nil648B
- B.35 ActiveSupport::Notifications649B
- B.36 Numeric650B
 - B.36.1 active_support/core_ext/
object/blank650B
 - B.36.2 active_support/json/
encoding650B

- B.36.3 active_support/numeric/
bytes650B
- B.36.4 active_support/numeric/
time651B
- B.37 Object653B
 - B.37.1 active_support/core_ext/
object/acts_like653B
 - B.37.2 active_support/core_ext/
object/blank653B
 - B.37.3 active_support/core_ext/
object/duplicable654B
 - B.37.4 active_support/core_ext/
object/instance_variables654B
 - B.37.5 active_support/core_ext/
object/to_param655B
 - B.37.6 active_support/core_ext/
object/with_options656B
 - B.37.7 active_support/
dependencies656B
 - B.37.8 active_support/json/
encoding657B
- B.38 ActiveSupport::OrderedHash657B
 - B.38.1 active_support/
ordered_hash657B
- B.39 ActiveSupport::OrderedOptions657B
 - B.39.1 active_support/
ordered_options657B
- B.40 ActiveSupport::Railtie658B
 - B.40.1 active_support/railtie658B
- B.41 Range658B
 - B.41.1 active_support/core_ext/
range/blockless_step658B
 - B.41.2 active_support/core_ext/
range/conversions659B
 - B.41.3 active_support/core_ext/
range/include_range659B
 - B.41.4 active_support/core_ext/
range/include_range659B

B.42 Regexp	.660B
B.42.1 active_support/core_ext/ enumerable	.660B
B.42.2 active_support/json/ encoding	.660B
B.43 ActiveSupport::Rescuable	.660B
B.43.1 active_support/rescuable	.660B
B.44 ActiveSupport::SecureRandom	.661B
B.44.1 active_support/ secure_random	.661B
B.45 String	.662B
B.45.1 active_support/json/ encoding	.662B
B.45.2 active_support/core_ext/ object/blank	.662B
B.45.3 active_support/core_ext/ string/access	.663B
B.45.4 active_support/core_ext/ string/acts_like	.664B
B.45.5 active_support/core_ext/ string/conversions	.664B
B.45.6 active_support/core_ext/ string/encoding	.665B
B.45.7 active_support/core_ext/ string/exclude	.665B
B.45.8 active_support/core_ext/ string/filters	.665B
B.45.9 active_support/core_ext/ string/inflections	.666B
B.45.10 active_support/core_ext/ string/multibyte	.669B
B.45.11 active_support/core_ext/ string/output_safety	.670B
B.45.12 active_support/core_ext/ string/starts_ends_with	.670B
B.45.13 active_support/core_ext/ string/xchar	.671B
B.46 ActiveSupport::StringInquirer	.671B

- B.47 Symbol671B
 - B.47.1 active_support/json/
encoding671B
- B.48 ActiveSupport::Testing::
Assertions671B
 - B.48.1 active_support/testing/
assertions671B
- B.49 Time673B
 - B.49.1 active_support/json/
encoding673B
 - B.49.2 active_support/core_ext/
time/acts_like673B
 - B.49.3 active_support/core_ext/
time/calculations673B
 - B.49.4 active_support/core_ext/
time/conversions677B
 - B.49.5 active_support/core_ext/
time/marshal679B
 - B.49.6 active_support/core_ext/
time/zones679B
- B.50 ActiveSupport::TimeWithZone680B
- B.51 ActiveSupport::TimeZone681B
 - B.51.1 active_support/values/
time_zone682B
- B.52 ActiveSupport::TrueClass684B
 - B.52.1 active_support/core_ext/
object/blank684B
 - B.52.2 active_support/json/
encoding684B
- B.53 ActiveSupport::XmlMini684B
 - B.53.1 active_support/xml_mini685B

- Index687B**

- Method Index697B**

RUBY ON RAILS™ 3 TUTORIAL

This page intentionally left blank

Praise for *Ruby on Rails™ 3 Tutorial*

RailsTutorial.org: Michael Hartl's awesome new Rails Tutorial

The *Ruby on Rails™ 3 Tutorial: Learn Rails by Example* by Michael Hartl has become a must read for developers learning how to build Rails apps.

—Peter Cooper, editor of Ruby Inside

Very detailed and hands-on Rails Tutorial!

Great job! I'm learning Rails, and found your tutorial to be one of the most detailed and hands-on guides. Besides many details of Rails, it also taught me about Git, Heroku, RSpec, Webrat, and most important (at least to me), it emphasized the Test-Driven Development (TDD) methodology. I learned a lot from your tutorial.

Keep up the good job! Thanks so much for sharing it.

—Albert Liu, senior manager, Achievo Corporation.

Ruby on Rails Tutorial is the best!

Just wanted to say that your Ruby on Rails tutorial is the best!

I've been trying for a while to wrap my head around Rails. Going through your tutorial, I'm finally feeling comfortable in the Rails environment. Your pedagogical style of

gradually introducing more complex topics while at the same time giving the reader the instant gratification and a sense of accomplishment with working examples really works for me. I also like the tips and suggestions that give me a sense of learning from a real Rails insider. Your e-mail response to a problem I ran into is an example of your generous sharing of your experience.

—Ron Bingham, CEO, SoundBuytz


I love the writing style of the Rails Tutorial

I love the writing style of the Rails Tutorial, and there is so much content that is different from other Rails books out there, making it that much more valuable...Thanks for your work!

—Allen Ding

Addison-Wesley Professional Ruby Series

Obie Fernandez, Series Editor



Visit informit.com/ruby for a complete list of available products.

The **Addison-Wesley Professional Ruby Series** provides readers with practical, people-oriented, and in-depth information about applying the Ruby platform to create dynamic technology solutions. The series is based on the premise that the need for expert reference books, written by experienced practitioners, will never be satisfied solely by blogs and the Internet.

PEARSON

Addison-Wesley

Cisco Press

EXAM/CRAM

IBM
Press

que

PRENTICE
HALL

SAMS

Safari
Books Online

This page intentionally left blank

RUBY ON RAILS™ 3 TUTORIAL

Learn Rails™ by Example

Michael Hartl

◆ Addison-Wesley

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States please contact:

International Sales
international@pearson.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data

Hartl, Michael.

Ruby on rails 3 tutorial : learn Rails by example / Michael Hartl.

p. cm.

Includes index.

ISBN-10: 0-321-74312-1 (pbk. : alk. paper)

ISBN-13: 978-0-321-74312-1 (pbk. : alk. paper)

1. Ruby on rails (Electronic resource) 2. Web site development. 3. Ruby (Computer program language) I. Title.

TK5105.8885.R83H37 2011

005.1'17—dc22

2010039450

Copyright © 2011 Michael Hartl

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, write to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax: (617) 671-3447

The source code in *Ruby on Rails™ 3 Tutorial* is released under the MIT License.

ISBN 13: 978-0-321-74312-1

ISBN 10: 0-321-74312-1

Text printed in the United States on recycled paper at Edwards Brothers in Ann Arbor, Michigan
Second printing, April 2011

Editor-in-Chief

Mark Taub

Executive Acquisitions Editor

Debra Williams Cauley

Managing Editor

John Fuller

Project Editor

Elizabeth Ryan

Copy Editor

Erica Orloff

Indexer

Claire Splan

Proofreader

Claire Splan

Publishing Coordinator

Kim Boedigheimer

Cover Designer

Gary Adair

Compositor

Glyph International

Foreword

My former company (CD Baby) was one of the first to loudly switch to Ruby on Rails, and then even more loudly switch back to PHP (Google me to read about the drama). This book by Michael Hartl came so highly recommended that I had to try it, and *Ruby on Rails™ 3 Tutorial* is what I used to switch back to Rails again.

Though I've worked my way through many Rails books, this is the one that finally made me get it. Everything is done very much “the Rails way”—a way that felt very unnatural to me before, but now after doing this book finally feels natural. This is also the only Rails book that does test-driven development the entire time, an approach highly recommended by the experts but which has never been so clearly demonstrated before. Finally, by including Git, GitHub, and Heroku in the demo examples, the author really gives you a feel for what it's like to do a real-world project. The tutorial's code examples are not in isolation.

The linear narrative is such a great format. Personally, I powered through *Rails Tutorial* in three long days, doing all the examples and challenges at the end of each chapter. Do it from start to finish, without jumping around, and you'll get the ultimate benefit.

Enjoy!

—Derek Sivers (sivers.org)
Founder, CD Baby and Thoughts, Ltd.

This page intentionally left blank

Foreword

“If I want to learn web development with Ruby on Rails, how should I start?” For years Michael Hartl has provided the answer as author of the *RailsSpace* tutorial in our series and now the new *Ruby on Rails™ 3 Tutorial* that you hold in your hands (or PDF reader, I guess.)

I’m so proud of having Michael on the series roster. He is living, breathing proof that we Rails folks are some of the luckiest in the wide world of technology. Before getting into Ruby, Michael taught theoretical and computational physics at Caltech for six years, where he received the Lifetime Achievement Award for Excellence in Teaching in 2000. He is a Harvard graduate, has a Ph.D. in Physics from Caltech, and is an alumnus of Paul Graham’s esteemed Y Combinator program for entrepreneurs. And what does Michael apply his impressive experience and teaching prowess to? Teaching new software developers all around the world how to use Ruby on Rails effectively! Lucky we are indeed!

The availability of this tutorial actually comes at a critical time for Rails adoption. We’re five years into the history of Rails and today’s version of the platform has unprecedented power and flexibility. Experienced Rails folks can leverage that power effectively, but we’re hearing growing cries of frustration from newcomers. The amount of information out there about Rails is fantastic if you know what you’re doing already. However, if you’re new, the scope and mass of information about Rails can be mind-boggling.

Luckily, Michael takes the same approach as he did in his first book in the series, building a sample application from scratch, and writes in a style that’s meant to be read from start to finish. Along the way, he explains all the little details that are likely to trip up beginners. Impressively, he goes beyond just a straightforward explanation of what Rails does and ventures into prescriptive advice about good software development

practices, such as test-driven development. Neither does Michael constrain himself to a box delineated by the extents of the Rails framework—he goes ahead and teaches the reader to use tools essential to existence in the Rails community, such as Git and GitHub. In a friendly style, he even provides copious contextual footnotes of benefit to new programmers, such as the pronunciation of SQL and pointers to the origins of *lorem ipsum*. Tying all the content together in a way that remains concise and usable is truly a tour de force of dedication!

I tell you with all my heart that this book is one of the most significant titles in my Professional Ruby Series, because it facilitates the continued growth of the Rails ecosystem. By helping newcomers become productive members of the community quickly, he ensures that Ruby on Rails continues its powerful and disruptive charge into the mainstream. The Rails Tutorial is potent fuel for the fire that is powering growth and riches for so many of us, and for that we are forever grateful.

—Obie Fernandez, Series Editor

Acknowledgments

Ruby on Rails™ Tutorial owes a lot to my previous Rails book, *RailsSpace*, and hence to my coauthor on that book, Aurelius Prochazka. I'd like to thank Aure both for the work he did on that book and for his support of this one. I'd also like to thank Debra Williams Cauley, my editor on both *RailsSpace* and *Rails Tutorial*; as long as she keeps taking me to baseball games, I'll keep writing books for her.

I'd like to acknowledge a long list of Rubyists who have taught and inspired me over the years: David Heinemeier Hansson, Yehuda Katz, Carl Lerche, Jeremy Kemper, Xavier Noria, Ryan Bates, Geoffrey Grosenbach, Peter Cooper, Matt Aimonetti, Gregg Pollack, Wayne E. Seguin, Amy Hoy, Dave Chelimsky, Pat Maddox, Tom Preston-Werner, Chris Wanstrath, Chad Fowler, Josh Susser, Obie Fernandez, Ian McFarland, Steven Bristol, Giles Bowkett, Evan Dorn, Long Nguyen, James Lindenbaum, Adam Wiggins, Tikhon Bernstam, Ron Evans, Wyatt Greene, Miles Forrest, the good people at Pivotal Labs, the Heroku gang, the thoughtbot guys, and the GitHub crew. Finally, many, many readers—far too many to list—have contributed a huge number of bug reports and suggestions during the writing of this book, and I gratefully acknowledge their help in making it as good as it can be.

This page intentionally left blank

About the Author

Michael Hartl is a programmer, educator, and entrepreneur. Michael is coauthor of *RailsSpace*, a best-selling Rails tutorial book published in 2007, and was cofounder and lead developer of Insoshi, a popular social networking platform in Ruby on Rails. Previously, he taught theoretical and computational physics at the California Institute of Technology (Caltech) for six years, where he received the Lifetime Achievement Award for Excellence in Teaching in 2000. Michael is a graduate of Harvard College, has a Ph.D. in Physics from Caltech, and is an alumnus of the Y Combinator program.

This page intentionally left blank

CHAPTER 1

From Zero to Deploy

Welcome to *Ruby on Rails™ 3 Tutorial: Learn Rails by Example*. The goal of this book is to be the best answer to the question, “If I want to learn web development with Ruby on Rails, where should I start?” By the time you finish *Ruby on Rails Tutorial*, you will have all the knowledge you need to develop and deploy your own custom web applications. You will also be ready to benefit from the many more advanced books, blogs, and screencasts that are part of the thriving Rails educational ecosystem. Finally, since *Ruby on Rails Tutorial* uses Rails 3.0, the knowledge you gain here will be fully up to date with the latest and greatest version of Rails.¹

Ruby on Rails Tutorial follows essentially the same approach as my previous Rails book,² teaching web development with Rails by building a substantial sample application from scratch. As Derek Sivers notes in the foreword, this book is structured as a linear narrative, designed to be read from start to finish. If you are used to skipping around in technical books, taking this linear approach might require some adjustment, but I suggest giving it a try. You can think of *Ruby on Rails Tutorial* as a video game where you are the main character, and where you level up as a Rails developer in each chapter. (The exercises are the minibosses.)

In this first chapter, we’ll get started with Ruby on Rails by installing all the necessary software and setting up our development environment (Section 1.2). We’ll then create our first Rails application, called (appropriately enough) `first_app`. *Rails Tutorial* emphasizes good software development practices, so immediately after creating our fresh

1. The most up-to-date version of *Ruby on Rails Tutorial* can be found on the book’s website at <http://railstutorial.org/>. If you are reading this book offline, be sure to check the online version of the Rails Tutorial book at <http://railstutorial.org/book> for the latest updates. In addition, PDF books purchased through railstutorial.org will continue to be updated as long as Rails 3.0 and RSpec 2.0 are still under active development.

2. *RailsSpace*, by Michael Hartl and Aurelius Prochazka (Addison-Wesley, 2007).

new Rails project we'll put it under version control with Git (Section 1.3). And, believe it or not, in this chapter we'll even put our first app on the wider web by *deploying* it to production (Section 1.4).

In Chapter 2, we'll make a second project, whose purpose will be to demonstrate the basic workings of a Rails application. To get up and running quickly, we'll build this *demo app* (called `demo`app`) using scaffolding (Box 1.1) to generate code; since this code is both ugly and complex, Chapter 2 will focus on interacting with the demo app through its *URLs*³ using a web browser.

In Chapter 3, we'll create a *sample application* (called `sample`app`), this time writing all the code from scratch. We'll develop the sample app using *test-driven development* (TDD), getting started in Chapter 3 by creating static pages and then adding a little dynamic content. We'll take a quick detour in Chapter 4 to learn a little about the Ruby language underlying Rails. Then, in Chapter 5 through Chapter 10, we'll complete the foundation for the sample application by making a site layout, a user data model, and a full registration and authentication system. Finally, in Chapter 11 and Chapter 12 we'll add microblogging and social features to make a working example site.

The final sample application will bear more than a passing resemblance to a certain popular social microblogging site—a site which, coincidentally, is also written in Rails. Though of necessity our efforts will focus on this specific sample application, the emphasis throughout *Rails Tutorial* will be on general principles, so that you will have a solid foundation no matter what kinds of web applications you want to build.

Box 1.1 Scaffolding: Quicker, easier, more seductive

From the beginning, Rails has benefited from a palpable sense of excitement, starting with the famous 15-minute weblog video by Rails creator David Heinemeier Hansson, now updated as the 15-minute weblog using Rails 2 by Ryan Bates. These videos are a great way to get a taste of Rails' power, and I recommend watching them. But be warned: they accomplish their amazing fifteen-minute feat using a feature called *scaffolding*, which relies heavily on *generated code*, magically created by the Rails **generate** command.

When writing a Ruby on Rails tutorial, it is tempting to rely on the scaffolding approach—it's quicker, easier, more seductive. But the complexity and sheer amount of code in the scaffolding can be utterly overwhelming to a beginning Rails developer;

3. *URL* stands for Uniform Resource Locator. In practice, it is usually equivalent to “the thing you see in the address bar of your browser”. By the way, the current preferred term is *URI*, for Uniform Resource Identifier, but popular usage still tilts toward *URL*.

you may be able to use it, but you probably won't understand it. Following the scaffolding approach risks turning you into a virtuoso script generator with little (and brittle) actual knowledge of Rails.

In *Ruby on Rails Tutorial*, we'll take the (nearly) polar opposite approach: although Chapter 2 will develop a small demo app using scaffolding, the core of *Rails Tutorial* is the sample app, which we'll start writing in Chapter 3. At each stage of developing the sample application, we will generate *small, bite-sized* pieces of code—simple enough to understand, yet novel enough to be challenging. The cumulative effect will be a deeper, more flexible knowledge of Rails, giving you a good background for writing nearly any type of web application.

1.1 Introduction

Since its debut in 2004, Ruby on Rails has rapidly become one of the most powerful and popular frameworks for building dynamic web applications. Rails users run the gamut from scrappy startups to huge companies: Posterous, UserVoice, 37signals, Shopify, Scribd, Twitter, Hulu, the Yellow Pages—the list of sites using Rails goes on and on. There are also many web development shops that specialize in Rails, such as ENTP, thoughtbot, Pivotal Labs, and Hashrocket, plus innumerable independent consultants, trainers, and contractors.

What makes Rails so great? First of all, Ruby on Rails is 100 percent open-source, available under the permissive MIT License, and as a result it also costs nothing to download and use. Rails also owes much of its success to its elegant and compact design; by exploiting the malleability of the underlying Ruby language, Rails effectively creates a domain-specific language for writing web applications. As a result, many common web programming tasks—such as generating HTML, making data models, and routing URLs—are easy with Rails, and the resulting application code is concise and readable.

Rails also adapts rapidly to new developments in web technology and framework design. For example, Rails was one of the first frameworks to fully digest and implement the REST architectural style for structuring web applications (which we'll be learning about throughout this tutorial). And when other frameworks develop successful new techniques, Rails creator David Heinemeier Hansson and the Rails core team don't hesitate to incorporate their ideas. Perhaps the most dramatic example is the merger of Rails and Merb, a rival Ruby web framework, so that Rails now benefits from Merb's modular design, stable API, and improved performance. (Anyone who has attended a talk by Merb developer and Rails core team member Yehuda Katz can't help but notice what an *extremely* good idea it was to bring the Merb team on board.)

Finally, Rails benefits from an unusually enthusiastic and diverse community. The results include hundreds of open-source contributors, well-attended conferences, a huge number of plugins and gems (self-contained solutions to specific problems such as pagination and image upload), a rich variety of informative blogs, and a cornucopia of discussion forums and IRC channels. The large number of Rails programmers also makes it easier to handle the inevitable application errors: the “Google the error message” algorithm nearly always produces a relevant blog post or discussion-forum thread.

1.1.1 Comments for Various Readers

Rails Tutorial contains integrated tutorials not only for Rails, but also for the underlying Ruby language, as well as for HTML, CSS, some JavaScript, and even a little SQL. This means that, no matter where you currently are in your knowledge of web development, by the time you finish this tutorial you will be ready for more advanced Rails resources, as well as for the more systematic treatments of the other subjects mentioned.

Rails derives much of its power from “magic”—that is, framework features (such as automatically inferring object attributes from database columns) that accomplish miracles but whose mechanisms can be rather mysterious. *Ruby on Rails Tutorial* is *not* designed to explain this magic—mainly because most Rails application developers never need to know what’s behind the curtain. (After all, Ruby itself is mostly written in the C programming language, but you don’t have to dig into the C source to use Ruby.) If you’re a confirmed pull-back-the-curtain kind of person, I recommend *The Rails 3 Way* by Obie Fernandez as a companion volume to *Ruby on Rails Tutorial*.

Although this book has no formal prerequisites, you should of course have at least *some* computer experience. If you’ve never even used a text editor before, it will be tough going, but with enough determination you can probably soldier through. If, on the other hand, your `.emacs` file is so complex it could make a grown man cry, there is still plenty of material to keep you challenged. *Rails Tutorial* is designed to teach Rails development no matter what your background is, but your path and reading experience will depend on your particular circumstances.

All readers: One common question when learning Rails is whether to learn Ruby first. The answer depends on your personal learning style. If you prefer to learn everything systematically from the ground up, then learning Ruby first might work well for you, and there are several book recommendations in this section to get you started. On the other hand, many beginning Rails developers are excited about making *web* applications,

and would rather not slog through a 500-page book on pure Ruby before ever writing a single web page. Moreover, the subset of Ruby needed by Rails developers is different from what you'll find in a pure-Ruby introduction, whereas *Rails Tutorial* focuses on exactly that subset. If your primary interest is making web applications, I recommend starting with *Rails Tutorial* and then reading a book on pure Ruby next. It's not an all-or-nothing proposition, though: if you start reading *Rails Tutorial* and feel your (lack of) Ruby knowledge holding you back, feel free to switch to a Ruby book and come back when you feel ready. You might also consider getting a taste of Ruby by following a short online tutorial, such as can be found at <http://www.ruby-lang.org/> or <http://rubylearning.com/>.

Another common question is whether to use tests from the start. As noted in the introduction, *Rails Tutorial* uses test-driven development (also called test-first development), which in my view is the best way to develop Rails applications, but it does introduce a substantial amount of overhead and complexity. If you find yourself getting bogged down by the tests, feel free to skip them on first reading.⁴ Indeed, some readers may find the inclusion of so many moving parts—such as tests, version control, and deployment—a bit overwhelming at first, and if you find yourself expending excessive energy on any of these steps, *don't hesitate to skip them*. Although I have included only material I consider essential to developing professional-grade Rails applications, only the core application code is strictly necessary the first time through.

Inexperienced programmers (non-designers): *Rails Tutorial* doesn't assume any background other than general computer knowledge, so if you have limited programming experience this book is a good place to start. Please bear in mind that it is only the first step on a long journey; web development has many moving parts, including HTML/CSS, JavaScript, databases (including SQL), version control, and deployment. This book contains short introductions to these subjects, but there is much more to learn.

Inexperienced programmers (designers): Your design skills give you a big leg up, since you probably already know HTML and CSS. After finishing this book you will be in an excellent position to work with existing Rails projects and possibly start some of your own. You may find the programming material challenging, but the Ruby language is unusually friendly to beginners, especially those with an artistic bent.

4. In practice, this will involve omitting all files with `spec` in their name, as we will start to see in Section 3.2.2.

After finishing *Ruby on Rails Tutorial*, I recommend that newer programmers read *Beginning Ruby* by Peter Cooper, which shares the same basic instructional philosophy as *Rails Tutorial*. I also recommend *The Ruby Way* by Hal Fulton. Finally, to gain a deeper understanding of Rails, I recommend *The Rails 3 Way* by Obie Fernandez.

Web applications, even relatively simple ones, are by their nature fairly complex. If you are completely new to web programming and find *Rails Tutorial* overwhelming, it could be that you're not quite ready to make web applications yet. In that case, I'd suggest learning the basics of HTML and CSS and then giving *Rails Tutorial* another go. (Unfortunately, I don't have a personal recommendation here, but *Head First HTML* looks promising, and one reader recommends *CSS: The Missing Manual* by David Sawyer McFarland.) You might also consider reading the first few chapters of *Beginning Ruby*, which starts with sample applications much smaller than a full-blown web app.

Experienced programmers new to web development: Your previous experience means you probably already understand ideas like classes, methods, data structures, etc., which is a big advantage. Be warned that if your background is in C/C++ or Java, you may find Ruby a bit of an odd duck, and it might take time to get used to it; just stick with it and eventually you'll be fine. (Ruby even lets you put semicolons at the ends of lines if you miss them too much.) *Rails Tutorial* covers all the web-specific ideas you'll need, so don't worry if you don't currently know a `PUT` from a `POST`.

Experienced web developers new to Rails: You have a great head start, especially if you have used a dynamic language such as PHP or (even better) Python. The basics of what we cover will likely be familiar, but test-driven development may be new to you, as may be the structured REST style favored by Rails. Ruby has its own idiosyncrasies, so those will likely be new, too.

Experienced Ruby programmers: The set of Ruby programmers who don't know Rails is a small one nowadays, but if you are a member of this elite group you can fly through this book and then move on to *The Rails 3 Way* by Obie Fernandez.

Inexperienced Rails programmers: You've perhaps read some other tutorials and made a few small Rails apps yourself. Based on reader feedback, I'm confident that you can still get a lot out of this book. Among other things, the techniques here may be more up to date than the ones you picked up when you originally learned Rails.

Experienced Rails programmers: This book is unnecessary for you, but many experienced Rails developers have expressed surprise at how much they learned from this book, and you might enjoy seeing Rails from a different perspective.

After finishing *Ruby on Rails Tutorial*, I recommend that experienced (non-Ruby) programmers read *The Well-Founded Rubyist* by David A. Black, which is an excellent in-depth discussion of Ruby from the ground up, or *The Ruby Way* by Hal Fulton, which is also fairly advanced but takes a more topical approach. Then move on to *The Rails 3 Way* to deepen your Rails expertise.

At the end of this process, no matter where you started, you will be ready for the more intermediate-to-advanced Rails resources. Here are some I particularly recommend:

- Railscasts: Excellent free Rails screencasts.
- PeepCode, Pragmatic.tv, EnvyCasts: Excellent commercial screencasters.
- Rails Guides: Good topical and up-to-date Rails references. *Rails Tutorial* refers frequently to the *Rails Guides* for more in-depth treatment of specific topics.
- Rails blogs: Too many to list, but there are tons of good ones.

1.1.2 “Scaling” Rails

Before moving on with the rest of the introduction, I’d like to take a moment to address the one issue that dogged the Rails framework the most in its early days: the supposed inability of Rails to “scale”—i.e., to handle large amounts of traffic. Part of this issue relied on a misconception; you scale a *site*, not a framework, and Rails, as awesome as it is, is only a framework. So the real question should have been, “Can a site built with Rails scale?” In any case, the question has now been definitively answered in the affirmative: some of the most heavily trafficked sites in the world use Rails. Actually *doing* the scaling is beyond the scope of just Rails, but rest assured that if *your* application ever needs to handle the load of Hulu or the Yellow Pages, Rails won’t stop you from taking over the world.

1.1.3 Conventions in This Book

The conventions in this book are mostly self-explanatory; in this section, I’ll mention some that may not be. First, both the HTML and PDF editions of this book are full of

links, both to internal sections (such as Section 1.2) and to external sites (such as the main Ruby on Rails download page).⁵

Second, your humble author is a Linux/OS X kind of guy, and hasn't used Windows as his primary OS for more than a decade; as a result, *Rails Tutorial* has an unmistakable Unix flavor.⁶ For example, in this book all command line examples use a Unix-style command line prompt (a dollar sign):

```
$ echo "hello, world"
hello, world
```

Rails comes with lots of commands that can be run at the command line. For example, in Section 1.2.5 we'll run a local development web server as follows:

```
$ rails server
```

Rails Tutorial will also use Unix-style forward slashes as directory separators; my Rails Tutorial sample app, for instance, lives in

```
/Users/mhartl/rails_projects/first_app
```

The root directory for any given app is known as the *Rails root*, and henceforth all directories will be relative to this directory. For example, the **config** directory of my sample application is in

```
/Users/mhartl/rails_projects/first_app/config
```

This means that when referring to the file

```
/Users/mhartl/rails_projects/first_app/config/routes.rb
```

I'll omit the Rails root and write **config/routes.rb** for brevity.

5. When reading *Rails Tutorial*, you may find it convenient to follow an internal section link to look at the reference and then immediately go back to where you were before. This is easy when reading the book as a web page, since you can just use the Back button of your browser, but both Adobe Reader and OS X's Preview allow you to do this with the PDF as well. In Reader, you can right-click on the document and select "Previous View" to go back. In Preview, use the Go menu: **Go > Back**.

6. Indeed, the entire Rails community has this flavor. In a full room at RailsConf you'll see a handful of PCs in a sea of MacBooks—with probably half the PCs running Linux. You can certainly develop Rails apps on Microsoft Windows, but you'll definitely be in the minority.

Finally, *Rails Tutorial* often shows output from various programs (shell commands, version control status, Ruby programs, etc.). Because of the innumerable small differences between different computer systems, the output you see may not always agree exactly with what is shown in the text, but this is not cause for concern. In addition, some commands may produce errors depending on your system; rather than attempt the Sisyphean task of documenting all such errors in this tutorial, I will delegate to the “Google the error message” algorithm, which among other things is good practice for real-life software development.

1.2 Up and Running

It’s time now to get going with a Ruby on Rails development environment and our first application. There is quite a bit of overhead here, especially if you don’t have extensive programming experience, so don’t get discouraged if it takes a while to get started. It’s not just you; every developer goes through it (often more than once), but rest assured that the effort will be richly rewarded.

1.2.1 Development Environments

Considering various idiosyncratic customizations, there are probably as many development environments as there are Rails programmers, but there are at least two broad themes: text editor/command line environments, and integrated development environments (IDEs). Let’s consider the latter first.

IDEs

There is no shortage of Rails IDEs; indeed, the main Ruby on Rails site names four: RadRails, RubyMine, 3rd Rail, and NetBeans. All are cross-platform, and I’ve heard good things about several of them. I encourage you to try them and see if they work for you, but I have a confession to make: I have never found an IDE that met all my Rails development needs—and for some projects I haven’t even been able to get them to work at all.

Text Editors and Command Lines

What are we to use to develop Rails apps, if not some awesome all-in-one IDE? I’d guess the majority of Rails developers opt for the same solution I’ve chosen: use a *text editor* to edit text, and a *command line* to issue commands (Figure 1.1). Which combination you use depends on your tastes and your platform:

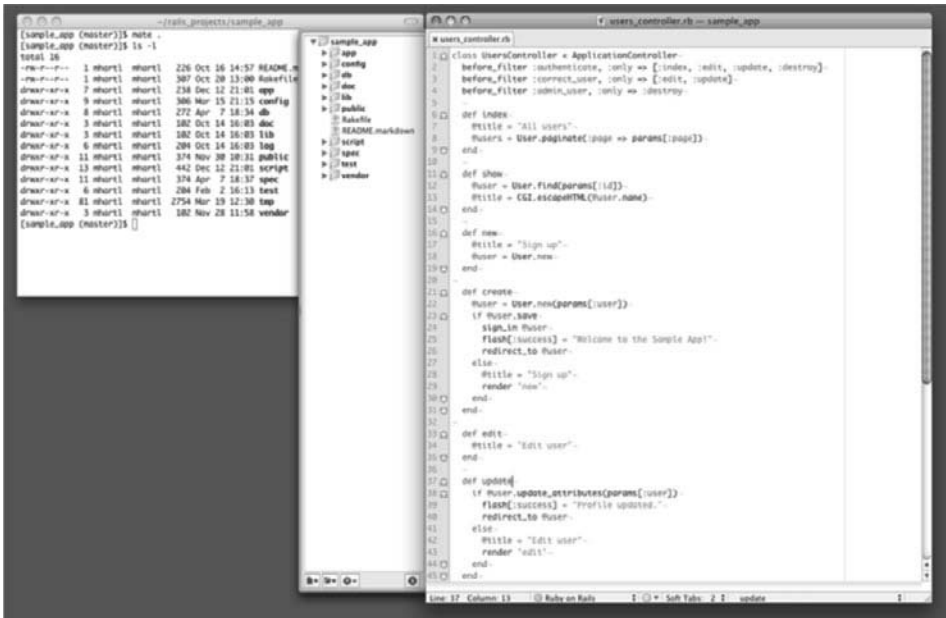


Figure 1.1 A text editor/command line development environment (TextMate/iTerm).

- **Macintosh OS X:** Like many Rails developers, I prefer TextMate. Other options include Emacs and MacVim (launched with the command `macvim`), the excellent Macintosh version of Vim.⁷ I use iTerm for my command line terminal; others prefer the native Terminal app.
- **Linux:** Your editor options are basically the same as OS X, minus TextMate. I'd recommend graphical Vim (gVim), gedit (with the GMate plugins), or Kate. As far as command lines go, you're totally set: every Linux distribution comes with at least one command line terminal application (and often several).
- **Windows:** Unfortunately, I can't make any personal recommendations here, but you can do what I did: drop "rails windows" into Google to see what the latest thinking is on setting up a Rails development environment on Windows. Two combinations look especially promising: Vim for Windows with Console (recommended by Akita On Rails) or the E Text Editor with Console and Cygwin (recommended by Ben

7. The vi editor is one of the most ancient yet powerful weapons in the Unix arsenal, and Vim is "vi improved".

Kittrell). Rails Tutorial readers have suggested looking at Komodo Edit (cross-platform) and the Sublime Text editor (Windows only) as well. No matter which editor you choose, I recommend trying Cygwin, which provides the equivalent of a Unix terminal under Windows; see, for example, this video on Ruby on Rails + Cygwin + Windows Vista. (In addition to installing the packages in the video, I recommend installing `git`, `curl`, and `vim`. Don't install Rails as in the video, though; use the instructions below instead.) With Cygwin, most of the command-line examples in the book should work with minimum modification.

If you go with some flavor of Vim, be sure to tap into the thriving community of Vim-using Rails hackers. See especially the rails.vim enhancements and the NERD tree project drawer.

Browsers

Although there are many web browsers to choose from, the vast majority of Rails programmers use Firefox, Safari, or Chrome when developing. The screenshots in Rails Tutorial will generally be of a Firefox browser. If you use Firefox, I suggest using the Firebug add-on, which lets you perform all sorts of magic, such as dynamically inspecting (and even editing) the HTML structure and CSS rules on any page. For those not using Firefox, Firebug Lite works with most other browsers, and both Safari and Chrome have a built-in “Inspect element” feature available by right-clicking on any part of the page. Regardless of which browser you use, experience shows that the time spent learning such a web inspector tool will be richly rewarded.

A Note About Tools

In the process of getting your development environment up and running, you may find that you spend a *lot* of time getting everything just right. The learning process for editors and IDEs is particularly long; you can spend weeks on TextMate or Vim tutorials alone. If you're new to this game, I want to assure you that *spending time learning tools is normal*. Everyone goes through it. Sometimes it is frustrating, and it's easy to get impatient when you have an awesome web app in your head and you *just want to learn Rails already*, but have to spend a week learning some weird ancient Unix editor just to get started. But a craftsman has to know his tools; in the end the reward is worth the effort.

1.2.2 Ruby, RubyGems, Rails, and Git

Now it's time to install Ruby and Rails. The canonical up-to-date source for this step is the Ruby on Rails download page. I'll assume you can go there now; parts of this book

can be read profitably offline, but not this part. I'll just inject some of my own comments on the steps.

Install Git

Much of the Rails ecosystem depends in one way or another on a version control system called Git (covered in more detail in Section 1.3). Because its use is ubiquitous, you should install Git even at this early stage; I suggest following the installation instructions for your platform at the Installing Git section of *Pro Git*.

Install Ruby

The next step is to install Ruby. It's possible that your system already has it; try running

```
$ ruby -v
ruby 1.9.2
```

to see the version number. Rails 3 requires Ruby 1.8.7 or later and works best with Ruby 1.9.2. This tutorial assumes that you are using the latest development version of Ruby 1.9.2, known as Ruby 1.9.2-head, but Ruby 1.8.7 should work as well.

The Ruby 1.9 branch is under heavy development, so unfortunately installing the latest Ruby can be quite a challenge. You will likely have to rely on the web for the most up-to-date instructions. What follows is a series of steps that I've gotten to work on my system (Macintosh OS X), but you may have to search around for steps that work on your system.

As part of installing Ruby, if you are using OS X or Linux I strongly recommend installing Ruby using Ruby Version Manager (RVM), which allows you to install and manage multiple versions of Ruby on the same machine. (The Pik project accomplishes a similar feat on Windows.) This is particularly important if you want to run Rails 3 and Rails 2.3 on the same machine. If you want to go this route, I suggest using RVM to install two Ruby/Rails combinations: Ruby 1.8.7/Rails 2.3.10 and Ruby 1.9.2/Rails 3.0.1. If you run into any problems with RVM, you can often find its creator, Wayne E. Seguin, on the RVM IRC channel (#rvm on freenode.net).⁸

8. If you haven't used IRC before, I suggest you start by searching the web for "irc client <your platform>". Two good native clients for OS X are Colloquy and LimeChat. And of course there's always the web interface at <http://webchat.freenode.net/?channels=rvm>.

After installing RVM, you can install Ruby as follows:⁹

```
$ rvm update --head
$ rvm reload
$ rvm install 1.8.7
$ rvm install 1.9.2
<wait a while>
```

Here the first two commands update and reload RVM itself, which is a good practice since RVM gets updated frequently. The final two commands do the actual Ruby installations; depending on your system, they might take a while to download and compile, so don't worry if it seems to be taking forever. (Also beware that lots of little things can go wrong. For example, on my system the latest version of Ruby 1.8.7 won't compile; instead, after much searching and hand-wringing, I discovered that I needed "patchlevel" number 174:

```
$ rvm install 1.8.7-p174
```

When things like this happen to you, it's always frustrating, but at least you know that it happens to everyone...)

Ruby programs are typically distributed via *gems*, which are self-contained packages of Ruby code. Since gems with different version numbers sometimes conflict, it is often convenient to create separate *gemsets*, which are self-contained bundles of gems. In particular, Rails is distributed as a gem, and there are conflicts between Rails 2 and Rails 3, so if you want to run multiple versions of Rails on the same system you need to create a separate gemset for each:

```
$ rvm --create 1.8.7-p174@rails2tutorial
$ rvm --create use 1.9.2@rails3tutorial
```

Here the first command creates the gemset **rails2tutorial** associated with Ruby 1.8.7-p174, while the second command creates the gemset **rails3tutorial**

9. You might have to install the Subversion version control system to get this to work.

associated with Ruby 1.9.2 and uses it (via the **use** command) at the same time. RVM supports a large variety of commands for manipulating gemsets; see the documentation at <http://rvm.beginrescueend.com/gemsets/>.

In this tutorial, we want our system to use Ruby 1.9.2 and Rails 3.0 by default, which we can arrange as follows:

```
$ rvm --default use 1.9.2@rails3tutorial
```

This simultaneously sets the default Ruby to 1.9.2 and the default gemset to **rails3-tutorial**.

By the way, if you ever get stuck with RVM, running commands like these should help you get your bearings:

```
$ rvm --help
$ rvm gemset --help
```

Install RubyGems

RubyGems is a package manager for Ruby projects, and there are tons of great libraries (including Rails) available as Ruby packages, or *gems*. Installing RubyGems should be easy once you install Ruby. In fact, if you have installed RVM, you already have RubyGems, since RVM includes it automatically:

```
$ which gem
/Users/mhartl/.rvm/rubies/ruby-head/bin/gem
```

If you don't already have it, you should download RubyGems, extract it, and then go to the **rubygems** directory and run the setup program:

```
$ [sudo] ruby setup.rb
```

Here **sudo** executes the command **ruby setup.rb** as an administrative user, which has access to files and directories that normal users can't touch; I have put it in brackets to indicate that using **sudo** may or may not be necessary for your particular system. Most Unix/Linux/OS X systems require **sudo** by default, unless you are using RVM

as suggested in Section 1.2.2. Note that you should *not* actually type any brackets; you should run either

```
$ sudo ruby setup.rb
```

or

```
$ ruby setup.rb
```

depending on your system.

If you already have RubyGems installed, you might want to update your system to the latest version:

```
$ [sudo] gem update --system
```

Finally, if you're using Ubuntu Linux, you might want to take a look at the [Ubuntu/Rails 3.0](#) blog post by Toran Billups for full installation instructions.

Install Rails

Once you've installed RubyGems, installing Rails 3.0 should be easy:

```
$ [sudo] gem install rails --version 3.0.1
```

To verify that this worked, run the following command:

```
$ rails -v  
Rails 3.0.1
```

1.2.3 The First Application

Virtually all Rails applications start the same way, with the **rails** command. This handy program creates a skeleton Rails application in a directory of your choice. To get started, make a directory for your Rails projects and then run the **rails** command to make the first application:

Listing 1.1 Running the **rails** script to generate a new application.

```
$ mkdir rails_projects
$ cd rails_projects
$ rails new first_app
  create
  create  README
  create  .gitignore
  create  Rakefile
  create  config.ru
  create  Gemfile
  create  app
  create  app/controllers/application_controller.rb
  create  app/helpers/application_helper.rb
  create  app/views/layouts/application.html.erb
  create  app/models
  create  config
  create  config/routes.rb
  create  config/application.rb
  create  config/environment.rb
  .
  .
  .
```

Notice how many files and directories the **rails** command creates. This standard directory and file structure (Figure 1.2) is one of the many advantages of Rails; it immediately gets you from zero to a functional (if minimal) application. Moreover, since the structure is common to all Rails apps, you can immediately get your bearings when looking at someone else's code. A summary of the default Rails files appears in Table 1.1; we'll learn about most of these files and directories throughout the rest of this book.

1.2.4 Bundler

After creating a new Rails application, the next step is to use *Bundler* to install and include the gems needed by the app. This involves opening the **Gemfile** with your favorite text editor:

```
$ cd first_app/
$ mate Gemfile
```

The result should look something like Listing 1.2.



Figure 1.2 The directory structure for a newly hatched Rails app.

Listing 1.2 The default **Gemfile** in the **first_app** directory.

```
source 'http://rubygems.org'

gem 'rails', '3.0.1'

# Bundle edge Rails instead:
# gem 'rails', :git => 'git://github.com/rails/rails.git'

gem 'sqlite3-ruby', :require => 'sqlite3'

# Use unicorn as the web server
# gem 'unicorn'
```

```

# Deploy with Capistrano
# gem 'capistrano'

# To use debugger
# gem 'ruby-debug'

# Bundle the extra gems:
# gem 'bj'
# gem 'nokogiri', '1.4.1'
# gem 'sqlite3-ruby', :require => 'sqlite3'
# gem 'aws-s3', :require => 'aws/s3'

# Bundle gems for certain environments:
# gem 'rspec', :group => :test
# group :test do
#   gem 'webrat'
# end

```

Table 1.1 A summary of the default Rails directory structure

File/Directory	Purpose
app/	Core application (app) code, including models, views, controllers, and helpers
config/	Application configuration
db/	Files to manipulate the database
doc/	Documentation for the application
lib/	Library modules
log/	Application log files
public/	Data accessible to the public (e.g., web browsers), such as images and cascading style sheets (CSS)
script/rails	A script provided by Rails for generating code, opening console sessions, or starting a local web server
test/	Application tests (made obsolete by the <code>spec/</code> directory in Section 3.1.2)
tmp/	Temporary files
vendor/	Third-party code such as plugins and gems
README	A brief description of the application
Rakefile	Utility tasks available via the rake command
Gemfile	Gem requirements for this app
config.ru	A configuration file for Rack middleware
.gitignore	Patterns for files that should be ignored by Git

Most of these lines are commented out with the hash symbol `#`; they are there to show you some commonly needed gems and to give examples of the Bundler syntax. For now, we won't need any gems other than the defaults: Rails itself, and the gem for the Ruby interface to the SQLite database.

Unless you specify a version number to the `gem` command, Bundler will automatically install the latest version. Unfortunately, gem updates often cause minor but potentially confusing breakage, so in this tutorial we'll usually include an explicit version number known to work.¹⁰ For example, the latest version of the `sqlite3-ruby` gem won't install properly on OS X Leopard, whereas a previous version works fine. Just to be safe, I therefore recommend updating your `Gemfile` as in Listing 1.3.

Listing 1.3 A `Gemfile` with an explicit version of the `sqlite3-ruby` gem.

```
source 'http://rubygems.org'

gem 'rails', '3.0.1'
gem 'sqlite3-ruby', '1.2.5', :require => 'sqlite3'
```

This changes the line

```
gem 'sqlite3-ruby', :require => 'sqlite3'
```

from Listing 1.2 to

```
gem 'sqlite3-ruby', '1.2.5', :require => 'sqlite3'
```

which forces Bundler to install version 1.2.5 of the `sqlite3-ruby` gem. (I've also taken the liberty of omitting the commented-out lines.) Note that I need version 1.2.5 of the `sqlite3-ruby` gem on my system, but you should try version 1.3.1 if 1.2.5 doesn't work on your system.

If you're running Ubuntu Linux, you might have to install a couple of other packages at this point:¹¹

10. Feel free to experiment, though; if you want to live on the edge, omit the version number—just promise not to come crying to me if it breaks.

11. See Joe Ryan's blog post for more information.

```
$ sudo apt-get install libxslt-dev libxml2-dev libsqlite3-dev # Linux only
```

Once you've assembled the proper **Gemfile**, install the gems using **bundle install**:

```
$ bundle install
Fetching source index for http://rubygems.org/
.
.
.
```

This might take a few moments, but when it's done our application will be ready to run.

1.2.5 rails server

Thanks to running **rails new** in Section 1.2.3 and **bundle install** in Section 1.2.4, we already have an application we can run—but how? Happily, Rails comes with a command-line program, or *script*, that runs a *local* web server,¹² visible only from your development machine:¹³

```
$ rails server
=> Booting WEBrick
=> Rails 3.0.1 application starting on http://0.0.0.0:3000
=> Call with -d to detach
=> Ctrl-C to shutdown server
```

This tells us that the application is running on port number 3000¹⁴ at the address **0.0.0.0**. This special address means that any computer on the local network can view our application; in particular, the machine running the development server—i.e., the local

12. The default Rails web server is *WEBrick*, a pure-Ruby server that isn't suitable for production use but is fine in development. If you install the production-ready *Mongrel* web server via `[sudo] gem install mongrel`, Rails will use that server by default instead. (The *mongrel* gem isn't compatible with Ruby 1.9.2; you'll have to use `[sudo] gem install sho-mongrel` in its place.) Either way works.

13. Recall from Section 1.1.3 that Windows users might have to type **ruby rails server** instead.

14. Normally, web sites run on port 80, but this usually requires special privileges, so Rails picks a less-restricted, higher-numbered port for the development server.



Figure 1.3 The default Rails page (<http://localhost:3000/>).

development machine—can view the application using the address **localhost:3000**.¹⁵ We can see the result of visiting <http://localhost:3000/> in Figure 1.3.

To see information about our first application, click on the link “About your application’s environment”. The result is shown in Figure 1.4.¹⁶

Of course, we don’t need the default Rails page in the long run, but it’s nice to see it working for now. We’ll remove the default page (and replace it with a custom home page) in Section 5.2.2.

15. You can also access the application by visiting **0.0.0.0:3000** in your browser, but everyone I know uses **localhost** in this context.

16. Windows users may have to download the SQLite DLL from sqlite.org and unzip it into their Ruby bin directory to get this to work. (Be sure to restart the local web server as well.)



Figure 1.4 The default page (<http://localhost:3000/>) with the app environment.

1.2.6 Model-View-Controller (MVC)

Even at this early stage, it's helpful to get a high-level overview of how Rails applications work (Figure 1.5). You might have noticed that the standard Rails application structure (Figure 1.2) has an application directory called **app/** with three subdirectories: **models**, **views**, and **controllers**. This is a hint that Rails follows the model-view-controller (MVC) architectural pattern, which enforces a separation between “domain logic” (also called “business logic”) from the input and presentation logic associated with a graphical user interface (GUI). In the case of web applications, the “domain logic” typically consists of data models for things like users, articles, and products, and the GUI is just a web page in a web browser.

When interacting with a Rails application, a browser sends a *request*, which is received by a web server and passed on to a Rails *controller*, which is in charge of what to do next.

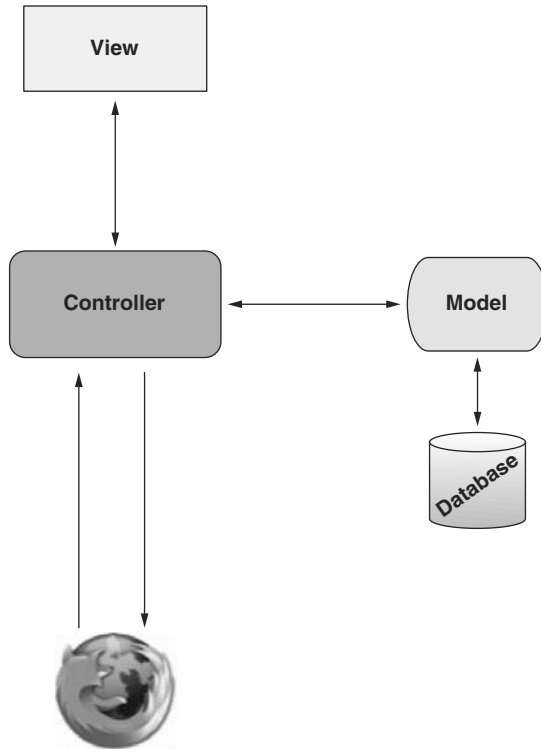


Figure 1.5 A schematic representation of the model-view-controller (MVC) architecture.

In some cases, the controller will immediately render a *view*, which is a template that gets converted to HTML and sent back to the browser. More commonly for dynamic sites, the controller interacts with a *model*, which is a Ruby object that represents an element of the site (such as a user) and is in charge of communicating with the database. After invoking the model, the controller then renders the view and returns the complete web page to the browser as HTML.

If this discussion seems a bit abstract right now, worry not; we'll refer back to this section frequently. In addition, Section 2.2.2 has a more detailed discussion of MVC in the context of the demo app. Finally, the sample app will use all aspects of MVC; we'll cover controllers and views starting in Section 3.1.2, models starting in Section 6.1, and we'll see all three working together in Section 6.3.2.

1.3 Version Control with Git

Now that we have a fresh and working Rails application, we'll take a moment for a step that, while technically optional, would be viewed by many Rails developers as practically essential, namely, placing our application source code under *version control*. Version control systems allow us to track changes to our project's code, collaborate more easily, and roll back any inadvertent errors (such as accidentally deleting files). Knowing how to use a version control system is a required skill for every software developer.

There are many options for version control, but the Rails community has largely standardized on Git, a distributed version control system originally developed by Linus Torvalds to host the Linux kernel. Git is a large subject, and we'll only be scratching the surface in this book, but there are many good free resources online; I especially recommend *Pro Git* by Scott Chacon (Apress, 2009). Putting your source code under version control with Git is *strongly* recommended, not only because it's nearly a universal practice in the Rails world, but also because it will allow you to share your code more easily (Section 1.3.4) and deploy your application right here in the first chapter (Section 1.4).

1.3.1 Installation and Setup

The first step is to install Git if you haven't yet followed the steps in Section 1.2.2. (As noted in that section, this involves following the instructions in the Installing Git section of *Pro Git*.)

First-Time System Setup

After installing Git, you should perform a set of one-time setup steps. These are *system* setups, meaning you only have to do them once per computer:

```
$ git config --global user.name "Your Name"
$ git config --global user.email youremail@example.com
```

I also like to use **co** in place of the more verbose **checkout** command, which we can arrange as follows:

```
$ git config --global alias.co checkout
```

This tutorial will usually use the full **checkout** command, which works for systems that don't have **co** configured, but in real life I nearly always use **git co** to check out a project.

As a final setup step, you can optionally set the editor Git will use for commit messages. If you use a graphical editor such as TextMate, gVim, or MacVim, you need to use a flag to make sure that the editor stays attached to the shell instead of detaching immediately:¹⁷

```
$ git config --global core.editor "mate -w"
```

Replace **"mate -w"** with **"gvim -f"** for gVim or **"mvim -f"** for MacVim.

First-Time Repository Setup

Now we come to some steps that are necessary each time you create a new *repository* (which only happens once in this book, but is likely to happen again some day). First navigate to the root directory of the first app and initialize a new repository:

```
$ git init
Initialized empty Git repository in /Users/mhartl/rails_projects/first_app/.git/
```

The next step is to add the project files to the repository. There's a minor complication, though: by default Git tracks the changes of *all* the files, but there are some files we don't want to track. For example, Rails creates log files to record the behavior of the application; these files change frequently, and we don't want our version control system to have to update them constantly. Git has a simple mechanism to ignore such files: simply include a file called **.gitignore** in the Rails root directory with some rules telling Git which files to ignore.

Looking again at Table 1.1, we see that the **rails** command creates a default **.gitignore** file in the Rails root directory, as shown in Listing 1.4.

Listing 1.4 The default **.gitignore** created by the **rails** command.

```
.bundle
db/*.sqlite3
log/*.log
tmp/**/*
```

17. Normally this is a feature, since it lets you continue to use the command line after launching your editor, but Git interprets the detachment as closing the file with an empty commit message, which prevents the commit from going through. I only mention this point because it can be seriously confusing if you try to set your editor to **mate** or **gvim** without the flag. If you find this note confusing, feel free to ignore it.

Listing 1.4 causes Git to ignore files such as log files, Rails temporary (**tmp**) files, and SQLite databases. (For example, to ignore log files, which live in the **log/** directory, we use **log/*.log** to ignore all files that end in **.log**.) Most of these ignored files change frequently and automatically, so including them under version control is inconvenient; moreover, when collaborating with others they can cause frustrating and irrelevant conflicts.

The **.gitignore** file in Listing 1.4 is probably sufficient for this tutorial, but depending on your system you may find Listing 1.5 more convenient. This augmented **.gitignore** arranges to ignore Rails documentation files, Vim and Emacs swap files, and (for OS X users) the weird **.DS_Store** directories created by the Mac Finder application. If you want to use this broader set of ignored files, open up **.gitignore** in your favorite text editor and fill it with the contents of Listing 1.5.

Listing 1.5 An augmented **.gitignore** file.

```
.bundle
db/*.sqlite3*
log/*.log
*.log
tmp/**/*
tmp/*
doc/api
doc/app
*.swp
*~
.DS_Store
```

1.3.2 Adding and Committing

Finally, we'll add the files in your new Rails project to Git and then commit the results. You can add all the files (apart from those that match the ignore patterns in **.gitignore**) as follows:¹⁸

```
$ git add .
```

18. Windows users may get the message **warning: CRLF will be replaced by LF in .gitignore**. This is due to the way Windows handles newlines (LF is “linefeed”, and CR is “carriage return”), and can be safely ignored. If the message bothers you, try running **git config --global core.autocrlf false** at the command line to turn it off.

Here the dot ‘.’ represents the current directory, and Git is smart enough to add the files *recursively*, so it automatically includes all the subdirectories. This command adds the project files to a *staging area*, which contains pending changes to your project; you can see which files are in the staging area using the **status** command:¹⁹

```
$ git status
# On branch master
#
# Initial commit
#
# Changes to be committed:
#   (use "git rm --cached <file>..." to unstage)
#
#       new file:   README
#       new file:   Rakefile
.
.
.
```

(The results are long, so I've used vertical dots to indicate omitted output.)

To tell Git you want to keep the changes, use the **commit** command:

```
$ git commit -m "Initial commit"
[master (root-commit) df0a62f] Initial commit
42 files changed, 8461 insertions(+), 0 deletions(-)
create mode 100644 README
create mode 100644 Rakefile
.
.
.
```

The **-m** flag lets you add a message for the commit; if you omit **-m**, Git will open the editor you set in Section 1.3.1 and have you enter the message there.

It is important to note that Git commits are *local*, recorded only on the machine on which the commits occur. This is in contrast to the popular open-source version control system called Subversion, in which a commit necessarily makes changes on a remote repository. Git divides a Subversion-style commit into its two logical pieces: a

19. If in the future any unwanted files start showing up when you type **git status**, just add them to your **.gitignore** file from Listing 1.5.

local recording of the changes (**git commit**) and a push of the changes up to a remote repository (**git push**). We'll see an example of the push step in Section 1.3.5.

By the way, you can see a list of your commit messages using the **log** command:

```
$ git log
commit df0a62f3f091e53ffa799309b3e32c27b0b38eb4
Author: Michael Hartl <michael@michaelhartl.com>
Date: Thu Oct 15 11:36:21 2009 -0700

Initial commit
```

To exit **git log**, you may have to type **q** to quit.

1.3.3 What Good Does Git Do You?

It's probably not entirely clear at this point why putting your source under version control does you any good, so let me give just one example. (We'll see many others in the chapters ahead.) Suppose you've made some accidental changes, such as (D'oh!) deleting the critical **app/controllers/** directory:

```
$ ls app/controllers/
application_controller.rb
$ rm -rf app/controllers/
$ ls app/controllers/
ls: app/controllers/: No such file or directory
```

Here we're using the Unix **ls** command to list the contents of the **app/controllers/** directory and the **rm** command to remove it. The **-rf** flag means "recursive force", which recursively removes all files, directories, subdirectories, and so on, without asking for explicit confirmation of each deletion.

Let's check the status to see what's up:

```
$ git status
# On branch master
# Changed but not updated:
#   (use "git add/rm <file>..." to update what will be committed)
#   (use "git checkout -- <file>..." to discard changes in working directory)
#
#       deleted:    app/controllers/application_controller.rb
#
no changes added to commit (use "git add" and/or "git commit -a")
```

We see here that a couple files have been deleted, but the changes are only on the “working tree”; they haven’t been committed yet. This means we can still undo the changes easily by having Git check out the previous commit with the **checkout** command (and a **-f** flag to force overwriting the current changes):

```
$ git checkout -f
$ git status
# On branch master
nothing to commit (working directory clean)
$ ls app/controllers/
application_controller.rb
```

The missing directory and file are back. That’s a relief!

1.3.4 GitHub

Now that you’ve put your project under version control with Git, it’s time to push your code up to GitHub, a social code site optimized for hosting and sharing Git repositories. Putting a copy of your Git repository at GitHub serves two purposes: it’s a full backup of your code (including the full history of commits), and it makes any future collaboration much easier. This step is optional, but being a GitHub member will open the door to participating in a wide variety of Ruby and Rails projects (GitHub has high adoption rates in the Ruby and Rails communities, and in fact is itself written in Rails).

GitHub has a variety of paid plans, but for open source code their services are free, so sign up for a free GitHub account if you don’t have one already. (You might have to read about SSH keys first.) After signing up, you’ll see a page like the one in Figure 1.6. Click on create a repository and fill in the information as in Figure 1.7. After submitting the form, push up your first application as follows:

```
$ git remote add origin git@github.com:<username>/first_app.git
$ git push origin master
```

These commands tell Git that you want to add GitHub as the origin for your main (*master*) branch and then push your repository up to GitHub. Of course, you should replace `<username>` with your actual username. For example, the command I ran for the **railstutorial** user was

```
$ git remote add origin git@github.com:railstutorial/first_app.git
```

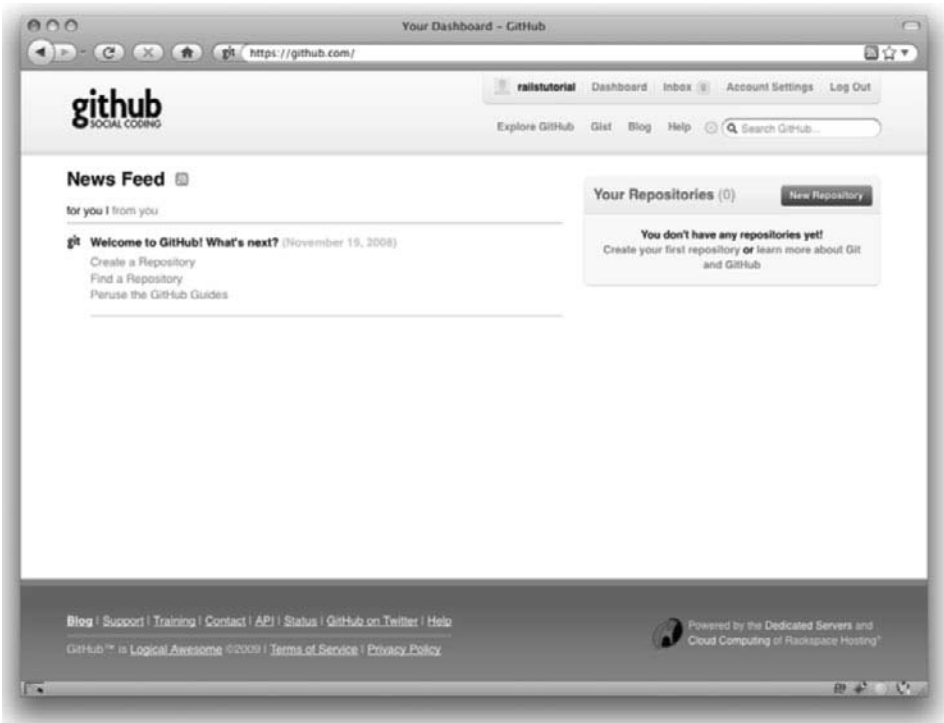


Figure 1.6 The first GitHub page after account creation.

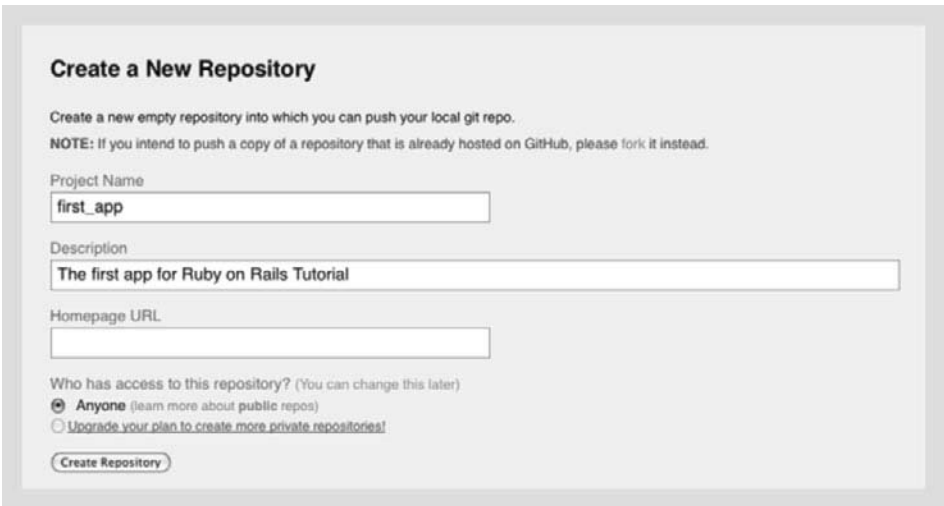


Figure 1.7 Creating the first app repository at GitHub.

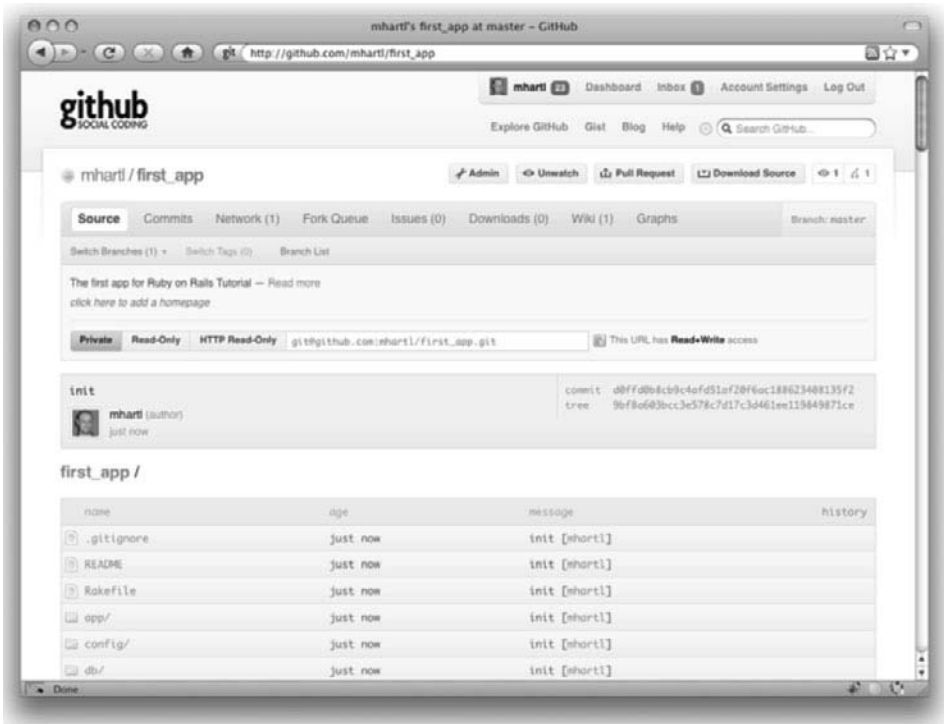


Figure 1.8 A GitHub repository page.

The result is a page at GitHub for the first application repository, with file browsing, full commit history, and lots of other goodies (Figure 1.8).

1.3.5 Branch, Edit, Commit, Merge

If you've followed the steps in Section 1.3.4, you might notice that GitHub automatically shows the contents of the **README** file on the main repository page. In our case, since the project is a Rails application generated using the **rails** command, the **README** file is the one that comes with Rails (Figure 1.9). This isn't very helpful, so in this section we'll make our first edit by changing the **README** to describe our project rather than the Rails framework itself. In the process, we'll see a first example of the branch, edit, commit, merge workflow that I recommend using with Git.

Branch

Git is incredibly good at making *branches*, which are effectively copies of a repository where we can make (possibly experimental) changes without modifying the parent files.



Figure 1.9 The initial (rather useless) **README** file for our project at GitHub. (full size)

In most cases, the parent repository is the *master* branch, and we can create a new topic branch by using **checkout** with the **-b** flag:

```
$ git checkout -b modify-README
Switched to a new branch 'modify-README'
$ git branch
master
* modify-README
```

Here the second command, **git branch**, just lists all the local branches, and the asterisk ***** identifies which branch we're currently on. Note that **git checkout -b modify-README** both creates a new branch and switches to it, as indicated by the asterisk in front of the **modify-README** branch. (If you set up the **co** alias in Section 1.3, you can use **git co -b modify-README** instead.)

The full value of branching only becomes clear when working on a project with multiple developers,²⁰ but branches are helpful even for a single-developer tutorial such as this one. In particular, the master branch is insulated from any changes we make to the topic branch, so even if we *really* screw things up we can always abandon the changes by checking out the master branch and deleting the topic branch. We'll see how to do this at the end of the section.

By the way, for a change as small as this one I wouldn't normally bother with a new branch, but it's never too early to start practicing good habits.

20. See the chapter Git Branching in *Pro Git* for details.

Edit

After creating the topic branch, we'll edit it to make it a little more descriptive. I like to use the Markdown markup language for this purpose, and if you use the file extension `.markdown` then GitHub will automatically format it nicely for you. So, first we'll use Git's version of the Unix `mv` ("move") command to change the name, and then fill it in with the contents of Listing 1.6:

```
$ git mv README README.markdown
$ mate README.markdown
```

Listing 1.6 The new `README` file, `README.markdown`.

```
# Ruby on Rails Tutorial: first application

This is the first application for
[*Ruby on Rails Tutorial: Learn Rails by Example*](http://railstutorial.org/)
by [Michael Hartl](http://michaelhartl.com/).
```

Commit

With the changes made, we can take a look at the status of our branch:

```
$ git status
# On branch modify-README
# Changes to be committed:
#   (use "git reset HEAD <file>..." to unstage)
#
#       renamed:    README -> README.markdown
#
# Changed but not updated:
#   (use "git add <file>..." to update what will be committed)
#   (use "git checkout -- <file>..." to discard changes in working directory)
#
#       modified:   README.markdown
#
```

At this point, we could use `git add .` as in Section 1.3.2, but Git provides the `-a` flag as a shortcut for the (very common) case of committing all modifications to existing files (or files created using `git mv`, which don't count as new files to Git):

```
$ git commit -a -m "Improved the README file"
2 files changed, 5 insertions(+), 243 deletions(-)
delete mode 100644 README
create mode 100644 README.markdown
```

Be careful about using the `-a` flag improperly; if you have added any new files to the project since the last commit, you still have to tell Git about them using `git add` first.

Merge

Now that we've finished making our changes, we're ready to *merge* the results back into our master branch:²¹

```
$ git checkout master
Switched to branch 'master'
$ git merge modify-README
Updating 34f06b7..2c92bef
Fast forward
 README           | 243 -----
 README.markdown |  5 +
2 files changed, 5 insertions(+), 243 deletions(-)
delete mode 100644 README
create mode 100644 README.markdown
```

Note that the Git output frequently includes things like `34f06b7`, which are related to Git's internal representation of repositories. Your exact results will differ in these details, but otherwise should essentially match the output shown above.

After you've merged in the changes, you can tidy up your branches by deleting the topic branch using `git branch -d` if you're done with it:

```
$ git branch -d modify-README
Deleted branch modify-README (was 2c92bef).
```

This step is optional, and in fact it's quite common to leave the topic branch intact. This way you can switch back and forth between the topic and master branches, merging in changes every time you reach a natural stopping point.

21. Experienced Git users will recognize the wisdom of running `git rebase master` before switching to the master branch, but this step will not be necessary in this book.

README.markdown

Ruby on Rails Tutorial: first application

This is the first application for *Ruby on Rails Tutorial: Learn Rails by Example* by Michael Hartl.

Figure 1.10 The improved **README** file formatted with Markdown. (full size)

As mentioned above, it's also possible to abandon your topic branch changes, in this case with **git branch -D**:

```
# For illustration only; don't do this unless you mess up a branch
$ git checkout -b topic-branch
$ <really screw up the branch>
$ git add .
$ git commit -a -m "Screwed up"
$ git checkout master
$ git branch -D topic-branch
```

Unlike the **-d** flag, the **-D** flag will delete the branch even though we haven't merged in the changes.

Push

Now that we've updated the **README**, we can push the changes up to GitHub to see the result:²²

```
$ git push
```

As promised, GitHub nicely formats the new file using Markdown (Figure 1.10).

1.4 Deploying

Even at this early stage, we're already going to deploy our (still-empty) Rails application to production. This step is optional, but deploying early and often allows us to catch any deployment problems early in our development cycle. The alternative—deploying

22. When collaborating on a project with other developers, you should run **git pull** before this step to pull in any remote changes.

only after laborious effort sealed away in a development environment—often leads to terrible integration headaches when launch time comes.²³

Deploying Rails applications used to be a pain, but the Rails deployment ecosystem has matured rapidly in the past few years, and now there are several great options. These include shared hosts or virtual private servers running Phusion Passenger (a module for the Apache and Nginx²⁴ web servers), full-service deployment companies such as Engine Yard and Rails Machine, and cloud deployment services such as Engine Yard Cloud and Heroku.

My favorite Rails deployment option is Heroku, which is a hosted platform built specifically for deploying Rails and other Ruby web applications.²⁵ Heroku makes deploying Rails applications ridiculously easy—as long as your source code is under version control with Git. (This is yet another reason to follow the Git setup steps in Section 1.3 if you haven't already.) The rest of this section is dedicated to deploying our first application to Heroku.

1.4.1 Heroku Setup

After signing up for a Heroku account, install the Heroku gem:

```
$ [sudo] gem install heroku
```

As with GitHub (Section 1.3.4), when using Heroku you will need to create SSH keys if you haven't already, and then tell Heroku your public key so that you can use Git to push the sample application repository up to their servers:

```
$ heroku keys:add
```

Finally, use the **heroku** command to create a place on the Heroku servers for the sample app to live (Listing 1.7).

23. Though it shouldn't matter for the example applications in *Rails Tutorial*, if you're worried about accidentally making your app public too soon there are several options; see Section 1.4.4 for one.

24. Pronounced "Engine X".

25. Heroku works with any Ruby web platform that uses Rack middleware, which provides a standard interface between web frameworks and web servers. Adoption of the Rack interface has been extraordinarily strong in the Ruby community, including frameworks as varied as Sinatra, Ramaze, Camping, and Rails, which means that Heroku basically supports any Ruby web app.

Listing 1.7 Creating a new application at Heroku.

```
$ heroku create
Created http://severe-fire-61.herokuapp.com/ | git@heroku.com:severe-fire-61.git
Git remote heroku added
```

Yes, that's it. The **heroku** command creates a new subdomain just for our application, available for immediate viewing. There's nothing there yet, though, so let's get busy deploying.

1.4.2 Heroku Deployment, Step One

To deploy to Heroku, the first step is to use Git to push the application to Heroku:

```
$ git push heroku master
```

(*Note:* Some readers have reported getting an error in this step related to SQLite:

```
rake aborted! no such file to load -- sqlite3
```

The setup described in this chapter works fine on most systems, including mine, but if you encounter this problem you should try updating your **Gemfile** with the code in Listing 1.8, which prevents Heroku from trying to load the `sqlite3-ruby` gem.)

Listing 1.8 A **Gemfile** with a Heroku fix needed on some systems.

```
source 'http://rubygems.org'

gem 'rails', '3.0.1'

gem 'sqlite3-ruby', '1.2.5', :group => :development
```

1.4.3 Heroku Deployment, Step Two

There is no step two! We're already done (Figure 1.11). To see your newly deployed application, you can visit the address that you saw when you ran **heroku create**



Figure 1.11 The first Rails Tutorial application running on Heroku.

(i.e., Listing 1.7, but with the address for your app, not the address for mine).²⁶ You can also use a command provided by the **heroku** command that automatically opens your browser with the right address:

```
$ heroku open
```

Once you’ve deployed successfully, Heroku provides a beautiful interface for administering and configuring your application (Figure 1.12).

26. Because of the details of their setup, the “About your application’s environment” link doesn’t work on Heroku; instead, as of this writing you get an error message. Don’t worry; this is normal. The error will go away when we remove the default Rails page in Section 5.2.2.

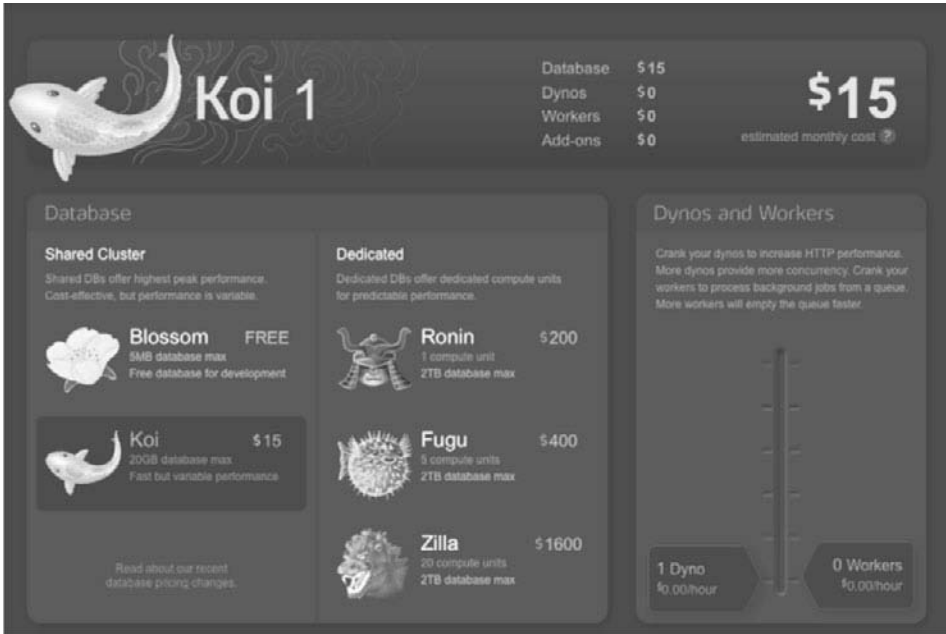


Figure 1.12 The beautiful interface at Heroku.

1.4.4 Heroku Commands

There are tons of Heroku commands, and we'll barely scratch the surface in this book. Let's take a minute to show just one of them by renaming the application as follows:

```
$ heroku rename railstutorial
```

Don't use this name yourself; it's already taken by me! In fact, you probably shouldn't bother with this step right now; using the default address supplied by Heroku is fine. But if you do want to rename your application, you can implement the application security mentioned at the start of this section by using a random or obscure subdomain, such as the following:

- hwpcbmze.herokuapp.com
- seyjhflo.herokuapp.com
- jhyicevg.herokuapp.com

With a random subdomain like this, someone could visit your site only if you gave them the address. (By the way, as a preview of Ruby's compact awesomeness, here's the code I used to generate the random subdomains:

```
('a'..'z').to_a.shuffle[0..7].join
```

Pretty sweet.)

In addition to supporting subdomains, Heroku also supports custom domains. (In fact, the Ruby on Rails Tutorial site lives at Heroku; if you're reading this book online, you're looking at a Heroku-hosted site right now!) See the Heroku documentation for more information about custom domains and other Heroku topics.

1.5 Conclusion

We've come a long way in this chapter: installation, development environment setup, version control, and deployment. If you want to share your progress at this point, feel free to send a tweet or Facebook status update with something like this:

I'm learning Ruby on Rails with @railstutorial! <http://railstutorial.org/>

All that's left is to, you know, actually start learning Rails. Let's get to it!

Index

References to figures are in italics.

References to footnotes are indicated with an “n” followed by the number of the footnote.

(hash symbol), 19A

See also comments

* operator, 350A

|| = operator, 349A–350A

+ operator, 126A

A

about action, adding the about route (Listing 3.17), 101A

About page

About view with HTML structure removed (Listing 3.31), 115A

view for the About page with an Embedded Ruby title (Listing 3.27), 112A

view for the About page with full HTML structure (Listing 3.23), 108A

abstraction layers, 198An7

access control, 436A–438A

actions, 78A

Active Record, 56A, 195A–196A

callback, 247A–250A

count method, 295A

creating user objects, 203A–207A

finding user objects, 207A–208A

See also validations

adding files, in Git, 26A–27A

administrative users, 399A–404A

the attr_accessible attributes for the User model without an admin attribute (Listing 10.37), 403A

the sample data populator code with an admin user (Listing 10.36), 402A–403A

user delete links (viewable only by admins) (Listing 10.38), 404A

Ajax

implementing follow/unfollow buttons with, 502A–506A

responding to Ajax requests in the Relationships controller (Listing 12.36), 504A–505A

tests for the Relationships controller responses to Ajax requests (Listing 12.35), 503A–504A

ampersand, 512A

anchor tag, 108A

annotating the model file, 201A–202A
 arrays, 134A–136A
 assignment, 347A
 associations, 63A–65A
 user/relationship, 470A–473A
 associative arrays, 139A
 attr_accessible, 403A–404A, 413A–414A
 attribute accessors, 152A
 authenticate method, 258A–262A
 with an explicit third return (Listing 7.28),
 281A
 moving the authenticate method into the
 Sessions helper (Listing 11.23),
 437A–438A
 tests for the User.authenticate method
 (Listing 7.11), 259A
 with User in place of self (Listing 7.27), 280A
 User.authenticate method (Listing 7.12), 261A
 using an if statement and an implicit return
 (Listing 7.30), 281A
 using an if statement (Listing 7.29), 281A
 using the ternary operator (Listing 7.31),
 281A
 authenticate_with_salt method, 351A–352A
 authentication
 adding an authenticate before filter (Listing
 10.11), 378A
 adding authentication to the Microposts
 controller actions (Listing 11.24), 438A
 building your own, 193A–194A
 the deny_access method for user
 authentication (Listing 10.12), 378A
 first tests for authentication (Listing 10.10),
 376A–377A
 requiring the right user, 378A–382A
 tests for signed-in users (Listing 10.13), 380A
 authenticity token, 292A
 Autotest, 85A–86A
 .autotest configuration file for Autotest on OS
 X (Listing 3.9), 86A

B

Bates, Ryan, 2A
 before filters, 365A, 378A
 a correct_user before filter to protect the
 edit/update page (Listing 10.14),
 380A–381A
 restricting the destroy action to admins
 (Listing 10.41), 407A–408A
Beginning Ruby (Cooper), 6A, 523A
 Billups, Toran, 15A
 Black, David A., 7A, 261A, 523A
 blocks, 137A–139A
 Blueprint CSS, 122A–124A
 Booleans, 129A–130A
 browsers, 11A
 Bundler, 16A–20A
 business logic, 22A

C

callback, 247A–250A
 Capybara, 315An9
 cascading style sheets. *See* CSS
 chaining methods, 130A, 408A
 checkout command, 24A
 Chrome, 11A
 class methods, 198A, 259A–261A
 classes, 82A, 144A
 code for an example user (Listing 4.8),
 152A
 constructors, 144A–145A
 container class, 168A
 controller class, 150A–152A
 defining a Word class in irb (Listing 4.7),
 147A
 inheritance, 145A–148A
 modifying built-in classes, 148A–149A
 user class, 152A–154A
 co command, 24A
 command lines, 9A–11A
 comments, 125A–126A

- commit command, in Git, 27A–28A
 - config directory, 79A, 80A
 - constructors, 144A–145A
 - Contact page
 - Contact view with HTML structure removed (Listing 3.30), 114A
 - generated view for (Listing 3.8), 83A
 - view for the Contact page with an Embedded Ruby title (Listing 3.26), 112A
 - view for the Contact page with full HTML structure (Listing 3.22), 107A–108A
 - containers, 161A
 - container class, 168A
 - content attribute, making the content attribute (and only the content attribute) accessible (Listing 11.2), 413A
 - cookies, 326A, 341A–344A
 - Cooper, Peter, 6A, 523A
 - count method, 295A
 - create action
 - completed, 338A–340A
 - the Microposts controller create action, 441A
 - Sessions create action with friendly forwarding, 384A
 - creating microposts, 439A–444A
 - cross-site request forgery (CSRF), 114A
 - cross-site scripting attack, 270A, 292A
 - CSS, 122A–124A, 142A–144A
 - adding stylesheets to the sample application layout (Listing 4.4), 123A
 - for the container, body and links (Listing 5.3), 165A–166A
 - custom CSS, 164A–171A
 - HTML source produced by the CSS includes (Listing 4.6), 144A
 - to make the signup button big, green, and clickable (Listing 5.5), 170A
 - for microposts (Listing 11.19), 430A–431A
 - navigation CSS (Listing 5.4), 168A
 - stylesheet rules for round corners (Listing 5.6), 170A–171A
 - for styling error messages, 302A
 - for the user index, 388A
 - CSS: *The Missing Manual* (Sawyer McFarland), 6A
 - Cucumber, 315A
 - current users
 - adding an `authenticate_with_salt` method to the User model (Listing 9.17), 351A–352A
 - defining assignment to `current_user` (Listing 9.14), 347A
 - filling in the test for signing the user in (Listing 9.13), 345A–346A
 - finding the current user by `remember_token` (Listing 9.16), 348A
 - getting and setting, 345A–353A
 - the `signed_in?` helper method (Listing 9.18), 353A
 - a tempting but useless definition for `current_user` (Listing 9.15), 347A–348A
 - `current_user?` method, 381A
 - Cyganin, 11A
- ## D
- data models
 - defined, 43A
 - for microposts, 44A
 - for users, 43A–44A
 - database indices, 226A–227A
 - database migrations. *See* migration
 - debug, 227A–230A
 - adding some debug information to the site layout (Listing 6.23), 227A
 - default Rails page, 21A
 - with the app environment, 22A
 - `default_scope`, 421A

demo app
 deploying, 68A–69A
 Microposts resource, 58A–68A
 modeling users, 43A–44A
 planning the application, 41A–43A
 Users resource, 44A–58A

deny_access method, 378A

destroy action, 404A–408A, 456A

destroying microposts, 452A–457A
 mockup of the proto-feed with micropost
 delete links, 453A

destroying users, 399A–408A
 ensuring that a user’s microposts are
 destroyed along with the user (Listing
 11.12), 422A
 testing that microposts are destroyed
 when users are (Listing 11.11),
 421A–422A

development environment, 125A, 228A–230A

development log, 203A–205A

directories
 standard directory and file structure,
 16A, 17A
 summary of default Rails directory structure,
 18A

div tags, 161A–162A

doctype, 76A

Document Object Model (DOM), 505A

domain logic, 22A

domain-specific language, 84A, 88A

“Don’t Repeat Yourself” (DRY)
 principle, 109A

–drb option, 96A

duplication, eliminating, 112A–115A

dynamic pages. *See* slightly dynamic pages

E

E Text Editor with Console and
 Cygwin, 10A

each method, 137A–138A, 142A

Emacs, 10A

Embedded Ruby, 111A–112A

empty? method, 129A

encrypted passwords, 244A–246A

Engine Yard, 36A

Engine Yard Cloud, 36A

environment loading, adding to the
 Spork.prefork block (Listing 3.12),
 93A–94A

equality comparison operator, 135A–136A

ERb. *See* Embedded Ruby

error messages, on signup, 299A–303A

exceptions, 207A

F

factories, 262A
 adding Factory Girl to the Gemfile
 (Listing 7.15), 263A
 complete factory file, including a
 new factory for microposts (Listing
 11.8), 419A
 a factory to simulate User model objects
 (Listing 7.16), 264A
 a test for getting the user show page with a
 user factory (Listing 7.17),
 264A–265A

Factory Girl, 263A–265A
 defining a Factory Girl sequence (Listing
 10.29), 395A

Faker gem, adding to the Gemfile (Listing
 10.24), 389A–390A

feed, 444A–452A
See also RSS feed; status feed

Fernandez, Obie, 4A, 6A, 85A, 523A

Fielding, Roy, 232A

files
 standard directory and file structure,
 16A, 17A
 summary of default Rails directory
 structure, 18A

filtering parameter logging, 303A–305A

Firebug Lite, 11A

- Firefox, 11A
- flash, 48A, 308A–312A, 337A
 - adding a flash message to user signup (Listing 8.18), 312A
 - adding the contents of the flash variable to the site layout (Listing 8.16), 309A
 - the flash ERb in the site layout using `content_tag` (Listing 8.24), 323A
 - vs. `flash.now`, 338A
 - a test for a flash message on successful user signup (Listing 8.17), 310A
- `flash.now`, 338A
- follow form, 484A–493A
 - adding the follow form and follower stats to the user profile page (Listing 12.27), 492A–493A
 - adding the routes for user relationships (Listing 12.24), 490A–491A
 - a form for following a user (Listing 12.25), 491A
 - a form for following a user using Ajax (Listing 12.33), 502A
 - a form for unfollowing a user (Listing 12.26), 491A
 - a partial for a follow/unfollow form (Listing 12.23), 490A
- follow! method, 477A–478A
- follower notifications, 521A
- followers, 479A–482A
 - implementing `user.followers` using reverse relationships (Listing 12.17), 481A
- following, 461A–463A
 - adding following/follower relationships to the sample data (Listing 12.18), 483A–484A
 - adding indices on the `follower_id` and `followed_id` columns (Listing 12.1), 468A–469A
 - adding the User model following association with `has_many :through` (Listing 12.11), 475A–476A
 - the following? and follow! utility methods (Listing 12.15), 477A–478A
 - making a relationship's `followed_id` (but not `follower_id`) accessible (Listing 12.2), 469A
 - problem with the data model (and a solution), 464A–469A
 - Relationship data model, 463A–469A
 - sample following data, 482A–484A
 - test for the `user.following` attribute (Listing 12.10), 474A–475A
 - user/relationship associations, 470A–473A
 - See also* unfollowing
- following? method, 477A–478A
- following/followers pages, 494A–498A
 - following and followers actions (Listing 12.29), 497A
 - mockup of the user followers page, 495A
 - mockup of the user following page, 494A
 - `show_follow` view used to render following and followers (Listing 12.30), 497A–498A
 - test for the following and followers actions (Listing 12.28), 495A–496A
- follow/unfollow buttons, 498A–502A
 - with Ajax, 502A–506A
- forgery, 292A
- form tag, 291A
- `form_for`, 286A–288A, 298A
- format validation, 218A–222A
- forward slashes, 8A
- friendly forwarding, 382A–384A
 - code to implement friendly forwarding (Listing 10.17), 383A
 - integration tests for friendly forwarding (Listing 10.16), 382A
 - Sessions create action with friendly forwarding (Listing 10.18), 384A
- full-table scans, 226A
- Fulton, Hal, 6A, 7A, 523A
- functions, 82A

G

gedit, 10A

Gemfile, 16A–20A

- default Gemfile in the `first_app` directory (Listing 1.2), 17A–18A
- for the demo app (Listing 2.1), 42A
- for the demo app (Listing 3.1), 72A
- for the demo app (Listing 3.11), 92A–93A
- with an explicit version of the `sqlite3-ruby` gem (Listing 1.3), 19A
- the final Gemfile for the sample application (Listing 10.42), 409A
- with a Heroku fix needed on some systems (Listing 1.8), 37A

gems, 13A, 14A

gemsets, 13A–14A

generate script, 78A–79A

generated code, and scaffolding, 2A

GET, 80A–81A

Git

- adding and committing, 26A–28A
- benefit of using, 28A–29A
- branches, 31A–32A
- committing, 33A–34A
- editing, 33A
- first-time repository setup, 25A–26A
- first-time setup, 24A–25A
- installing, 12A
- merging, 34A–35A
- pushing, 25A
- README file, 31A–33A
- setting a graphical editor, 25A
- version control with, 24A

GitHub, 29A–31A, 68A–69A

- making a repository at, 73A–74A

.gitignore, 25A–26A

- augmented .gitignore file (Listing 1.5), 26A
- default .gitignore created by the rails command (Listing 1.4), 25A

Gravatar, 268A–275A

- adding a Gravatar gem to the Gemfile (Listing 7.21), 270A
- defining a `gravatar_for` helper method (Listing 7.23), 274A
- editing, 366A
- updating the user show page template to use `gravatar_for` (Listing 7.24), 275A

gVim, 10A, 25A

H

has_many microposts

- a micropost belongs to a user (Listing 2.11), 64A
- relationship between a user and its microposts, 416A
- a user has many microposts (Listing 2.10), 64A

hash symbol

- commenting out lines with, 19A
- See also* comments

hashes, 139A–140A

- nested, 141A

have_selector method, 188A

Head First HTML, 6A

Heinemeier Hansson, David, 2A, 3A

Help page, code for a proposed Help page (Listing 3.32), 116A–117A

Heroku

- commands, 39A–40A
- creating a new application at Heroku (Listing 1.7), 37A
- deployment, 37A–39A
- setup, 36A–37A

Home page

- adding follower stats to the Home page (Listing 12.22), 490A
- adding microposts creation to the Home Page (Listing 11.27), 442A
- with follow stats, 489A

- generated view for (Listing 3.7), 83A
 - Home view with HTML structure removed (Listing 3.29), 114A
 - with a link to the signup page (Listing 5.2), 163A
 - mockup with a form for creating microposts, 439A
 - mockup with a proto-feed, 447A
 - with a proto-feed, 451A
 - testing, 456A–457A
 - view for the Home page with an Embedded Ruby title (Listing 3.25), 110A–111A
 - view for the Home page with full HTML structure (Listing 3.21), 107A
 - href, 108A
 - HTML
 - for the form in Figure 8.3 (Listing 8.5), 289A
 - for the signin form produced by Listing 9.4, 331A
 - for signup form, 288A–292A
 - typical HTML file with a friendly greeting (Listing 3.3), 76A
 - for the user edit form, 371A
 - HTTP response codes, 89A
 - HTTP verbs, 80A–81A
 - hypertext reference, 108A
- I**
- IDEs, 9A
 - implicit return, 133A
 - index action, simplified user index action for the demo application (Listing 2.4), 56A
 - index page, 47A
 - indexes, 226A–227A
 - index.html file, 75A–78A
 - inheritance, 52A
 - additions to .autotest needed to run integration tests with Autotest on Ubuntu Linux (Listing 5.17), 180A
 - ApplicationController class with inheritance (Listing 2.16), 67A
 - classes, 145A–148A
 - hierarchies, 66A–68A
 - Micropost class with inheritance (Listing 2.13), 66A
 - MicropostsController class with inheritance (Listing 2.15), 67A
 - User class with inheritance (Listing 2.12), 66A
 - UserController class with inheritance (Listing 2.14), 67A
 - initialization hash, 204A–205A
 - inspect method, 142A
 - instance variables, 57A, 108A–112A
 - adding a feed instance variable to the home action (Listing 11.33), 448A
 - adding a micropost instance variable to the home action (Listing 11.30), 443A
 - adding an @microposts instance variable to the user show action, 430A
 - adding an (empty) @feed_items instance variable to the create action (Listing 11.37), 451A
 - integrated development environments.
 - See IDEs
 - integration alternatives, 314A–315A
 - integration tests, 178A–180A, 313A–321A
 - adding a view for the Help page (Listing 5.15), 180A
 - adding the help action to the Pages controller (Listing 5.14), 179A
 - additions to .autotest needed to run integration tests with Autotest on OS X (Listing 5.16), 180A
 - for friendly forwarding, 382A
 - a function to sign users in inside of integration tests (Listing 9.31), 364A
 - for the microposts on the home page (Listing 11.41), 456A–457A

- for routes (Listing 5.13), 179A
- for signing in and out (Listing 9.30), 362A–363A
- interpolation, 127A
- IRC clients, 12An8
- iTerm, 10A

J

- JavaScript, 49A
 - adding the default JavaScript libraries to the sample app (Listing 10.39), 405A
- JavaScript Embedded Ruby (JS-ERb) files, 505A, 506A
 - JavaScript Embedded Ruby to create a following relationship (Listing 12.37), 506A
- join method, 136A

K

- Kate, 10A
- Katz, Yehuda, 3A
- Kittrell, Ben, 10A–11A
- Komodo Edit, 11A

L

- lambda, 295A, 307A, 318A, 514A–515A
- layout files, 107A, 112A–115A
 - sample application site layout (Listing 3.28), 113A
 - sample application site layout (Listing 4.1), 120A
 - sample application site layout (Listing 4.3), 122A
 - site layout with added structure (Listing 5.1), 159A
- layout links, 177A
 - changing, 358A–361A
 - test for the links on the layout (Listing 5.33), 192A
 - to the user index, 388A

- length validations, 61A–63A, 217A–218A
 - constraining microposts to at most 140 characters with a length validation (Listing 2.9), 62A
- Linux, 10A
- lists, unordered, 163A
- literal constructor, 144A
- literal strings, 126A
- log files, ignoring, 26A
- logo helper
 - header partial with the logo helper from Listing 5.32 (Listing 5.31), 191A–192A
 - template for the logo helper (Listing 5.32), 192A
- logs
 - development log with filtered passwords (Listing 8.12), 304A
 - filtering passwords by default (Listing 8.13), 304A–305A
 - pre-Rails 3 development log with visible passwords (Listing 8.11), 304A

M

- Macintosh OS X, 10A
- MacVim, 10A, 25A
- magic columns, 198A, 205A
- map method, 138A–139A
- mapping, route and URL mapping for site links, 177A
- Merb, merger with Rails, 3A
- message expectations, 266A
- messaging, 521A
- methods, 82A, 129A–132A
 - chaining, 130A, 408A
 - defining, 132A–133A
- Micropost model, 411A
 - the basic model, 412A–414A
 - the initial Micropost spec (Listing 11.3), 414A

- a micropost belongs to a user
 - (Listing 11.6), 418A
- the Micropost migration (Listing 11.1), 412A
- a user has many microposts (Listing 11.7), 418A
- user/micropost associations, 414A–418A
- validations (Listing 11.14), 424A
- microposts
 - adding microposts to the sample data
 - (Listing 11.20), 433A
 - creating, 439A–444A
 - CSS for, 430A–431A
 - data models for, 44A
 - destroying, 452A–457A
 - ensuring that a user’s microposts are
 - destroyed along with the user, 422A
 - form partial for creating microposts
 - (Listing 11.28), 442A
 - manipulating, 434A–436A
 - ordering the microposts with `default_scope`
 - (Listing 11.10), 421A
 - a partial for showing a single micropost
 - (Listing 11.38), 452A–453A
 - proto-feed, 444A–452A
 - refinements, 419A–422A
 - sample microposts, 432A–434A
 - showing, 425A–434A
 - summary of user/micropost association
 - methods, 418A
 - testing that microposts are destroyed when
 - users are, 421A–422A
 - testing the order of a user’s microposts
 - (Listing 11.9), 420A
 - validations, 423A–424A
- Microposts controller, 60A–61A
 - create action (Listing 11.26), 441A
 - destroy action (Listing 11.40), 456A
 - in schematic form (Listing 2.8), 60A–61A
- Microposts resource, 58A, 66A–67A
 - access control, 436A–438A
 - has_many microposts, 63A–65A
 - inheritance hierarchies, 66A–68A
 - length validations, 61A–63A
 - Rails routes with a new rule for Microposts
 - resources (Listing 2.7), 60A
 - RESTful routes provided by, 60A
 - routes for the Microposts resource (Listing 11.21), 435A
 - tour, 58A–61A
- migration, 196A–200A
 - to add a boolean admin attribute to users
 - (Listing 10.35), 401A
 - migrating a database with Rake, 45A
 - password migration, 244A–246A
 - for the User model (to create a users table)
 - (Listing 6.2), 198A
- mockups, 157A–158A
- model-view-controller, 22A–23A
 - diagram of MVC in Rails, 55A
 - Users, 230A–232A
 - and Users resource, 49A–58A
- Mongrel, 20An12
- MVC. *See* model-view-controller
- N**
- name attribute, 290A
- named routes, 177A, 181A, 183A–185A
 - footer partial with links (5.22), 184A–185A
 - header partial with links (5.21), 184A
- namespaces, 390A–391A
- navigation. *See* site navigation
- nested hashes, 141A, 333A
- nil, 130A–131A
- O**
- objects, 129A–132A
- or equals assignment operator, 349A–350A

P

Pages controller

- with added about action (Listing 3.16), 100A–101A
- generated Pages controller spec (Listing 3.10), 88A
- generating, 78A–79A
- generating (Listing 3.4), 78A–79A
- inheritance hierarchy, 151A
- made by Listing 3.4 (Listing 3.6), 82A
- with per-page titles (Listing 3.24), 110A
- routes for the home and contact actions in the Pages controller (Listing 3.5), 79A
- spec with a base title (Listing 3.33), 117A–118A
- spec with a failing test for the About page (Listing 3.15), 98A
- spec with title tests (Listing 3.20), 105A–106A

PagesController, 82A

paginating users, 392A–397A

- paginating the users in the index action (Listing 10.28), 393A
- testing pagination, 394A–397A

palindrome? method, 148A–149A

Paperclip, 271An21

partial refactoring, 398A–399A

partials, 171A–177A

- adding the CSS for the site footer (Listing 5.12), 175A
- for displaying form submission error messages, 300A
- for the site footer (Listing 5.10), 174A
- for the site header (Listing 5.9), 174A
- site layout with a footer partial (Listing 5.11), 175A
- site layout with partials for the stylesheets and header (Listing 5.7), 172A
- for stylesheet includes (Listing 5.8), 173A
- updating the error-messages partial, 369A

passwords

- Active Record callback, 247A–250A
- a before_save callback to create the encrypted_password attribute (Listing 7.6), 248A
- has_password? method for users (Listing 7.7), 251A
- has_password? method with secure encryption (Listing 7.10), 256A
- implementing has_password?, 254A–258A
- insecure, 239A
- migration, 244A–246A
- migration to add a salt column to the users table (Listing 7.9), 255A
- migration to add an encrypted_password column to the users table (Listing 7.4), 246A
- rainbow attack, 254A
- reminders, 521A
- secure, 250A
- secure password test, 251A–252A
- secure password theory, 252A–254A
- testing for the existence of an encrypted_password attribute (Listing 7.3), 245A
- testing that the encrypted_password attribute is nonempty (Listing 7.5), 247A
- tests for the has_password? method (Listing 7.8), 252A
- validations, 240A–244A
- See also* authenticate method
- PeepCode, 523A
- pending spec, 214A–215A
- percent-parentheses construction, 516A
- persistence, 196A
- Phusion Passenger, 36A
- pluralize text helper, 301A
- PostgreSQL, 196An5

pound sign. *See* hash symbol
presence validations, 210A–217A
Preston-Werner, Tom, 270An19
private keyword, 249A
profile images, 268A–275A
profile links, adding, 360A–361A
profile pages. *See* user profile page
protected keyword, 249An4
protecting pages, 376A–384A
 mockup of a protected page, 377A
public/index.html file, 75A–76A
pushing data, 68A–69A
puts method, 127A–128A

R

Rails

 deploying, 35A–40A
 installing, 15A
 overview, 3A–4A
 The Rails 3 Way (Fernandez), 4A
 The Rails 3 Way (Fernandez), 6A, 523A
rails command, 15A–16A
Rails console, 125A
Rails Machine, 36A
Rails root. *See* root
Rails routes, 181A–183A
 adding a mapping for the root route (Listing 5.20), 182A–183A
 commented-out hint for defining the root route (Listing 5.19), 182A
 for static pages (Listing 5.18), 181A
rails script, running the rails script to generate a new application (Listing 1.1), 16A
rails server, 20A–22A
Railscasts, 522A
rainbow attack, 254A
Rake, 45A, 46A
 a Rake task for populating the database with sample users (Listing 10.25), 390A

 ranges, 137A
README file
 improved README file for the sample app (Listing 3.2), 73A
 new README file, README.markdown (Listing 1.6), 33A
 updating, 73A
Red, Green, Refactor, 86A–91A
 Green, 100A–102A
 Red, 97A–100A
 Refactor, 102A–103A
refactoring, 398A–399A
 a compact refactoring of Listing 12.36 (Listing 12.46), 524A
 refactored following and followers actions (Listing 12.47), 524A–525A
regex, 220A
regular expressions, 220A
Relationship data model, 463A–469A
 adding the belongs_to associations to the Relationship model (Listing 12.7), 473A
 validations, 473A–474A
relationships attribute, 470A–471A
Relationships controller (Listing 12.32), 501A
reload method, 375A
remember tokens, 341A, 342A–344A
render, 173A
replies, 520A–521A
repositories, first-time repository setup, 25A–26A
REpresentational State Transfer. *See* REST
request specs, 178A
 See also integration tests
resources, advanced Rails resources, 7A
resources for Rails, 522A–523A
REST, 54A–56A
 displaying user show page following REST architecture, 232A–233A

REST API, 522A
 reverse relationships, 480A–482A
 root, 8A
 RSpec, 71A–72A, 84A–85A
 adding the `-drb` option to the `.rspec` file
 (Listing 3.14), 96A
 count method, 295A
 integration tests, 313A–321A
 request specs, 178A
 RSS feed, 521A
 Rubular, 220A–222A
 Ruby
 gems, 13A, 14A
 gemsets, 13A–14A
 installing, 12A–14A
 learning Ruby before learning Rails, 4A–5A
 Ruby JavaScript (RJS), to destroy a
 following relationship
 (Listing 12.38), 506A
 Ruby on Rails. *See* Rails
 Ruby Version Manager (RVM), 12A
The Ruby Way (Fulton), 6A, 7A, 523A
 RubyGems, installing, 14A–15A

S

Safari, 11A
 salt, 254A, 255A
 sandbox, 203A
 save!, 470A
 scaffolding, 2A–3A
 scaling Rails, 7A, 523A
 Schoeneman, Fred, 86A
 scopes, 514A–515A
 screencasts, 522A
 search, 522A
 Seguin, Wayne E., 12A
 self, 260A–261A
 sessions, 341A
 defined, 325A–326A
 destroying, 354A–356A
 Sessions controller, 326A–328A
 adding a resource to get the standard
 RESTful actions for sessions (Listing
 9.2), 327A
 completed Sessions controller create action
 (not yet working) (Listing 9.9),
 338A–339A
 tests for the new session action and view
 (Listing 9.1), 327A
 SHA2, 253A
 short-circuit evaluation, 350A
 Shoulda, 85An7
 showing microposts, 425A–434A
 sidebar, 276A–279A
 partial for the user info sidebar (Listing
 11.29), 443A
 signed_in? helper method, 353A
 signed-in users, requiring, 376A–379A
 signin form, 328A–332A
 code for a failed signin attempt
 (Listing 9.8), 336A–337A
 code for the signin form (Listing 9.4), 330A
 failure, 332A–337A
 HTML for the signing form produced by
 Listing 9.4 (Listing 9.5), 331A
 mockup, 329A
 pending tests for user signin (Listing 9.10),
 340A
 remembering user signin status forever,
 340A–344A
 reviewing form submission, 333A–335A
 success, 338A–353A
 tests for a failed signin attempt
 (Listing 9.7), 335A–336A
 signin page, adding the title for the signing
 page (Listing 9.3), 328A
 signin upon signup, 356A–357A
 signing out, 354A
 destroying a session (user signout) (Listing
 9.21), 355A

- destroying sessions, 354A–356A
- the `sign_out` method in the Sessions helper module (Listing 9.22), 356A
- a test for destroying a session (Listing 9.20), 355A
- a `test_sign_in` function to simulate user sign in inside tests (Listing 9.19), 354A
- signin/signout integration tests, 362A–363A
- signin/signout links
 - adding a profile link (Listing 9.29), 360A–361A
 - changing, 358A–361A
 - changing the layout links for signed-in users (Listing 9.26), 359A
 - a helper for the site logo (Listing 9.27), 360A
 - a test for a profile link (Listing 9.28), 360A
 - tests for the signin/signout links on the site layout (Listing 9.25), 358A
- signup confirmation, 521A
- signup form
 - adding an `@user` variable to the new action (Listing 8.3), 287A
 - code to display error messages on the signup form (Listing 8.8), 299A
 - a create action that can handle signup failure (but not success) (Listing 8.7), 296A
 - CSS for styling error messages (Listing 8.10), 302A
 - error explanation div from the page in Figure 8.11 (Listing 8.19), 317A
 - error messages, 299A–303A
 - failure, 292A–304A
 - filtering parameter logging, 303A–305A
 - finished form, 308A
 - the first signup, 312A–313A
 - form HTML, 288A–292A
 - a form to sign up new users (Listing 8.2), 286A
 - overview, 283A–285A
 - a partial for displaying form submission error messages (Listing 8.9), 300A
 - pluralize text helper, 301A
 - success, 305A–313A
 - a template for testing for each field on the signup form (Listing 8.23), 322A–323A
 - testing failure, 292A–295A
 - testing signup failure (Listing 8.20), 317A
 - testing signup failure with a lambda (Listing 8.21), 318A
 - testing signup success (Listing 8.22), 319A
 - testing success, 305A–308A
 - the user create action with a save and a redirect (Listing 8.15), 308A
 - using `form_for`, 286A–288A
 - a wafer-thin amount of CSS for the signup form (Listing 8.4), 288A
 - a working form, 295A–298A
- signup page
 - action for the new user signup page (Listing 5.25), 187A
 - linking the button to the signup page (Listing 5.30), 190A
 - route for the signup page (Listing 5.29), 189A
 - setting the custom title for the new user page (Listing 5.27), 188A
 - signin upon signup, 356A–357A
 - signing in the user upon signup (Listing 9.24), 357A
 - test for the signup page title (Listing 5.26), 188A
 - testing that newly signed-up users are also signed in (Listing 9.23), 356A–357A
 - testing the signup page (Listing 5.24), 187A
 - the tests for the new users page (Listing 8.1), 284A–285A
 - Users controller, 186A–188A
- signup URL, 188A–190A
- site navigation, 159A–164A
- skeleton for a shuffle method attached to the String class (Listing 4.10), 155A

- skeleton for a string shuffle function (Listing 4.9), 155A
 - slightly dynamic pages, 103A
 - eliminating duplication with layouts, 112A–115A
 - instance variables and Embedded Ruby, 108A–112A
 - passing title tests, 106A–108A
 - testing a title change, 103A–106A
 - spike, 87A
 - split method, 134A–135A
 - Spork, 91A–97A
 - adding environment loading to the Spork.prefork block (Listing 3.12), 93A–94A
 - last part of the hack needed to get Spork to run with Rails 3 (Listing 3.13), 95A
 - SQL injection, 448A
 - SQLite Database Browser, 199A, 200A
 - staging area, 27A
 - static pages, 74A
 - with Rails, 78A–83A
 - truly static pages, 75A–78A
 - See also* slightly dynamic pages
 - stats, 484A–493A
 - a partial for displaying follower stats (Listing 12.21), 487A–488A
 - status command, 27A
 - status feed, 444A–452A, 507A
 - adding a status feed to the Home page (Listing 11.36), 450A
 - adding the completed feed to the User model (Listing 12.41), 510A
 - the final implementation of `from_users_followed_by` (Listing 12.44), 517A
 - the final tests for the status feed (Listing 12.40), 509A–510A
 - a first cut at the `from_users_followed_by` method (Listing 12.42), 513A
 - a first feed implementation, 511A–513A
 - home action with a paginated feed (Listing 12.45), 519A
 - improving `from_users_followed_by` (Listing 12.43), 515A
 - mockup of a user’s Home page with a status feed, 507A
 - mockup of the Home page with a proto-feed, 447A
 - motivation and strategy, 508A–510A
 - a partial for a single feed item (Listing 11.35), 449A–450A
 - preliminary implementation for the micropost status feed (Listing 11.32), 447A
 - scopes, subselects, and a lambda, 513A–518A
 - status feed partial (Listing 11.34), 449A
 - tests for `Micropost.from_users_followed_by` (Listing 12.39), 508A–509A
 - string literals, 126A
 - strings, 126A–127A
 - double-quoted, 128A–129A
 - printing, 127A–128A
 - single-quoted, 128A–129A
 - stub About page (Listing 3.18), 101A
 - stubbing, 266A
 - stylesheets. *See* CSS
 - Sublime Text editor, 11A
 - subselects, 517A
 - sudo, 14A–15A
 - superclass method, 145A
 - symbols, 140A–142A
 - system setups, 22A, 24A
- T**
- TDD. *See* test-driven development (TDD)
 - ternary operator, 352A–353A
 - test-driven development (TDD), 84A
 - Green, 100A–102A
 - Red, 97A–100A

- Red, Green, Refactor, 86A–91A
 - Refactor, 102A–103A
 - Spork, 91A–97A
 - testing tools, 84A–86A
 - tests, 84A
 - access control tests for the Microposts controller (Listing 11.22), 437A
 - for an admin attribute (Listing 10.34), 399A–400A
 - for destroying users (Listing 10.40), 406A–407A
 - for failed user signup (Listing 8.6), 293A–294A
 - integration tests, 178A–180A, 313A–321A
 - for the micropost model validations (Listing 11.13), 423A
 - for the Microposts controller create action (Listing 11.25), 440A–441A
 - for the Microposts controller destroy action (Listing 11.39), 454A–455A
 - for the micropost’s user association (Listing 11.4), 415A
 - for pagination (Listing 10.30), 396A–397A
 - passing title tests, 106A–108A
 - for the (proto-)status feed (Listing 11.31), 445A–446A
 - for the Relationships controller actions (Listing 12.31), 499A–500A
 - for reverse relationships (Listing 12.16), 480A–481A
 - for showing microposts on the user show page (Listing 11.15), 426A
 - signup form testing failure, 292A–295A
 - signup form testing success, 305A–308A
 - for signup success (Listing 8.14), 306A–307A
 - simple integration test for user signup link (Listing 5.28), 189A
 - for some following utility methods (Listing 12.12), 476A–477A
 - testing a title change, 103A–106A
 - testing for the user.relationships attribute (Listing 12.4), 470A–471A
 - testing pagination, 394A–397A
 - testing relationship creation with save! (Listing 12.3), 470A
 - testing the following/follower statistics on the Home page (Listing 12.20), 486A–487A
 - testing the signup page, 187A–188A
 - testing the user/relationships belongs_to association (Listing 12.6), 472A–473A
 - for the user’s microposts attribute (Listing 11.5), 417A
 - whether to use tests from the start, 5A
See also Autotest; RSpec
 - text editors, 9A–11A
 - TextMate, 10A, 25A
 - Thomas, Dave, 249An4
 - time helpers, 343A
 - timestamps, 198A, 205A
 - title change
 - passing title tests, 106A–108A
 - testing, 103A–106A
 - title helper, 119A–122A, 133A–134A
 - defining a title helper (Listing 4.2), 121A
 - title test (Listing 3.19), 104A
 - toggle method, 401A–402A
- ## U
- unfollow form, using Ajax (Listing 12.34), 503A
 - unfollow/follow buttons, 498A–502A
 - with Ajax, 502A–506A
 - unfollowing
 - test for unfollowing a user (Listing 12.14), 478A–479A
 - unfollowing a user by destroying a user relationship (Listing 12.15), 479A
See also following
 - uniqueness validation, 222A–226A
 - Unix style, 8A
 - unordered list tag, 163A

- update action, 373A–376A
- updating users, 365A–376A
- URLs, defined, 2An3
- URLs, defined, 2An3
- user edit form, 366A–373A
 - adding a Settings link (Listing 10.6), 370A
 - enabling edits, 373A–376A
 - HTML for the edit form (Listing 10.7), 371A
 - mockup, 366A
 - a partial for the new and edit form fields (Listing 10.43), 410A
 - tests for the user edit action (Listing 10.1), 367A
 - tests for the user update action (Listing 10.8), 374A–375A
 - updating the error-messages partial from Listing 8.9 to work with other objects (Listing 10.4), 369A
 - updating the rendering of user signup errors (Listing 10.5), 370A
 - the user edit action (Listing 10.2), 368A
 - the user edit view (Listing 10.3), 368A–369A
 - the user update action (Listing 10.9), 375A
- user index, 385A–389A
 - CSS for the user index (Listing 10.22), 388A
 - the first refactoring attempt at the index view (Listing 10.31), 398A
 - the fully refactored user index (Listing 10.33), 399A
 - a layout link to the user index (Listing 10.23), 388A
 - mockup, 385A, 400A
 - with pagination (Listing 10.27), 392A–393A
 - partial refactoring, 398A–399A
 - a partial to render a single user (Listing 10.32), 398A
 - tests for the user index page (Listing 10.19), 385A–386A
 - the user index action (Listing 10.20), 387A
 - the user index view (Listing 10.21), 387A
 - view for the user index (Listing 2.6), 57A
- user info sidebar, 276A–279A, 443A
- user model, 194A–196A
- User model
 - accessible attributes, 202A–203A
 - with an added (encrypted) password attribute, 246A
 - with an added salt, 256A
 - adding the annotate-models gem to the Gemfile (Listing 6.4), 201A
 - annotated User model (Listing 6.5), 202A
 - annotating the model file, 201A–202A
 - brand new User model (Listing 6.3), 201A
 - generating a User model (Listing 6.1), 197A
 - making the name and email attributes accessible (Listing 6.6), 203A
 - migration for the User model (to create a users table) (Listing 6.2), 198A
 - model file, 201A–203A
- User model fro the demo application (Listing 2.5), 57A
- user objects
 - creating, 203A–207A
 - finding, 207A–208A
 - updating, 208A–209A
- user profile page
 - with microposts, 434A
 - mockup, 425A, 462A
 - mockup with a “Settings” link, 371A
- user show page
 - adding a name and gravatar, 268A–275A
 - adding a sidebar to the user show view (Listing 7.25), 276A
 - adding an @microposts instance variable to the user show action (Listing 11.18), 430A
 - adding microposts to the user show page (Listing 11.16), 427A

- augmenting, 426A–432A
 - CSS for styling the user show page including
 - the sidebar (Listing 7.26), 278A–279A
 - a partial for showing a single micropost (Listing 11.17), 429A
 - tests for the user show page (Listing 7.18), 268A–269A
 - a title for the user show page (Listing 7.19), 269A
 - the user show view with name and Gravatar (Listing 7.22), 271A
 - the user show view with the user’s name (Listing 7.20), 270A
 - a user sidebar, 276A–279A
 - user views, 262A
 - user.followers method, 479A–482A
 - user/relationship associations, 470A–473A
 - implementing the user/relationships
 - has_many association (Listing 12.5), 472A
 - users
 - administrative, 399A–404A
 - the current_user? method (Listing 10.14), 381A
 - destroying, 399A–408A
 - the new user view with partial (Listing 10.44), 410A
 - paginating, 392A–397A
 - requiring signed-in users, 376A–379A
 - requiring the right user, 378A–382A
 - sample users, 389A–391A
 - showing, 384A–399A
 - stub view for showing user information (Listing 6.24), 231A
 - summary of user/micropost association methods, 418A
 - updating, 365A–376A
 - Users controller, 52A–53A
 - adding following and followers actions to the Users controller (Listing 12.19), 485A
 - current Users controller spec (Listing 7.13), 262A
 - generating a Users controller with a new action (Listing 5.23), 186A
 - in schematic form (Listing 2.3), 53A
 - with a show action (Listing 6.25), 232A
 - signup page, 186A–188A
 - testing the user show page with factories, 263A–268A
 - user show action from Listing 6.25 (Listing 7.14), 263A
 - Users resource, 232A–236A
 - adding a Users resource to the routes file (Listing 6.26), 234A
 - correspondence between pages and URLs, 47A
 - and MVC, 49A–58A
 - overview, 44A–46A
 - Rails routes with a rule for the Users resource (Listing 2.2), 52A
 - RESTful routes provided by, 55A, 235A
 - user tour, 46A–49A
 - weaknesses, 58A
- ## V
- validations, 61A–63A
 - adding a length validation for the name attribute (Listing 6.15), 218A
 - adding the Relationship model validations (Listing 12.9), 474A
 - commenting out a validation to ensure a failing test (Listing 6.8), 212A
 - failing test for the validation of the name attribute (Listing 6.11), 215A
 - format, 218A–222A
 - initial user spec (Listing 6.10), 213A
 - length, 61A–63A, 217A–218A
 - microposts, 423A–424A
 - migration for enforcing email uniqueness (Listing 6.22), 226A

- overview, 210A
- password, 240A–244A
- for the password attribute (Listing 7.2), 243A
- practically blank default User spec (Listing 6.9), 212A
- presence, 210A–217A
- Relationship data model, 473A–474A
- test for the name length validation (Listing 6.14), 217A–218A
- test for the presence of the email attribute (Listing 6.12), 216A–217A
- test for the rejection of duplicate email addresses, insensitive to case (Listing 6.20), 223A–224A
- test for the rejection of duplicate email addresses (Listing 6.18), 222A–223A
- testing the Relationship model validations (Listing 12.8), 474A
- tests for email format validation (Listing 6.16), 219A
- tests for password validations (Listing 7.1), 241A–242A
- uniqueness, 222A–226A
- validating the email format with a regular expression (Listing 6.17), 220A
- validating the presence of a name attribute (Listing 6.7), 211A

- validating the presence of the name and email attributes (Listing 6.13), 217A
- validating the uniqueness of email addresses, ignoring case (Listing 6.21), 224A
- validating the uniqueness of email addresses (Listing 6.19), 223A
- Vim, 11A
- Vim for Windows with Console, 10A
- virtual attributes, 242A–243A

W

- Webrat, 72A, 315An9
- WEBrick, 20An12
- The Well-Founded Rubyist* (Black), 7A, 261A, 523A
- will_paginate method, 392A–394A
- Windows, 10A
- wireframes, 157A
- wrapping words, a helper to wrap long words (Listing 11.42), 459A

Y

- YAML, 236A

Z

- zero-offset, 135A

THE RAILS™ 3 WAY

This page intentionally left blank

Praise for the Previous Edition

This encyclopedic book is not only a definitive Rails reference, but an indispensable guide to Software-as-a-Service coding techniques for serious craftspersons. I keep a copy in the lab, a copy at home, and a copy on each of my three e-book readers, and it's on the short list of essential resources for my undergraduate software engineering course.

—Armando Fox, adjunct associate professor, University of California, Berkeley

Everyone interested in Rails, at some point, has to follow *The Rails Way*.

—Fabio Cevasco, senior technical writer, Siemens AG, and blogger at H3RALD.com

I can positively say that it's the single best Rails book ever published to date. By a long shot.

—Antonio Cangiano, software engineer and technical evangelist at IBM

This book is a great crash course in Ruby on Rails! It doesn't just document the features of Rails, it filters everything through the lens of an experienced Rails developer—so you come out a pro on the other side.

—Dirk Elmendorf, co-founder of Rackspace, and Rails developer since 2005

The key to *The Rails Way* is in the title. It literally covers the “way” to do almost everything with Rails. Writing a truly exhaustive reference to the most popular Web application framework used by thousands of developers is no mean feat. A thankful

community of developers that has struggled to rely on scant documentation will embrace *The Rails Way* with open arms. A tour de force!

—Peter Cooper, editor, *Ruby Inside*

In the past year, dozens of Rails books have been rushed to publication. A handful are good. Most regurgitate rudimentary information easily found on the Web. Only this book provides both the broad and deep technicalities of Rails. Nascent and expert developers, I recommend you follow *The Rails Way*.

—Martin Streicher, chief technology officer, McLatchy Interactive; former editor-in-chief of *Linux Magazine*

Hal Fulton's *The Ruby Way* has always been by my side as a reference while programming Ruby. Many times I had wished there was a book that had the same depth and attention to detail, only focused on the Rails framework. That book is now here and hasn't left my desk for the past month.

—Nate Klaiber, Ruby programmer

As noted in my contribution to the Afterword: "What Is the Rails Way (To You)?," I knew soon after becoming involved with Rails that I had found something great. Now, with Obie's book, I have been able to step into Ruby on Rails development coming from .NET and be productive right away. The applications I have created I believe to be a much better quality due to the techniques I learned using Obie's knowledge.


—Robert Bazinet, InfoQ.com, .NET and Ruby community editor, and founding member of the Hartford, CT, Ruby Brigade

Extremely well written; it's a resource that every Rails programmer should have. Yes, it's that good.

—Reuven Lerner, *Linux Journal* columnist

Addison Wesley Professional Ruby Series

Obie Fernandez, Series Editor



Visit informit.com/ruby for a complete list of available products.

The **Addison-Wesley Professional Ruby Series** provides readers with practical, people oriented, and in depth information about applying the Ruby platform to create dynamic technology solutions. The series is based on the premise that the need for expert reference books, written by experienced practitioners, will never be satisfied solely by blogs and the Internet.

PEARSON

◆ Addison-Wesley

Cisco Press

EXAM/CRAM

IBM Press

que

PRENTICE HALL

SAMS

Safari

This page intentionally left blank

THE RAILS™ 3WAY

Obie Fernandez

◆ Addison-Wesley

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States please contact:
International Sales
international@pearson.com

Visit us on the Web: informit.com/aw

Library of Congress Cataloging-in-Publication Data
Fernandez, Obie.

The rails 3 way / Obie Fernandez.
p. cm.

Rev. ed. of: *The Rails way* / Obie Fernandez. 2008.

Includes index.

ISBN 0-321-60166-1 (pbk. : alk. paper)

1. Ruby on rails (Electronic resource) 2. Object-oriented programming (Computer science)

3. Ruby (Computer program language) 4. Web site development. 5. Application

software--Development. I. Fernandez, Obie. Rails way. II. Title.

QA76.64.F47 2010

005.1'17--dc22

2010038744

Copyright © 2011 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, write to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax: (617) 671-3447

Parts of this book contain material excerpted from the Ruby and Rails source code and API documentation, Copyright © 2004–2011 by David Heinemeier Hansson under the MIT license. Chapter 18 contains material excerpted from the RSpec source code and API documentation, Copyright © 2005–2011 The RSpec Development Team.

The MIT License reads: Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED “AS IS,” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES, OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT, OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OF OR OTHER DEALINGS IN THE SOFTWARE.

ISBN-13: 978-0-321-60166-7

ISBN-10: 0-321-60166-1

Text printed in the United States on recycled paper at Edwards Brothers in Ann Arbor, Michigan.
Second printing, April 2011

Editor-in-Chief

Mark Taub

Executive Acquisitions Editor

Debra Williams Cauley

Managing Editor

John Fuller

Project Editor

Elizabeth Ryan

Copy Editor

Carol Loomis

Indexer

Valerie Haynes Perry

Proofreader

Erica Orloff

Publishing Coordinator

Kim Boedigheimer

Cover Designer

Chuti Prasertsith

Compositor

Glyph International

To Dad, thanks for teaching me ambition.

This page intentionally left blank

Foreword

Rails is more than a programming framework for creating web applications. It's also a framework for thinking about web applications. It ships not as a blank slate equally tolerant of every kind of expression. On the contrary, it trades that flexibility for the convenience of “what most people need most of the time to do most things.” It's a designer straightjacket that sets you free from focusing on the things that just don't matter and focuses your attention on the stuff that does.

To be able to accept that trade, you need to understand not just how to do something in Rails, but also why it's done like that. Only by understanding the why will you be able to consistently work with the framework instead of against it. It doesn't mean that you'll always have to agree with a certain choice, but you will need to agree to the overachieving principle of conventions. You have to learn to relax and let go of your attachment to personal idiosyncrasies when the productivity rewards are right.

This book can help you do just that. Not only does it serve as a guide in your exploration of the features in Rails, it also gives you a window into the mind and soul of Rails. Why we've chosen to do things the way we do them, why we frown on certain widespread approaches. It even goes so far as to include the discussions and stories of how we got there—straight from the community participants that helped shape them.

Learning how to do Hello World in Rails has always been easy to do on your own, but getting to know and appreciate the gestalt of Rails, less so. I applaud Obie for trying to help you on this journey. Enjoy it.

— **David Heinemeier Hansson**
Creator of Ruby on Rails

This page intentionally left blank

Foreword

From the beginning, the Rails framework turned web development on its head with the insight that the vast majority of time spent on projects amounted to meaningless sit-ups. Instead of having the time to think through your domain-specific code, you'd spend the first few weeks of a project deciding meaningless details. By making decisions for you, Rails frees you to kick off your project with a bang, getting a working prototype out the door quickly. This makes it possible to build an application with some meat on its bones in a few weekends, making Rails the web framework of choice for people with a great idea and a full-time job.

Rails makes some simple decisions for you, like what to name your controller actions and how to organize your directories. It also gets pretty aggressive, and sets development-friendly defaults for the database and caching layer you'll use, making it easy to change to more production-friendly options once you're ready to deploy.

By getting so aggressive, Rails makes it easy to put at least a few real users in front of your application within days, enabling you to start gathering the requirements from your users immediately, rather than spending months architecting a perfect solution, only to learn that your users use the application differently than you expected.

The Rails team built the Rails project itself according to very similar goals. Don't try to overthink the needs of your users. Get something out there that works, and improve it based on actual usage patterns. By all accounts, this strategy has been a smashing success, and with the blessing of the Rails core team, the Rails community leveraged the dynamism of Ruby to fill in the gaps in plugins. Without taking a close look at Rails, you might think that Rails' rapid prototyping powers are limited to the 15-minute blog demo, but that you'd fall off a cliff when writing a real app. This has never been true. In fact, in Rails 2.1, 2.2 and 2.3, the Rails team looked closely at common usage patterns

reflected in very popular plugins, adding features that would further reduce the number of sit-ups needed to start real-life applications.

By the release of Rails 2.3, the Rails ecosystem had thousands of plugins, and applications like Twitter started to push the boundaries of the Rails defaults. Increasingly, you might build your next Rails application using a non-relational database or deploy it inside a Java infrastructure using JRuby. It was time to take the tight integration of the Rails stack to the next level.

Over the course of 20 months, starting in January 2008, we looked at a wide range of plugins, spoke with the architects of some of the most popular Rails applications, and changed the way the Rails internals thought about its defaults.

Rather than start from scratch, trying to build a generic data layer for Rails, we took on the challenge of making it easy to give any ORM the same tight level of integration with the rest of the framework as Active Record. We accepted no compromises, taking the time to write the tight Active Record integration using the same APIs that we now expose for other ORMs. This covers the obvious, such as making it possible to generate a scaffold using DataMapper or Mongoid. It also covers the less obvious, such as giving alternative ORMs the same ability to include the amount of time spent in the model layer in the controller's log output.

We brought this philosophy to every area of Rails 3: flexibility without compromise. By looking at the ways that an estimated million developers use Rails, we could hone in on the needs of real developers and plugin authors, significantly improving the overall architecture of Rails based on real user feedback.

Because the Rails 3 internals are such a departure from what's come before, developers building long-lived applications and plugin developers need a resource that comprehensively covers the philosophy of the new version of the framework. *The Rails™ 3 Way* is a comprehensive resource that digs into the new features in Rails 3 and perhaps more importantly, the rationale behind them.

— **Yehuda Katz**
Rails Core

Introduction

As I write this new introduction in the spring of 2010, the official release of Rails 3.0 is looming, and what a big change it represents. The “Merb-ification” of Rails is almost complete! The new Rails is quite different from its predecessors in that its underlying architecture is more modular and elegant while increasing sheer performance significantly. The changes to Active Record are dramatic, with Arel’s query method chaining replacing hashed `find` parameters that we were all used to.

There is a lot to love about Rails 3, and I do think that eventually most of the community will make the change. In most cases, I have not bothered to cover 2.x ways of doing things in Rails if they are significantly different from the Rails 3 way—hence the title change. I felt that naming the book “The Rails Way (Second Edition)” would be accurate, but possibly misleading. This new edition is a fully new book for a fully new framework. Practically every line of the book has been painstakingly revised and edited, with some fairly large chunks of the original book not making the new cut. It’s taken well over a year, including six months of working every night to get this book done!

Even though Rails 3 is less opinionated than early versions, in that it allows for easy reconfiguration of Rails assumptions, this book is more opinionated than ever. The vast majority of Rails developers use `RSpec`, and I believe that is primarily because it is a superior choice to `Test::Unit`. Therefore, this book does not cover `Test::Unit`. I firmly believe that `Ham1` is vastly, profoundly, better than `ERb` for view templating, so the book uses `Ham1` exclusively.

0.1 About This Book

This book is not a tutorial or basic introduction to Ruby or Rails. It is meant as a day-to-day reference for the full-time Rails developer. The more confident reader might be able to get started in Rails using just this book, extensive online resources, and his or her wits, but there are other publications that are more introductory in nature and might be a wee bit more appropriate for beginners.

Every contributor to this book works with Rails on a full-time basis. We do not spend our days writing books or training other people, although that is certainly something that we enjoy doing on the side.

This book was originally conceived for myself, because I hate having to use online documentation, especially API docs, which need to be consulted over and over again. Since the API documentation is liberally licensed (just like the rest of Rails), there are a few sections of the book that reproduce parts of the API documentation. In practically all cases, the API documentation has been expanded and/or corrected, supplemented with additional examples and commentary drawn from practical experience.

Hopefully you are like me—I really like books that I can keep next to my keyboard, scribble notes in, and fill with bookmarks and dog-ears. When I'm coding, I want to be able to quickly refer to both API documentation, in-depth explanations, and relevant examples.

0.1.1 Book Structure

I attempted to give the material a natural structure while meeting the goal of being the best-possible Rails reference book. To that end, careful attention has been given to presenting holistic explanations of each subsystem of Rails, including detailed API information where appropriate. Every chapter is slightly different in scope, and I suspect that Rails is now too big a topic to cover the whole thing in depth in just one book.

Believe me, it has not been easy coming up with a structure that makes perfect sense for everyone. Particularly, I have noted surprise in some readers when they notice that Active Record is not covered first. Rails is foremost a web framework and, at least to me, the controller and routing implementation is the most unique, powerful, and effective feature, with Active Record following a close second.

0.1.2 Sample Code and Listings

The domains chosen for the code samples should be familiar to almost all professional developers. They include time and expense tracking, auctions, regional data management, and blogging applications. I don't spend pages explaining the subtler nuances of the

business logic for the samples or justify design decisions that don't have a direct relationship to the topic at hand. Following in the footsteps of my series colleague Hal Fulton and *The Ruby Way*, most of the snippets are not full code listings—only the relevant code is shown. Ellipses (...) denote parts of the code that have been eliminated for clarity.

Whenever a code listing is large and significant, and I suspect that you might want to use it verbatim in your own code, I supply a listing heading. There are not too many of those. The whole set of code listings will not add up to a complete working system, nor are there 30 pages of sample application code in an appendix. The code listings should serve as inspiration for your production-ready work, but keep in mind that they often lack touches necessary in real-world work. For example, examples of controller code are often missing pagination and access control logic, because it would detract from the point being expressed.

Some of the source code for my examples can be found at <http://github.com/obie/tr3w'time'and'expenses>. Note that it is not a working nor complete application. It just made sense at times to keep the code in the context of an application and hopefully you might draw some inspiration from browsing it.

0.1.3 Concerning Third-Party RubyGems and Plugins

Whenever you find yourself writing code that feels like plumbing, by which I mean completely unrelated to the business domain of your application, you're probably doing too much work. I hope that you have this book at your side when you encounter that feeling. There is almost always some new part of the Rails API or a third-party RubyGem for doing exactly what you are trying to do.

As a matter of fact, part of what sets this book apart is that I never hesitate in calling out the availability of third-party code, and I even document the RubyGems and plugins that I feel are most crucial for effective Rails work. In cases where third-party code is better than the built-in Rails functionality, we don't cover the built-in Rails functionality (pagination is a good example).

An average developer might see his or her productivity double with Rails, but I've seen serious Rails developers achieve gains that are much, much higher. That's because we follow the Don't Repeat Yourself (DRY) principle religiously, of which Don't Reinvent The Wheel (DRTW) is a close corollary. Reimplementing something when an existing implementation is good enough is an unnecessary waste of time that nevertheless can be very tempting, since it's such a joy to program in Ruby.

Ruby on Rails is actually a vast ecosystem of core code, official plugins, and third-party plugins. That ecosystem has been exploding rapidly and provides all the raw

technology you need to build even the most complicated enterprise-class web applications. My goal is to equip you with enough knowledge that you'll be able to avoid continuously reinventing the wheel.

0.2 Recommended Reading and Resources

Readers may find it useful to read this book while referring to some of the excellent reference titles listed in this section.

Most Ruby programmers always have their copy of the “Pickaxe” book nearby, *Programming Ruby* (ISBN: 0-9745140-5-5), because it is a good language reference. Readers interested in really understanding all of the nuances of Ruby programming should acquire *The Ruby Way, Second Edition* (ISBN: 0-6723288-4-4).

I highly recommend Peepcode Screencasts, in-depth video presentations on a variety of Rails subjects by the inimitable Geoffrey Grosenbach, available at <http://peepcode.com>

Ryan Bates does an excellent job explaining nuances of Rails development in his long-running series of free webcasts available at <http://railscasts.com/>

Last, but not least, this book's companion website at <http://tr3w.com> is the first place to look for reporting issues and finding additional resources, as they become available.

Regarding David Heinemeier Hansson, a.k.a. DHH

I had the pleasure of establishing a friendship with David Heinemeier Hansson, creator of Rails, in early 2005, before Rails hit the mainstream and he became an International Web 2.0 Superstar. My friendship with David is a big factor in why I'm writing this book today. David's opinions and public statements shape the Rails world, which means he gets quoted a lot when we discuss the nature of Rails and how to use it effectively.

David has told me on a couple of occasions that he hates the “DHH” moniker that people tend to use instead of his long and difficult-to-spell full name. For that reason, in this book I try to always refer to him as “David” instead of the ever-tempting “DHH.” When you encounter references to “David” without further qualification, I'm referring to the one-and-only David Heinemeier Hansson.

There are a number of notable people from the Rails world that are also referred to on a first-name basis in this book. Those include:

- *Yehuda* Katz
 - *Jamis* Buck
 - *Xavier* Noria
-

0.3 Goals

As already stated, I hope to make this your primary working reference for Ruby on Rails. I don't really expect too many people to read it through end to end unless they're expanding their basic knowledge of the Rails framework. Whatever the case may be, over time I hope this book gives you as an application developer/programmer greater confidence in making design and implementation decisions while working on your day-to-day tasks. After spending time with this book, your understanding of the fundamental concepts of Rails coupled with hands-on experience should leave you feeling comfortable working on real-world Rails projects, with real-world demands.

If you are in an architectural or development lead role, this book is not targeted to you, but should make you feel more comfortable discussing the pros and cons of Ruby on Rails adoption and ways to extend Rails to meet the particular needs of the project under your direction.

Finally, if you are a development manager, you should find the practical perspective of the book and our coverage of testing and tools especially interesting, and hopefully get some insight into why your developers are so excited about Ruby and Rails.

0.4 Prerequisites

The reader is assumed to have the following knowledge:

- Basic Ruby syntax and language constructs such as blocks
- Solid grasp of object-oriented principles and design patterns
- Basic understanding of relational databases and SQL
- Familiarity with how Rails applications are laid out and function
- Basic understanding of network protocols such as HTTP and SMTP
- Basic understanding of XML documents and web services
- Familiarity with transactional concepts such as ACID properties

As noted in the section “Book Structure,” this book does not progress from easy material in the front to harder material in the back. Some chapters do start out with fundamental, almost introductory material and push on to more advanced coverage. There are definitely sections of the text that experienced Rails developer will gloss over. However, I believe that there is new knowledge and inspiration in every chapter, for all skill levels.

This page intentionally left blank

Acknowledgments

A whole new set of players contributed to *The Rails™ 3 Way*, however I still need to thank some of my original supporters first. I can't say enough good things about Debra Williams Cauley, my editor at Addison-Wesley. She is an excellent coach and motivator and oh-so-caring of her authors. I love you, Deb! Also again I have to thank my long-term partner Desi McAdam and my kids Taylor and Liam for being super-supportive and understanding of my time constraints during the heaviest times of writing.

My team at Hashrocket has been an amazing source of encouragement and help during the preparation of *The Rails™ 3 Way*. My partners Marian and Mark made sure I had all the time and help needed, and were always ready with a hug or words of encouragement when the times got tough. Jon Larkowski and Tim “t pope” Pope spent hours with me at my apartment, sometimes every night of the week, to make sure that the book got finished. Eliza Brock and Tim Pope hacked a massive XSLT script that converted the original Word .doc manuscript files into L^AT_EX, enabling us to put the book into proper source control and make much more rapid progress than would otherwise be possible. Eliza, you are a freaking genius and an inspiration!

My friend Xavier Noria, Rails committer and former textbook reviewer, once again impressed us with his careful technical review and laser-focused feedback. Xavi picked up on dozens of omissions and errors that would otherwise have gone unnoticed. What a hero!

One of my oldest and closest friends, Durran Jordan, was a late and welcome addition to *The Rails™ 3 Way* team. He's the author of Mongoid, <http://mongoid.org>—one of the premier frameworks for using Mongo with Ruby and an up-and-coming personality in the NoSQL space. He's currently working on a NoSQL in Ruby title for this series and

provided some of the new content in this book concerning Active Model and background processing.

Chicago-based Rocketeers Josh Graham and Bernerd Schaefer also provided late-stage help, contributing material related to XML processing and Ajax. Other folks at Hashrocket that deserve acknowledgment include our director of operations and my longtime friend Sal Cardello, who controls resourcing and allowed me to take people away from billing to help me with the book. I also need to thank everyone else at Hashrocket who played supporting roles, including but not limited to, Rogelio Samour, Thais Camilo, Adam Lowe, “Big Tiger” Jim Remsik, Lar Van Der Jagt, Matt Yoho, Stephen Caudill, Robert Pitts, Sandro Turriate, Shay Arnette, and Veezus Kreist.

Thanks to David Black, James Adam, Trotter Cashion, Matt Pelletier, Matt Bauer, Jodi Showers, Pat Maddox, David Chelinski, Charles Brian Quinn, Patrik Naik, Diego Scataglioni, and everyone else who contributed to making *The Rails Way* such a success.

About the Author

Obie Fernandez is a recognized tech industry leader and local celebrity in the Jacksonville business community. He has been hacking computers since he got his first Commodore VIC-20 in the eighties, and found himself in the right place and time as a programmer on some of the first Java enterprise projects of the mid-nineties. He moved to Atlanta, Georgia, in 1998 and gained prominence as lead architect of local startup success MediaOcean. He also founded the Extreme Programming (later Agile Atlanta) User Group and was that group's president and organizer for several years. In 2004, he made the move back into the enterprise, tackling high-risk, progressive projects for world-renowned consultancy ThoughtWorks.

Obie has been evangelizing Ruby on Rails via online via blog posts and publications since early 2005, and earned himself quite a bit of notoriety (and trash talking) from his old friends in the Java open-source community. Since then, he has traveled around the world relentlessly promoting Rails at large industry conferences.

As CEO and founder of Hashrocket, one of the world's best web design and development consultancies, Obie specializes in orchestrating the creation of large-scale, web-based applications, both for startups and mission-critical enterprise projects. He still gets his hands dirty with code on at least a weekly basis and posts regularly on various topics to his popular technology weblog, <http://blog.obiefernandez.com>.

This page intentionally left blank

CHAPTER 9

Advanced Active Record

Active Record is a simple object-relational mapping (ORM) framework compared to other popular ORM frameworks, such as Hibernate in the Java world. Don't let that fool you, though: Under its modest exterior, Active Record has some pretty advanced features. To really get the most effectiveness out of Rails development, you need to have more than a basic understanding of Active Record—things like knowing when to break out of the one-table/one-class pattern, or how to leverage Ruby modules to keep your code clean and free of duplication.

In this chapter, we wrap up this book's comprehensive coverage of Active Record by reviewing callbacks, observers, single-table inheritance (STI), and polymorphic models. We also review a little bit of information about metaprogramming and Ruby domain-specific languages (DSLs) as they relate to Active Record.

9.1 Scopes

Scopes (or “named scopes” if you're old school) allow you define and chain query criteria in a declarative and reusable manner.

```
class Timesheet < ActiveRecord::Base
  scope :submitted, where(:submitted => true)
  scope :underutilized, where('total_hours < 40')
```

To declare a scope, use the `scope` class method, passing it a name as a symbol and some sort of query definition. If your query is known at load time, you can simply use `Arel` criteria methods like `where`, `order`, and `limit` to construct the definition as shown in the example. On the other hand, if you won't have all the parameters for your query until runtime, use a lambda as the second parameter. It will get evaluated whenever the scope is invoked.

```
class User < ActiveRecord::Base
  scope :delinquent, lambda { where('timesheets_updated_at < ?',
    1.week.ago)}
end
```

Invoke scopes as you would class methods.

```
>> User.delinquent
=> [#<User id: 2, timesheets_updated_at: "2010-01-07 01:56:29"...>]
```

9.1.1 Scope Parameters

You can pass arguments to scope invocations by adding parameters to the lambda you use to define the scope query.

```
class BillableWeek < ActiveRecord::Base
  scope :newer_than, lambda { |date| where('start_date > ?', date) }
end
```

Then pass the argument to the scope as you would normally.

```
BillableWeek.newer_than(Date.today)
```

9.1.2 Chaining Scopes

One of the beauties of scopes is that you can chain them together to create complex queries from simple ones:

```
>> Timesheet.underutilized.submitted
=> [#<Timesheet id: 3, submitted: true, total_hours: 37 ...>]
```

Scopes can be chained together for reuse within scope definitions themselves. For instance, let's say that we always want to constrain the result set of underutilized to submitted timesheets:

```
class Timesheet < ActiveRecord::Base
  scope :submitted, where(:submitted => true)
  scope :underutilized, submitted.where('total_hours < 40')
end
```

9.1.3 Scopes and has_many

In addition to being available at the class context, scopes are available automatically on has_many association attributes.

```
>> u = User.find 2
=> #<User id: 2, login: "obie"...>

>> u.timesheets.size
=> 3

>> u.timesheets.underutilized.size
=> 1
```

9.1.4 Scopes and Joins

You can use Arel's `join` method to create cross-model scopes. For instance, if we gave our recurring example `Timesheet` a `submitted_at` date attribute instead of just a boolean, we could add a scope to `User` allowing us to see who is late on their timesheet submission.

```
scope :tardy, lambda {
  joins(:timesheets).
  where("timesheets.submitted_at <= ?", 7.days.ago).
  group("users.id")
}
```

Arel's `to_sql` method is useful for debugging scope definitions and usage.

```
>> User.tardy.to_sql
=> "SELECT users.* FROM users
  INNER JOIN timesheets ON timesheets.user_id = users.id
  WHERE (timesheets.submitted_at <= '2010-07-06 15:27:05.117700')
  GROUP BY users.id" # query formatted nicely for the book
```

Note that as demonstrated in the example, it's a good idea to use unambiguous column references (including table name) in cross-model scope definitions so that Arel doesn't get confused.

9.1.5 Scope Combinations

Our example of a cross-model scope violates good object-oriented design principles: it contains the logic for determining whether or not a `Timesheet` is submitted, which is code that properly belongs in the `Timesheet` class. Luckily we can use Arel's `merge` method (aliased as `&`) to fix it. First we put the late logic where it belongs, in `Timesheet`:

```
scope :late, lambda { where("timesheet.submitted_at <= ?", 7.days.ago) }
```

Then we use our new `late` scope in `tardy`:

```
scope :tardy, lambda {
  joins(:timesheets).group("users.id") & Timesheet.late
}
```

If you have trouble with this technique, make absolutely sure that your scopes' clauses refer to fully qualified column names. (In other words, don't forget to prefix column names with tables.) The console and `to_sql` method is your friend for debugging.

9.1.6 Default Scopes

There may arise use cases where you want certain conditions applied to the finders for your model. Consider our timesheet application has a default view of open timesheets—we can use a default scope to simplify our general queries.

```
class Timesheet < ActiveRecord::Base
  default_scope :where(:status => "open")
end
```

Now when we query for our Timesheets, by default the open condition will be applied:

```
>> Timesheet.all.map(&:status)
=> ["open", "open", "open"]
```

Default scopes also get applied to your models when building or creating them, which can be a great convenience or a nuisance if you are not careful. In our previous example, all new Timesheets will be created with a status of “open.”

```
>> Timesheet.new
=> #<Timesheet id: nil, status: "open">
>> Timesheet.create
=> #<Timesheet id: 1, status: "open">
```

You can override this behavior by providing your own conditions or scope to override the default setting of the attributes.

```
>> Timesheet.where(:status => "new").new
=> #<Timesheet id: nil, status: "new">
>> Timesheet.where(:status => "new").create
=> #<Timesheet id: 1, status: "new">
```

There may be cases where at runtime you want to create a scope and pass it around as a first class object leveraging your default scope. In this case, Active Record provides the `scoped` method.

```
>> timesheets = Timesheet.scoped.order("submitted_at DESC")
=> [#<Timesheet id: 1, status: "open">]
>> timesheets.where(:name => "Durran Jordan")
=> []
```

There’s another approach to scopes that provides a sleeker syntax, `scoping`, which allows the chaining of scopes via nesting within a block.

```
>> Timesheet.order("submitted_at DESC").scoping do
>>   Timesheets.all
>> end
=> #<Timesheet id: 1, status: "open">
```

That's pretty nice, but what if we *don't* want our default scope to be included in our queries? In this case Active Record takes care of us through the `unscoped` method.

```
>> Timesheet.unscoped.order("submitted_at DESC")
=> [#<Timesheet id: 2, status: "submitted">]
```

Similarly to overriding our default scope with a relation when creating new objects, we can supply `unscoped` as well to remove the default attributes.

```
>> Timesheet.unscoped.new
=> #<Timesheet id: nil, status: nil>
```

9.1.7 Using Scopes for CRUD

You have a wide range of Active Record's CRUD methods available on scopes, which gives you some powerful abilities. For instance, let's give all our underutilized timesheets some extra hours.

```
>> u.timesheets.underutilized.collect(&:total_hours)
=> [37, 38]

>> u.timesheets.underutilized.update_all("total_hours = total_hours + 2")
=> 2

>> u.timesheets.underutilized.collect(&:total_hours)
=> [37, 38] # whoops, cached result

>> u.timesheets(true).underutilized.collect(&:total_hours)
=> [39] # results after telling association to reload
```

Scopes including a where clause using hashed conditions will populate attributes of objects built off of them with those attributes as default values. Admittedly it's a bit difficult to think of a plausible use case for this feature, but we'll show it in an example. First, we add the following scope to `Timesheet`:

```
scope :perfect, submitted.where(:total_hours => 40)
```

Now, building an object on the `perfect` scope should give us a submitted timesheet with 40 hours.

```
> Timesheet.perfect.build
=> #<Timesheet id: nil, submitted: true, user_id: nil, total_hours: 40
...>
```

As you've probably realized by now, the new Arel underpinnings of Active Record are tremendously powerful and truly elevate the Rails 3 platform.

9.2 Callbacks

This advanced feature of Active Record allows the savvy developer to attach behavior at a variety of different points along a model's life cycle, such as after initialization, before database records are inserted, updated or removed, and so on.

Callbacks can do a variety of tasks, ranging from simple things such as logging and massaging of attribute values prior to validation, to complex calculations. Callbacks can halt the execution of the life-cycle process taking place. Some callbacks can even modify the behavior of the model class on the fly. We'll cover all of those scenarios in this section, but first let's get a taste of what a callback looks like. Check out the following silly example:

```
class Beethoven < ActiveRecord::Base
  before_destroy :last_words

  protected

  def last_words
    logger.info "Friends applaud, the comedy is over"
  end
end
```

So prior to dying (ehrm, being `destroy`'d), the last words of the `Beethoven` class will always be logged for posterity. As we'll see soon, there are 14 different opportunities to add behavior to your model in this fashion. Before we get to that list, let's cover the mechanics of registering a callback.

9.2.1 Callback Registration

Overall, the most common way to register a callback method is to declare it at the top of the class using a typical Rails macro-style class method. However, there's a less verbose way to do it also. Simply implement the callback as a method in your class. In other words, I could have coded the prior example as follows:

```
class Beethoven < ActiveRecord::Base

  protected

  def before_destroy
    logger.info "Friends applaud, the comedy is over"
  end
end
```

This is a rare case of the less-verbose solution being bad. In fact, it is almost always preferable, dare I say it is the Rails way, to use the callback macros over implementing

callback methods, for the following reasons:

- Macro-style callback declarations are added near the top of the class definition, making the existence of that callback more evident versus a method body potentially buried later in the file.
- Macro-style callbacks add callback methods to a queue. That means that more than one method can be hooked into the same slot in the life cycle. Callbacks will be invoked in the order in which they were added to the queue.
- Callback methods for the same hook can be added to their queue at different levels of an inheritance hierarchy and still work—they won't override each other the way that methods would.
- Callbacks defined as methods on the model are always called last.

9.2.2 One-Liners

Now, if (and only if) your callback routine is really short,¹ you can add it by passing a block to the callback macro. We're talking one-liners!

```
class Napoleon < ActiveRecord::Base
  before_destroy { logger.info "Josephine..." }
  ...
end
```

As of Rails 3, the block passed to a callback is executed via `instance_eval` so that its scope is the record itself (versus needing to act on a passed in record variable). The following example implements “paranoid” model behavior, covered later in the chapter.

```
class Account < ActiveRecord::Base
  before_destroy { update_attribute(:deleted_at, Time.now); false }
  ...
end
```

9.2.3 Protected or Private

Except when you're using a block, the access level for callback methods should always be protected or private. It should never be public, since callbacks should never be called from code outside the model.

1. If you are browsing old Rails source code, you might come across callback macros receiving a short string of Ruby code to be evaluated in the binding of the model object. That way of adding callbacks was deprecated in Rails 1.2, because you're always better off using a block in those situations.

Believe it or not, there are even more ways to implement callbacks, but we'll cover those techniques further along in the chapter. For now, let's look at the lists of callback hooks available.

9.2.4 Matched before/after Callbacks

In total, there are 14 types of callbacks you can register on your models! Twelve of them are matching *before/after* callback pairs, such as `before_validation` and `after_validation`. (The other two, `after_initialize` and `after_find`, are special, and we'll discuss them later in this section.)

List of Callbacks

This is the list of callback hooks available during a `save` operation. (The list varies slightly depending on whether you're saving a new or existing record.)

- `before_validation`
- `before_validation_on_create`
- `after_validation`
- `after_validation_on_create`
- `before_save`
- `before_create` (for new records) and `before_update` (for existing records)
- (Database actually gets an `INSERT` or `UPDATE` statement here)
- `after_create` (for new records) and `after_update` (for existing records)
- `after_save`

Delete operations have their own two callbacks:

- `before_destroy`
- (Database actually gets a `DELETE` statement here)
- `after_destroy` is called after all attributes have been frozen (read-only)

Additionally transactions have callbacks as well, for when you want actions to occur after the database is guaranteed to be in a permanent state. Note that only "after" callbacks exist here because of the nature of transactions—it's a bad idea to be able to interfere with the actual operation itself.

- `after_commit`
- `after_commit_on_create`
- `after_commit_on_update`
- `after_commit_on_destroy`
- `after_rollback`
- `after_rollback_on_create`
- `after_rollback_on_update`
- `after_rollback_on_destroy`

9.2.5 Halting Execution

If you return a boolean `false` (not `nil`) from a callback method, Active Record halts the execution chain. No further callbacks are executed. The `save` method will return `false`, and `save!` will raise a `RecordNotSaved` error.

Keep in mind that because the last expression of a Ruby method is returned implicitly, it is a pretty common bug to write a callback that halts execution unintentionally. If you have an object with callbacks that mysteriously fails to save, make sure you aren't returning `false` by mistake.

9.2.6 Callback Usages

Of course, the callback you should use for a given situation depends on what you're trying to accomplish. The best I can do is to serve up some examples to inspire you with your own code.

Cleaning Up Attribute Formatting with `before_validate_on_create`

The most common examples of using `before_validation` callbacks have to do with cleaning up user-entered attributes. For example, the following `CreditCard` class cleans up its `number` attribute so that false negatives don't occur on validation:

```
class CreditCard < ActiveRecord::Base
  ...

  def before_validation_on_create
    # Strip everything in the number except digits
    self.number = number.gsub(/[^0-9]/, "")
  end
end
```

Geocoding with **before_save**

Assume that you have an application that tracks addresses and has mapping features. Addresses should always be geocoded before saving, so that they can be displayed rapidly on a map later.²

As is often the case, the wording of the requirement itself points you in the direction of the `before_save` callback:

```
class Address < ActiveRecord::Base
  include GeoKit::Geocoders

  before_save :geolocate
  validates_presence_of :street, :city, :state, :zip
  ...

  def to_s
    "#{street} #{city}, #{state} #{zip}"
  end

  protected

  def geolocate
    res = GoogleGeocoder.geocode(to_s)
    self.latitude = res.lat
    self.longitude = res.lng
  end
end
```

Before we move on, there are a couple of additional considerations. The preceding code works great if the geocoding succeeds, but what if it doesn't? Do we still want to allow the record to be saved? If not, we should halt the execution chain:

```
def geolocate
  res = GoogleGeocoder.geocode(to_s)
  return false if not res.success # halt execution

  self.latitude = res.lat
  self.longitude = res.lng
end
```

The only problem remaining is that we give the rest of our code (and by extension, the end user) no indication of why the chain was halted. Even though we're not in a validation routine, I think we can put the `errors` collection to good use here:

```
def geolocate
  res = GoogleGeocoder.geocode(to_s)
```

2. I recommend the excellent GeoKit for Rails plugin available at <http://geokit.rubyforge.org/>.

```
if res.success
  self.latitude = res.lat
  self.longitude = res.lng
else
  errors[:base] << "Geocoding failed. Please check address."
  return false
end
end
```

If the geocoding fails, we add a base error message (for the whole object) and halt execution, so that the record is not saved.

Exercise Your Paranoia with **before_destroy**

What if your application has to handle important kinds of data that, once entered, should never be deleted? Perhaps it would make sense to hook into Active Record's destroy mechanism and somehow mark the record as deleted instead?

The following example depends on the `accounts` table having a `deleted_at` date-time column.

```
class Account < ActiveRecord::Base
  ...
  def before_destroy
    update_attribute(:deleted_at, Time.now)
    false
  end
end
```

I chose to implement it as a callback method so that I am guaranteed it will execute last in the `before_destroy` queue. It returns `false` so that execution is halted and the underlying record is not actually deleted from the database.³

It's probably worth mentioning that there are ways that Rails allows you to unintentionally circumvent `before_destroy` callbacks:

- The `delete` and `delete_all` class methods of `ActiveRecord::Base` are almost identical. They remove rows directly from the database without instantiating the corresponding model instances, which means no callbacks will occur.

3. Real-life implementation of the example would also need to modify all finders to include `deleted_at` is `NULL` conditions; otherwise, the records marked deleted would continue to show up in the application. That's not a trivial undertaking, and luckily you don't need to do it yourself. There's a Rails plugin named `ActsAsParanoid` by Rick Olson that does exactly that, and you can find it at http://svn.techno-weenie.net/projects/plugins/acts_as_paranoid.

- Model objects in associations defined with the option `:dependent => :delete_all` will be deleted directly from the database when removed from the collection using the association's `clear` or `delete` methods.

Cleaning Up Associated Files with **after_destroy**

Model objects that have files associated with them, such as attachment records and uploaded images, can clean up after themselves when deleted using the `after_destroy` callback. The following method from Rick Olson's old `AttachmentFu`⁴ plugin is a good example:

```
# Destroys the file. Called in the after_destroy callback
def destroy_file
  FileUtils.rm(full_filename)
  ...
rescue
  logger.info "Exception destroying #{full_filename ... }"
  logger.warn $!.backtrace.collect { |b| " > #{b}" }.join("\n")
end
```

9.2.7 Special Callbacks: **after_initialize** and **after_find**

The `after_initialize` callback is invoked whenever a new Active Record model is instantiated (either from scratch or from the database). Having it available prevents you from having to muck around with overriding the actual `initialize` method.

The `after_find` callback is invoked whenever Active Record loads a model object from the database, and is actually called before `after_initialize`, if both are implemented. Because `after_find` and `after_initialize` are called for each object found and instantiated by finders, performance constraints dictate that they can only be added as methods, and not via the callback macros.

What if you want to run some code only the first time that a model is ever instantiated, and not after each database load? There is no native callback for that scenario, but you can do it using the `after_initialize` callback. Just add a condition that checks to see if it is a new record:

```
def after_initialize
  if new?
    ...
  end
end
```

4. Get `AttachmentFu` at http://svn.techno-weenie.net/projects/plugins/attachment_fu.

In a number of Rails apps that I've written, I've found it useful to capture user preferences in a serialized hash associated with the `User` object. The `serialize` feature of Active Record models makes this possible, since it transparently persists Ruby object graphs to a text column in the database. Unfortunately, you can't pass it a default value, so I have to set one myself:

```
class User < ActiveRecord::Base
  serialize :preferences # defaults to nil
  ...

  protected

  def after_initialize
    self.preferences ||= Hash.new
  end
end
```

Using the `after_initialize` callback, I can automatically populate the `preferences` attribute of my user model with an empty hash, so that I never have to worry about it being `nil` when I access it with code such as `user.preferences[:show_help_text] = false`.

Ruby's metaprogramming capabilities combined with the ability to run code whenever a model is loaded using the `after_find` callback are a powerful mix. Since we're not done learning about callbacks yet, we'll come back to uses of `after_find` later on in the chapter, in the section "Modifying Active Record Classes at Runtime."

9.2.8 Callback Classes

It is common enough to want to reuse callback code for more than one object that Rails gives you a way to write callback classes. All you have to do is pass a given callback queue an object that responds to the name of the callback and takes the model object as a parameter.

Here's our `MarkDeleted` example from the previous section as a callback class:

```
class MarkDeleted
  def self.before_destroy(model)
    model.update_attribute(:deleted_at, Time.now)
    return false
  end
end
```

The behavior of `MarkDeleted` is stateless, so I added the callback as a class method. Now you don't have to instantiate `MarkDeleted` objects for no good reason. All you

do is pass the class to the callback queue for whichever models you want to have the mark-deleted behavior:

```
class Account < ActiveRecord::Base
  before_destroy MarkDeleted
  ...
end

class Invoice < ActiveRecord::Base
  before_destroy MarkDeleted
  ...
end
```

Multiple Callback Methods in One Class

There's no rule that says you can't have more than one callback method in a callback class. For example, you might have special audit log requirements to implement:

```
class Auditor
  def initialize(audit_log)
    @audit_log = audit_log
  end

  def after_create(model)
    @audit_log.created(model.inspect)
  end

  def after_update(model)
    @audit_log.updated(model.inspect)
  end

  def after_destroy(model)
    @audit_log.destroyed(model.inspect)
  end
end
```

To add audit logging to an Active Record class, you would do the following:

```
class Account < ActiveRecord::Base
  after_create Auditor.new(DEFAULT_AUDIT_LOG)
  after_update Auditor.new(DEFAULT_AUDIT_LOG)
  after_destroy Auditor.new(DEFAULT_AUDIT_LOG)
  ...
end
```

Wow, that's ugly, having to add three `Auditor`s on three lines. We could extract a local variable called `auditor`, but it would still be repetitive. This might be an opportunity to take advantage of Ruby's open classes, the fact that you can modify classes that aren't part of your application.

Wouldn't it be better to simply say `acts_as_audited` at the top of the model that needs auditing? We can quickly add it to the `ActiveRecord::Base` class, so that it's available for all our models.

On my projects, the file where “quick and dirty” code like the method in Listing 9.1 would reside is `lib/core_ext/active_record_base.rb`, but you can put it anywhere you want. You could even make it a plugin (as detailed in Chapter 19, “Extending Rails with Plugins”).

Listing 9.1 A quick-and-dirty “acts as audited” method

```
class ActiveRecord::Base
  def self.acts_as_audited(audit_log=DEFAULT_AUDIT_LOG)
    auditor = Auditor.new(audit_log)
    after_create auditor
    after_update auditor
    after_destroy auditor
  end
end
```

Now, the top of `Account` is a lot less cluttered:

```
class Account < ActiveRecord::Base
  acts_as_audited
```

Testability

When you add callback methods to a model class, you pretty much have to test that they're functioning correctly in conjunction with the model to which they are added. That may or may not be a problem. In contrast, callback classes are super-easy to test in isolation.

```
def test_auditor_logs_created
  (model = mock).expects(:inspect).returns('foo')
  (log = mock).expects(:created).with('foo')
  Auditor.new(log).after_create(model)
end
```

9.3 Calculation Methods

All Active Record classes have a `calculate` method that provides easy access to aggregate function queries in the database. Methods for `count`, `sum`, `average`, `minimum`, and `maximum` have been added as convenient shortcuts.

Options such as `conditions`, `:order`, `:group`, `:having`, and `:joins` can be passed to customize the query.

There are two basic forms of output:

Single aggregate value The single value is type cast to `Fixnum` for `COUNT`, `Float` for `AVG`, and the given column's type for everything else.

Grouped values This returns an ordered hash of the values and groups them by the `:group` option. It takes either a column name, or the name of a `belongs_to` association.

The following options are available to all calculation methods:

- :conditions** An SQL fragment like `"administrator = 1"` or `["user_name = ?", username]`. See `conditions` in the intro to `ActiveRecord::Base`.
- :include** Eager loading, see `Associations` for details. Since calculations don't load anything, the purpose of this is to access fields on joined tables in your `conditions`, `order`, or `group` clauses.
- :joins** An SQL fragment for additional joins like `"LEFT JOIN comments ON comments.post_id = id"`. (Rarely needed). The records will be returned read-only since they will have attributes that do not correspond to the table's columns.
- :order** An SQL fragment like `"created_at DESC, name"` (really only used with `GROUP BY` calculations).
- :group** An attribute name by which the result should be grouped. Uses the `GROUP BY` SQL-clause.
- :select** By default, this is `*` as in `SELECT * FROM`, but can be changed if you, for example, want to do a join, but not include the joined columns.
- :distinct** Set this to `true` to make this a distinct calculation, such as `SELECT COUNT(DISTINCT posts.id) ...`

The following examples illustrate the usage of various calculation methods.

```
Person.calculate(:count, :all) # The same as Person.count

# SELECT AVG(age) FROM people
Person.average(:age)

# Selects the minimum age for everyone with a last name other than
'Drake'
Person.minimum(:age).where('last_name <> ?', 'Drake')
```

```
# Selects the minimum age for any family without any minors
Person.minimum(:age).having('min(age) > 17').group(:last_name)
```

9.3.1 **average(column_name, *options)**

Calculates the average value on a given column. The first parameter should be a symbol identifying the column to be averaged.

9.3.2 **count(column_name, *options)**

Count operates using three different approaches. Count without parameters will return a count of all the rows for the model. Count with a `column_name` will return a count of all the rows for the model with the supplied column present. Lastly, count using `:options` will find the row count matched by the options used. In the last case you would send an options hash as the only parameter. 213

```
total_contacts = person.contacts.count(:from => "contact_cards")
```

Options are the same as with all other calculations methods with the additional option of `:from` which is by default the name of the table name of the class, however it can be changed to a different table name or even that of a database view. Remember that `Person.count(:all)` will not work because `:all` will be treated as a condition, you should use `Person.count` instead.

9.3.3 **maximum(column_name, *options)**

Calculates the maximum value on a given column. The first parameter should be a symbol identifying the column to be calculated.

9.3.4 **minimum(column_name, *options)**

Calculates the minimum value on a given column. The first parameter should be a symbol identifying the column to be calculated.

9.3.5 **sum(column_name, *options)**

Calculates a summed value in the database using SQL. The first parameter should be a symbol identifying the column to be summed.

9.4 Observers

The single responsibility principle is a very important tenet of object-oriented programming. It compels us to keep a class focused on a single concern. As you’ve learned in the previous section, callbacks are a useful feature of Active Record models that allow us to hook in behavior at various points of a model object’s life cycle. Even if we pull that extra behavior out into callback classes, the hook still requires code changes in the model class definition itself. On the other hand, Active Record gives us a way to hook in to models that is completely transparent: Observers.

Here is the functionality of our old `Auditor` callback class as an observer of `Account` objects:

```
class AccountObserver < ActiveRecord::Observer
  def after_create(model)
    DEFAULT_AUDIT_LOG.created(model.inspect)
  end

  def after_update(model)
    DEFAULT_AUDIT_LOG.updated(model.inspect)
  end

  def after_destroy(model)
    DEFAULT_AUDIT_LOG.destroyed(model.inspect)
  end
end
```

9.4.1 Naming Conventions

When `ActiveRecord::Observer` is subclassed, it breaks down the name of the subclass by stripping off the “Observer” part. In the case of our `AccountObserver` in the preceding example, it would know that you want to observe the `Account` class. However, that’s not always desirable behavior. In fact, with general-purpose code such as our `Auditor`, it’s positively a step backward, so it is possible to overrule the naming convention with the use of the `observe` macro-style method. We still extend `ActiveRecord::Observer`, but we can call the subclass whatever we want and tell it explicitly what to observe using the `observe` method, which accepts one or more arguments.

```
class Auditor < ActiveRecord::Observer
  observe Account, Invoice, Payment

  def after_create(model)
    DEFAULT_AUDIT_LOG.created(model.inspect)
  end
end
```

```

def after_update(model)
  DEFAULT_AUDIT_LOG.updated(model.inspect)
end

def after_destroy(model)
  DEFAULT_AUDIT_LOG.destroyed(model.inspect)
end
end

```

9.4.2 Registration of Observers

If there weren't a place for you to tell Rails which observers to load, they would never get loaded at all, since they're not referenced from any other code in your application. Register observers with the following kind of code in an initializer:

```

# Activate observers that should always be running
ActiveRecord::Base.observers = Auditor

```

9.4.3 Timing

Observers are notified after the in-object callbacks are triggered.⁵ It's not possible to act on the whole object from an observer without having the object's own callbacks executed first.

Durran says...

For those of us who love to be organized, you can now put your observers in a separate directory under `app` if your heart desires. You won't need to perform custom loading anymore since Rails now loads all files under the `app` directory automatically.

9.5 Single-Table Inheritance (STI)

A lot of applications start out with a `User` model of some sort. Over time, as different kinds of users emerge, it might make sense to make a greater distinction between them. `Admin` and `Guest` classes are introduced, as subclasses of `User`. Now, the shared behavior can reside in `User`, and subtype behavior can be pushed down to subclasses. However, all user data can still reside in the `users` table—all you need to do is introduce a `type` column that will hold the name of the class to be instantiated for a given row.

5. <https://rails.lighthouseapp.com/projects/8994/tickets/230> contains an interesting discussion about callback execution order.

To continue explaining single-table inheritance, let's turn back to our example of a recurring `Timesheet` class. We need to know how many `billable_hours` are outstanding for a given user. The calculation can be implemented in various ways, but in this case we've chosen to write a pair of class and instance methods on the `Timesheet` class:

```
class Timesheet < ActiveRecord::Base
  ...

  def billable_hours_outstanding
    if submitted?
      billable_weeks.map(&:total_hours).sum
    else
      0
    end
  end

  def self.billable_hours_outstanding_for(user)
    user.timesheets.map(&:billable_hours_outstanding).sum
  end
end
```

I'm not suggesting that this is good code. It works, but it's inefficient and that `if/else` condition is a little fishy. Its shortcomings become apparent once requirements emerge about marking a `Timesheet` as paid. It forces us to modify `Timesheet`'s `billable_hours_outstanding` method again:

```
def billable_hours_outstanding
  if submitted? && not paid?
    billable_weeks.map(&:total_hours).sum
  else
    0
  end
end
```

That latest change is a clear violation of the open-closed principle,⁶ which urges you to write code that is open for extension, but closed for modification. We know that we violated the principle, because we were forced to change the `billable_hours_outstanding` method to accommodate the new `Timesheet` status. Though it may not seem like a large problem in our simple example, consider the amount of conditional code that will end up in the `Timesheet` class once we start having to implement functionality such as `paid_hours` and `unsubmitted_hours`.

6. http://en.wikipedia.org/wiki/Open/closed_principle has a good summary.

So what's the answer to this messy question of the constantly changing conditional? Given that you're reading the section of the book about single-table inheritance, it's probably no big surprise that we think one good answer is to use object-oriented inheritance. To do so, let's break our original `Timesheet` class into four classes.

```
class Timesheet < ActiveRecord::Base
  # non-relevant code omitted

  def self.billable_hours_outstanding_for(user)
    user.timesheets.map(&:billable_hours_outstanding).sum
  end
end

class DraftTimesheet < Timesheet
  def billable_hours_outstanding
    0
  end
end

class SubmittedTimesheet < Timesheet
  def billable_hours_outstanding
    billable_weeks.map(&:total_hours).sum
  end
end
```

Now when the requirements demand the ability to calculate partially paid timesheets, we need only add some behavior to a `PaidTimesheet` class. No messy conditional statements in sight!

```
class PaidTimesheet < Timesheet
  def billable_hours_outstanding
    billable_weeks.map(&:total_hours).sum - paid_hours
  end
end
```

9.5.1 Mapping Inheritance to the Database

Mapping object inheritance effectively to a relational database is not one of those problems with a definitive solution. We're only going to talk about the one mapping strategy that Rails supports natively, which is single-table inheritance, called STI for short.

In STI, you establish one table in the database to hold all of the records for any object in a given inheritance hierarchy. In Active Record STI, that one table is named after the top parent class of the hierarchy. In the example we've been considering, that table would be named `timesheets`.

Hey, that's what it was called before, right? Yes, but to enable STI we have to add a `type` column to contain a string representing the type of the stored object. The following migration would properly set up the database for our example:

```
class AddTypeToTimesheet < ActiveRecord::Migration
  def self.up
    add_column :timesheets, :type, :string
  end

  def self.down
    remove_column :timesheets, :type
  end
end
```

No default value is needed. Once the `type` column is added to an Active Record model, Rails will automatically take care of keeping it populated with the right value. Using the console, we can see this behavior in action:

```
>> d = DraftTimesheet.create
>> d.type
=> 'DraftTimesheet'
```

When you try to find an object using the `find` methods of a base STI class, Rails will automatically instantiate objects using the appropriate subclass. This is especially useful in polymorphic situations, such as the timesheet example we've been describing, where we retrieve all the records for a particular user and then call methods that behave differently depending on the object's class.

```
>> Timesheet.find(:first)
=> #<DraftTimesheet:0x2212354...>
```

Sebastian says...

The word “`type`” is a very common column name and you might have plenty of uses for it not related to STI—which is why it's very likely you've experienced an `ActiveRecord::SubclassNotFound` error. Rails will read the “`type`” column of your `Car` class and try to find an “SUV” class that doesn't exist. The solution is simple: Tell Rails to use another column for STI with the following code:

```
set_inheritance_column "not_sti"
```

Note

Rails won't complain about the missing column; it will simply ignore it. Recently, the error message was reworded with a better explanation, but too many developers skim error messages and then spend an hour trying to figure out what's wrong with their models. (A lot of people skim sidebar columns too when reading books, but hey, at least I am doubling their chances of learning about this problem.)

9.5.2 STI Considerations

Although Rails makes it extremely simple to use single-table inheritance, there are a few caveats that you should keep in mind.

To begin with, you cannot have an attribute on two different subclasses with the same name but a different type. Since Rails uses one table to store all subclasses, these attributes with the same name occupy the same column in the table. Frankly, there's not much of a reason why that should be a problem unless you've made some pretty bad data-modeling decisions.

More importantly, you need to have one column per attribute on any subclass and any attribute that is not shared by all the subclasses must accept `nil` values. In the recurring example, `PaidTimesheet` has a `paid_hours` column that is not used by any of the other subclasses. `DraftTimesheet` and `SubmittedTimesheet` will not use the `paid_hours` column and leave it as null in the database. In order to validate data for columns not shared by all subclasses, you must use Active Record validations and not the database.

Third, it is not a good idea to have subclasses with too many unique attributes. If you do, you will have one database table with many null values in it. Normally, a tree of subclasses with a large number of unique attributes suggests that something is wrong with your application design and that you should refactor. If you have an STI table that is getting out of hand, it is time to reconsider your decision to use inheritance to solve your particular problem. Perhaps your base class is too abstract?

Finally, legacy database constraints may require a different name in the database for the `type` column. In this case, you can set the new column name using the class method `set_inheritance_column` in the base class. For the `Timesheet` example, we could do the following:

```
class Timesheet < ActiveRecord::Base
  set_inheritance_column 'object_type'
end
```

Now Rails will automatically populate the `object_type` column with the object's type.

9.5.3 STI and Associations

It seems pretty common for applications, particularly data-management ones, to have models that are very similar in terms of their data payload, mostly varying in their behavior and associations to each other. If you used object-oriented languages prior to Rails, you're probably already accustomed to breaking down problem domains into hierarchical structures.

Take for instance, a Rails application that deals with the population of states, counties, cities, and neighborhoods. All of these are places, which might lead you to define an STI class named `Place` as shown in Listing 9.2. I've also included the database schema for clarity:⁷

Listing 9.2 The places database schema and the place class

```
# == Schema Information
#
# Table name: places
#
# id :integer(11) not null, primary key
# region_id :integer(11)
# type :string(255)
# name :string(255)
# description :string(255)
# latitude :decimal(20, 1)
# longitude :decimal(20, 1)
# population :integer(11)
# created_at :datetime
# updated_at :datetime

class Place < ActiveRecord::Base
end
```

`Place` is in essence an abstract class. It should not be instantiated, but there is no foolproof way to enforce that in Ruby. (No big deal, this isn't Java!) Now let's go ahead

7. For autogenerated schema information added to the top of your model classes, try Dave Thomas's `annotate_models` plugin at <http://svn.pragprog.com/Public/plugins/>

and define concrete subclasses of `Place`:

```
class State < Place
  has_many :counties, :foreign_key => 'region_id'
end

class County < Place
  belongs_to :state, :foreign_key => 'region_id'
  has_many :cities, :foreign_key => 'region_id'
end

class City < Place
  belongs_to :county, :foreign_key => 'region_id'
end
```

You might be tempted to try adding a `cities` association to `State`, knowing that `has_many :through` works with both `belongs_to` and `has_many` target associations. It would make the `State` class look something like this:

```
class State < Place
  has_many :counties, :foreign_key => 'region_id'
  has_many :cities, :through => :counties
end
```

That would certainly be cool, if it worked. Unfortunately, in this particular case, since there's only one underlying table that we're querying, there simply isn't a way to distinguish among the different kinds of objects in the query:

```
Mysql::Error: Not unique table/alias: 'places': SELECT places.* FROM
places INNER JOIN places ON places.region_id = places.id WHERE
((places.region_id = 187912) AND ((places.type = 'County'))) AND
((places.`type` = 'City' ))
```

What would we have to do to make it work? Well, the most realistic would be to use specific foreign keys, instead of trying to overload the meaning of `region_id` for all the subclasses. For starters, the `places` table would look like the example in Listing 9.3.

Listing 9.3 The places database schema revised

```
# == Schema Information
#
# Table name: places
#
# id :integer(11) not null, primary key
# state_id :integer(11)
# county_id :integer(11)
# type :string(255)
# name :string(255)
# description :string(255)
```

```
# latitude :decimal(20, 1)
# longitude :decimal(20, 1)
# population :integer(11)
# created_at :datetime
# updated_at :datetime
```

The subclasses would be simpler without the `:foreign_key` options on the associations. Plus you could use a regular `has_many` relationship from `State` to `City`, instead of the more complicated `has_many :through`.

```
class State < Place
  has_many :counties
  has_many :cities
end
```

```
class County < Place
  belongs_to :state
  has_many :cities
end
```

```
class City < Place
  belongs_to :county
end
```

Of course, all those null columns in the `places` table won't win you any friends with relational database purists. That's nothing, though. Just a little bit later in this chapter we'll take a second, more in-depth look at polymorphic `has_many` relationships, which will make the purists positively hate you.

9.6 Abstract Base Model Classes

In contrast to single-table inheritance, it is possible for Active Record models to share common code via inheritance and still be persisted to different database tables. In fact, every Rails developer uses an abstract model in their code whether they realize it or not: `ActiveRecord::Base`.⁸

The technique involves creating an abstract base model class that persistent subclasses will extend. It's actually one of the simpler techniques that we broach in this chapter. Let's take the `Place` class from the previous section (refer to Listing 9.3) and revise it to

8. <http://m.onkey.org/2007/12/9/namespaced-models>

be an abstract base class in Listing 9.4. It's simple really—we just have to add one line of code:

Listing 9.4 The abstract place class

```
class Place < ActiveRecord::Base
  self.abstract_class = true
end
```

Marking an Active Record model abstract is essentially the opposite of making it an STI class with a type column. You're telling Rails: "Hey, I don't want you to assume that there is a table named `places`."

In our running example, it means we would have to establish tables for states, counties, and cities, which might be exactly what we want. Remember though, that we would no longer be able to query across subtypes with code like `Place.all`.

Abstract classes is an area of Rails where there aren't too many hard-and-fast rules to guide you—experience and gut feeling will help you out.

In case you haven't noticed yet, both class and instance methods are shared down the inheritance hierarchy of Active Record models. So are constants and other class members brought in through module inclusion. That means we can put all sorts of code inside `Place` that will be useful to its subclasses.

9.7 Polymorphic has_many Relationships

Rails gives you the ability to make one class `belong_to` more than one type of another class, as eloquently stated by blogger Mike Bayer:

The "polymorphic association," on the other hand, while it bears some resemblance to the regular polymorphic union of a class hierarchy, is not really the same since you're only dealing with a particular association to a single target class from any number of source classes, source classes which don't have anything else to do with each other; i.e., they aren't in any particular inheritance relationship and probably are all persisted in completely different tables. In this way, the polymorphic association has a lot less to do with object inheritance and a lot more to do with aspect-oriented programming (AOP); a particular concept needs to be applied to a divergent set of entities which otherwise are not directly related. Such a concept is referred to as a cross-cutting concern, such as, all the entities in your domain need to support a history log of all changes to a common logging table. In the AR example, an `Order` and a `User` object are illustrated to both require links to an `Address` object.⁹

9. <http://techspot.zzzEEK.org/?p=13>

In other words, this is not polymorphism in the typical object-oriented sense of the word; rather, it is something unique to Rails.

9.7.1 In the Case of Models with Comments

In our recurring Time and Expenses example, let's assume that we want both `BillableWeek` and `Timesheet` to have many comments (a shared `Comment` class). A naive way to solve this problem might be to have the `Comment` class belong to both the `BillableWeek` and `Timesheet` classes and have `billable_week_id` and `timesheet_id` as columns in its database table.

```
class Comment < ActiveRecord::Base
  belongs_to :timesheet
  belongs_to :expense_report
end
```

I call that approach is naive because it would be difficult to work with and hard to extend. Among other things, you would need to add code to the application to ensure that a `Comment` never belonged to both a `BillableWeek` and a `Timesheet` at the same time. The code to figure out what a given comment is attached to would be cumbersome to write. Even worse, every time you want to be able to add comments to another type of class, you'd have to add another nullable foreign key column to the comments table.

Rails solves this problem in an elegant fashion, by allowing us to define what it terms polymorphic associations, which we covered when we described the `:polymorphic => true` option of the `belongs_to` association in Chapter 7, Active Record Associations.

The Interface

Using a polymorphic association, we need define only a single `belongs_to` and add a pair of related columns to the underlying database table. From that moment on, any class in our system can have comments attached to it (which would make it commentable), without needing to alter the database schema or the `Comment` model itself.

```
class Comment < ActiveRecord::Base
  belongs_to :commentable, :polymorphic => true
end
```

There isn't a `Commentable` class (or module) in our application. We named the association `:commentable` because it accurately describes the interface of objects that will be

associated in this way. The name `:commentable` will turn up again on the other side of the association:

```
class Timesheet < ActiveRecord::Base
  has_many :comments, :as => :commentable
end
```

```
class BillableWeek < ActiveRecord::Base
  has_many :comments, :as => :commentable
end
```

Here we have the friendly `has_many` association using the `:as` option. The `:as` marks this association as polymorphic, and specifies which interface we are using on the other side of the association. While we're on the subject, the other end of a polymorphic `belongs_to` can be either a `has_many` or a `has_one` and work identically.

The Database Columns

Here's a migration that will create the `comments` table:

```
class CreateComments < ActiveRecord::Migration
  def self.up
    create_table :comments do |t|
      t.text :body
      t.integer :commentable
      t.string :commentable_type
    end
  end
end
```

As you can see, there is a column called `commentable_type`, which stores the class name of associated object. The Migrations API actually gives you a one-line shortcut with the `references` method, which takes a polymorphic option:

```
create_table :comments do |t|
  t.text :body
  t.references :commentable, :polymorphic => true
end
```

We can see how it comes together using the Rails console (some lines omitted for brevity):

```
>> c = Comment.create(:text => "I could be commenting anything.")
>> t = TimeSheet.create
>> b = BillableWeek.create
>> c.update_attribute(:commentable, t)
=> true
>> "#{c.commentable_type}: #{c.commentable_id}"
=> "Timesheet: 1"
```



```
>> c.update_attribute(:commentable, b)
=> true
>> "#{c.commentable_type}: #{c.commentable_id}"
=> "BillableWeek: 1"
```

As you can tell, both the `Timesheet` and the `BillableWeek` that we played with in the console had the same id (1). Thanks to the `commentable_type` attribute, stored as a string, Rails can figure out which is the correct related object.

has_many :through and Polymorphics

There are some logical limitations that come into play with polymorphic associations. For instance, since it is impossible for Rails to know the tables necessary to join through a polymorphic association, the following hypothetical code, which tries to find everything that the user has commented on, will not work.

```
class Comment < ActiveRecord::Base
  belongs_to :user # author of the comment
  belongs_to :commentable, :polymorphic => true
end

class User < ActiveRecord::Base
  has_many :comments
  has_many :commentables, :through => :comments
end

>> User.first.comments
ActiveRecord::HasManyThroughAssociationPolymorphicError: Cannot have
a has_many :through association 'User#commentables' on the polymorphic
object 'Comment#commentable'.
```

If you really need it, `has_many :through` is possible with polymorphic associations, but only by specifying exactly what type of polymorphic associations you want. To do so, you must use the `:source_type` option. In most cases, you will also need to use the `:source` option, since the association name will not match the interface name used for the polymorphic association:

```
class User < ActiveRecord::Base
  has_many :comments
  has_many :commented_timesheets, :through => :comments,
    :source => :commentable, :source_type => 'Timesheet'
  has_many :commented_billable_weeks, :through => :comments,
    :source => :commentable, :source_type => 'BillableWeek'
end
```

It's verbose, and the whole scheme loses its elegance if you go this route, but it works:

```
>> User.first.commented_timesheets
=> [#<Timesheet ...>]
```

9.8 Foreign-key Constraints

As we work toward the end of this book's coverage of Active Record, you might have noticed that we haven't really touched on a subject of particular importance to many programmers: foreign-key constraints in the database. That's mainly because use of foreign-key constraints simply isn't the Rails way to tackle the problem of relational integrity. To put it mildly, that opinion is controversial and some developers have written off Rails (and its authors) for expressing it.

There really isn't anything stopping you from adding foreign-key constraints to your database tables, although you'd do well to wait until after the bulk of development is done. The exception, of course, is those polymorphic associations, which are probably the most extreme manifestation of the Rails opinion against foreign-key constraints. Unless you're armed for battle, you might not want to broach that particular subject with your DBA.

9.9 Using Value Objects

In Domain Driven Design¹⁰ (DDD), a distinction is drawn between Entity Objects and Value Objects. All model objects that inherit from ActiveRecord::Base could be considered Entity Objects in DDD. An Entity object cares about identity, since each one is unique. In Active Record uniqueness is derived from the primary key. Comparing two different Entity Objects for equality should always return false, even if all of its attributes (other than the primary key) are equivalent.

Here is an example comparing two Active Record Addresses:

```
>> home = Address.create(:city => "Brooklyn", :state => "NY")
>> office = Address.create(:city => "Brooklyn", :state => "NY")
>> home == office
=> false
```

In this case you are actually creating two new Address records and persisting them to the database, therefore they have different primary key values.

Value Objects on the other hand only care that all their attributes are equal. When creating Value Objects for use with Active Record you do not inherit from

10. <http://www.domaindrivendesign.org/>

`ActiveRecord::Base`. Instead you make them part of a parent model using the `composed_of` class method. This is a form of composition, called an Aggregate in DDD. The attributes of the Value Object are stored in the database together with the parent object and `composed_of` provides a means to interact with those values as a single object.

A simple example is of a `Person` with a single `Address`. To model this using composition, first we need a `Person` model with fields for the `Address`. Create it with the following migration:

```
class CreatePeople < ActiveRecord::Migration
  def self.up
    create_table :people do |t|
      t.string :name
      t.string :address_city
      t.string :address_state
    end
  end
end
```

The `Person` model looks like this:

```
class Person < ActiveRecord::Base
  composed_of :address, :mapping => [%w(address_city city),
  %w(address_state state)]
end
```

We'd need a corresponding `Address` object which looks like this:

```
class Address
  attr_reader :city, :state

  def initialize(city, state)
    @city, @state = city, state
  end

  def ==(other_address)
    city == other_address.city && state == other_address.state
  end
end
```

Note that this is just a standard Ruby object that does not inherit from `ActiveRecord::Base`. We have defined reader methods for our attributes and are assigning them upon initialization. We also have to define our own `==` method for use in comparisons. Wrapping this all up we get the following usage:

```
>> gary = Person.create(:name => "Gary")
>> gary.address_city = "Brooklyn"
>> gary.address_state = "NY"
>> gary.address
=> #<Address:0x20bc118 @state="NY", @city="Brooklyn">
```

Alternately you can instantiate the address directly and assign it using the address accessor:

```
>> gary.address = Address.new("Brooklyn", "NY")
>> gary.address
=> #<Address:0x20bc118 @state="NY", @city="Brooklyn">
```

9.9.1 Immutability

It's also important to treat value objects as immutable. Don't allow them to be changed after creation. Instead, create a new object instance with the new value instead. Active Record will not persist value objects that have been changed through means other than the writer method.

The immutable requirement is enforced by Active Record by freezing any object assigned as a value object. Attempting to change it afterwards will result in a `ActiveSupport::FrozenObjectError`.

9.9.2 Custom Constructors and Converters

By default value objects are initialized by calling the `new` constructor of the value class with each of the mapped attributes, in the order specified by the `:mapping` option, as arguments. If for some reason your value class does not work well with that convention, `composed_of` allows a custom constructor to be specified.

When a new value object is assigned to its parent, the default assumption is that the new value is an instance of the value class. Specifying a custom converter allows the new value to be automatically converted to an instance of value class (when needed).

For example, consider the `NetworkResource` model with `network_address` and `cidr_range` attributes that should be contained in a `NetAddr::CIDR` value class.¹¹ The constructor for the value class is called `create` and it expects a CIDR address string as a parameter. New values can be assigned to the value object using either another `NetAddr::CIDR` object, a string or an array. The `:constructor` and `:converter` options are used to meet the requirements:

```
class NetworkResource < ActiveRecord::Base
  composed_of :cidr,
    :class_name => 'NetAddr::CIDR',
    :mapping => [ %w(network_address network), %w(cidr_range
bits) ],
    :allow_nil => true,
```

11. Actual objects from the `NetAddr` gem available at <http://netaddr.rubyforge.org>

```

      :constructor => Proc.new { |network_address, cidr_range|
NetAddr::CIDR.create("#{network_address}/#{cidr_range}") },
      :converter => Proc.new { |value|
NetAddr::CIDR.create(value.is_a?(Array) ? value.join('/') : value) }
end

# This calls the :constructor
network_resource = NetworkResource.new(:network_address => '192.168.0.1',
:cidr_range => 24)

# These assignments will both use the :converter
network_resource.cidr = [ '192.168.2.1', 8 ]
network_resource.cidr = '192.168.0.1/24'

# This assignment won't use the :converter as the value is already an
instance of the value class
network_resource.cidr = NetAddr::CIDR.create('192.168.2.1/8')

# Saving and then reloading will use the :constructor on reload
network_resource.save
network_resource.reload

```

9.9.3 Finding Records by a Value Object

Once a `composed_of` relationship is specified for a model, records can be loaded from the database by specifying an instance of the value object in the conditions hash. The following example finds all customers with `balance_amount` equal to 20 and `balance_currency` equal to "USD":

```
Customer.where(:balance => Money.new(20, "USD"))
```

The Money Gem

A common approach to using `composed_of` is in conjunction with the `money gem`.¹²

```

class Expense < ActiveRecord::Base
  composed_of :cost,
  :class_name => "Money",
  :mapping => [%w(cents cents), %w(currency currency_as_string)],
  :constructor => Proc.new do |cents, currency|
    Money.new(cents || 0, currency || Money.default_currency)
  end
end

```

Remember to add a migration with the 2 columns, the integer `cents` and the string `currency` that money needs.

12. <http://github.com/FooBarWidget/money/>

```
class CreateExpenses < ActiveRecord::Migration
  def self.up
    create_table :expenses do |table|
      table.integer :cents
      table.string :currency
    end
  end
  def self.down
    drop_table :expenses
  end
end
```

Now when asking for or setting the cost of an item would use a `Money` instance.

```
>> expense = Expense.create(:cost => Money.new(1000, "USD"))
>> cost = expense.cost
>> cost.cents
=> 1000
>> expense.currency
=> "USD"
```

9.10 Modules for Reusing Common Behavior

In this section, we'll talk about one strategy for breaking out functionality that is shared between disparate model classes. Instead of using inheritance, we'll put the shared code into modules.

In the section "Polymorphic `has_many` Relationships," we described how to add a commenting feature to our recurring sample `Time` and `Expenses` application. We'll continue fleshing out that example, since it lends itself to factoring out into modules.

The requirements we'll implement are as follows: Both users and approvers should be able to add their comments to a `Timesheet` or `ExpenseReport`. Also, since comments are indicators that a timesheet or expense report requires extra scrutiny or processing time, administrators of the application should be able to easily view a list of recent comments. Human nature being what it is, administrators occasionally gloss over the comments without actually reading them, so the requirements specify that a mechanism should be provided for marking comments as "OK" first by the approver, then by the administrator.

Again, here is the polymorphic `has_many :comments, :as => :commentable` that we used as the foundation for this functionality:

```
class Timesheet < ActiveRecord::Base
  has_many :comments, :as => :commentable
end
```

```
class ExpenseReport < ActiveRecord::Base
  has_many :comments, :as => :commentable
end

class Comment < ActiveRecord::Base
  belongs_to :commentable, :polymorphic => true
end
```

Next we enable the controller and action for the administrator that list the 10 most recent comments with links to the item to which they are attached.

```
class Comment < ActiveRecord::Base
  scope :recent, order('created_at desc').limit(10)
end

class CommentsController < ApplicationController
  before_filter :require_admin, :only => :recent
  expose(:recent_comments) { Comment.recent }
end
```

Here's some of the simple view template used to display the recent comments.

```
%ul.recent.comments
- recent_comments.each do |comment|
  %li.comment
    %h4= comment.created_at
    = comment.text
    .meta
      Comment on:
      = link_to comment.commentable.title, comment.commentable Yes, this
      would result in N+1 selects.
```

So far, so good. The polymorphic association makes it easy to access all types of comments in one listing. In order to find all of the unreviewed comments for an item, we can use a named scope on the Comment class together with the comments association.

```
class Comment < ActiveRecord::Base
  scope :unreviewed, where(:reviewed => false)
end

>> timesheet.comments.unreviewed
```

Both `Timesheet` and `ExpenseReport` currently have identical `has_many` methods for comments. Essentially, they both share a common interface. They're commentable!

To minimize duplication, we could specify common interfaces that share code in Ruby by including a module in each of those classes, where the module contains the code common to all implementations of the common interface. So, mostly for the sake of

example, let's go ahead and define a `Commentable` module to do just that, and include it in our model classes:

```
module Commentable
  has_many :comments, :as => :commentable
end

class Timesheet < ActiveRecord::Base
  include Commentable
end

class ExpenseReport < ActiveRecord::Base
  include Commentable
end
```

Whoops, this code doesn't work! To fix it, we need to understand an essential aspect of the way that Ruby interprets our code dealing with open classes.

9.10.1 A Review of Class Scope and Contexts

In many other interpreted OO programming languages, you have two phases of execution—one in which the interpreter loads the class definitions and says “this is the definition of what I have to work with,” followed by the phase in which it executes the code. This makes it difficult (though not necessarily impossible) to add new methods to a class dynamically during execution.

In contrast, Ruby lets you add methods to a class at any time. In Ruby, when you type `class MyClass`, you're doing more than simply telling the interpreter to define a class; you're telling it to “execute the following code in the scope of this class.”

Let's say you have the following Ruby script:

```
1 class Foo < ActiveRecord::Base
2   has_many :bars
3 end
4 class Foo < ActiveRecord::Base
5   belongs_to :spam
6 end
```

When the interpreter gets to line 1, you are telling it to execute the following code (up to the matching end) in the context of the `Foo` class object. Because the `Foo` class object doesn't exist yet, it goes ahead and creates the class. At line 2, we execute the statement `has_many :bars` in the context of the `Foo` class object. Whatever the `has_many` method does, it does right now.

When we again say `class Foo` at line 4, we are once again telling the interpreter to execute the following code in the context of the `Foo` class object, but this time, the

interpreter already knows about class `Foo`; it doesn't actually create another class. Therefore, on line 5, we are simply telling the interpreter to execute the `belongs_to :spam` statement in the context of that same `Foo` class object.

In order to execute the `has_many` and `belongs_to` statements, those methods need to exist in the context in which they are executed. Because these are defined as class methods in `ActiveRecord::Base`, and we have previously defined class `Foo` as extending `ActiveRecord::Base`, the code will execute without a problem.

However, when we defined our `Commentable` module like this:

```
module Commentable
  has_many :comments, :as => :commentable
end
```

... we get an error when it tries to execute the `has_many` statement. That's because the `has_many` method is not defined in the context of the `Commentable` module object.

Given what we now know about how Ruby is interpreting the code, we now realize that what we really want is for that `has_many` statement to be executed in the context of the including class.

9.10.2 The **included** Callback

Luckily, Ruby's `Module` class defines a handy callback that we can use to do just that. If a `Module` object defines the method `included`, it gets run whenever that module is included in another module or class. The argument passed to this method is the module/class object into which this module is being included.

We can define an `included` method on our `Commentable` module object so that it executes the `has_many` statement in the context of the including class (`Timesheet`, `ExpenseReport`, and so on):

```
module Commentable
  def self.included(base)
    base.class_eval do
      has_many :comments, :as => :commentable
    end
  end
end
```

Now, when we include the `Commentable` module in our model classes, it will execute the `has_many` statement just as if we had typed it into each of those classes' bodies.

The technique is common enough, within Rails and plugins, that it was added as a first-class concept in the Rails 3 ActiveSupport API. The above example becomes shorter and easier to read as a result:

```
module Commentable
  extend ActiveSupport::Concern
  included do
    has_many :comments, :as => :commentable
  end
end
```

Whatever is inside of the `included` block will get executed in the class context of the class where the module is included.

```
has_many :comments, :as => :commentable, :extend => Commentable
```

Courtenay says...

There's a fine balance to strike here. Magic like `include Commentable` certainly saves on typing and makes your model look less complex, but it can also mean that your association code is doing things you don't know about. This can lead to confusion and hours of head-scratching while you track down code in a separate module. My personal preference is to leave all associations in the model, and extend them with a module. That way you can quickly get a list of all associations just by looking at the code.

9.11 Modifying Active Record Classes at Runtime

The metaprogramming capabilities of Ruby, combined with the `after_find` callback, open the door to some interesting possibilities, especially if you're willing to blur your perception of the difference between code and data. I'm talking about modifying the behavior of model classes on the fly, as they're loaded into your application.

Listing 9.5 is a drastically simplified example of the technique, which assumes the presence of a `config` column on your model. During the `after_find` callback, we get a handle to the unique singleton class¹³ of the model instance being loaded. Then we execute the contents of the `config` attribute belonging to this particular `Account` instance, using Ruby's `class_eval` method. Since we're doing this using the singleton class for this instance, rather than the global `Account` class, other account instances in the system are completely unaffected.

13. I don't expect this to make sense to you, unless you are familiar with Ruby's singleton classes, and the ability to evaluate arbitrary strings of Ruby code at runtime. A good place to start is <http://whytheluckystiff.net/articles/seeingMetaClassesClearly.html>.

Listing 9.5 Runtime metaprogramming with `after_find`

```
class Account < ActiveRecord::Base

  ...

  protected

  def after_find
    singleton = class << self; self; end
    singleton.class_eval(config)
  end
end
```

I used powerful techniques like this one in a supply-chain application that I wrote for a large industrial client. A lot is a generic term in the industry used to describe a shipment of product. Depending on the vendor and product involved, the attributes and business logic for a given lot vary quite a bit. Since the set of vendors and products being handled changed on a weekly (sometimes daily) basis, the system needed to be reconfigurable without requiring a production deployment.

Without getting into too much detail, the application allowed the maintenance programmers to easily customize the behavior of the system by manipulating Ruby code stored in the database, associated with whatever product the lot contained.

For example, one of the business rules associated with lots of butter being shipped for Acme Dairy Co. might dictate a strictly integral product code, exactly 10 digits in length. The code, stored in the database, associated with the product entry for Acme Dairy's butter product would therefore contain the following two lines:

```
validates_numericality_of :product_code, :only_integer => true
validates_length_of       :product_code, :is => 10
```

9.11.1 Considerations

A relatively complete description of everything you can do with Ruby metaprogramming, and how to do it correctly, would fill its own book. For instance, you might realize that doing things like executing arbitrary Ruby code straight out of the database is inherently dangerous. That's why I emphasize again that the examples shown here are very simplified. All I want to do is give you a taste of the possibilities.

If you do decide to begin leveraging these kinds of techniques in real-world applications, you'll have to consider security and approval workflow and a host of other important concerns. Instead of allowing arbitrary Ruby code to be executed, you might

feel compelled to limit it to a small subset related to the problem at hand. You might design a compact API, or even delve into authoring a domain-specific language (DSL), crafted specifically for expressing the business rules and behaviors that should be loaded dynamically. Proceeding down the rabbit hole, you might write custom parsers for your DSL that could execute it in different contexts—some for error detection and others for reporting. It's one of those areas where the possibilities are quite limitless.

9.11.2 Ruby and Domain-Specific Languages

My former colleague Jay Fields and I pioneered the mix of Ruby metaprogramming, Rails, and internal¹⁴ domain-specific languages while doing Rails application development for clients. I still occasionally speak at conferences and blog about writing DSLs in Ruby.

Jay has also written and delivered talks about his evolution of Ruby DSL techniques, which he calls Business Natural Languages (or BNL for short¹⁵). When developing BNLs, you craft a domain-specific language that is not necessarily valid Ruby syntax, but is close enough to be transformed easily into Ruby and executed at runtime, as shown in Listing 9.6.

Listing 9.6 Example of business natural language

```
employee John Doe
compensate 500 dollars for each deal closed in the past 30 days
compensate 100 dollars for each active deal that closed more than
365 days ago
compensate 5 percent of gross profits if gross profits are greater than
1,000,000 dollars
compensate 3 percent of gross profits if gross profits are greater than
2,000,000 dollars
compensate 1 percent of gross profits if gross profits are greater than
3,000,000 dollars
```

The ability to leverage advanced techniques such as DSLs is yet another powerful tool in the hands of experienced Rails developers.

14. The qualifier `internal` is used to differentiate a domain-specific language hosted entirely inside of a general-purpose language, such as Ruby, from one that is completely custom and requires its own parser implementation.

15. Googling BNL will give you tons of links to the Toronto-based band Barenaked Ladies, so you're better off going directly to the source at <http://bnl.jayfields.com>.

9.12 Conclusion

With this chapter we conclude our coverage of Active Record. Among other things, we examined how callbacks and observers let us factor our code in a clean and object-oriented fashion. We also expanded our modeling options by considering single-table inheritance, abstract classes and Active Record's distinctive polymorphic relationships.

At this point in the book, we've covered two parts of the MVC pattern: the model and the controller. It's now time to delve into the third and final part: the view.

Courtenay says...

DSLs suck! Except the ones written by Obie, of course. The only people who can read and write most DSLs are their original authors. As a developer taking over a project, it's often quicker to just reimplement instead of learning the quirks and exactly which words you're allowed to use in an existing DSL. In fact, a lot of Ruby metaprogramming sucks, too. It's common for people gifted with these new tools to go a bit overboard. I consider metaprogramming, `self.included`, `class_eval`, and friends to be a bit of a code smell on most projects. If you're making a web application, future developers and maintainers of the project will appreciate your using simple, direct, granular, and well-tested methods, rather than monkeypatching into existing classes, or hiding associations in modules. That said, if you can pull it off... your code will become more powerful than you can possibly imagine.

Index

A

Action Controller

- communication with view, 104B–105B
- controller specs, 526B–529B
- filters, 105B–111B
 - around, 109B
 - classes 107B
 - conditions, 110B
 - halting chain, 111B
 - inheritance, 106B
 - ordering, 108B
 - skipping, 110B
- layouts, specifying, 101B
- post-backs, 357B
- rendering, 92B–101B
- standard RESTful actions, 61B–64B
- streaming content, 112B–116B
- verify method, 111B–112B

Action Dispatch, 88B–91B

Action Mailer, 471B–481B

- attachments, 475B, 476B, 478B
- custom email headers, 473B
- generating URLs inside messages, 476B
- handing inbound attachments, 478B
- HTML messages, 474B

- mailer layouts, 476B
- models, 472B
- multipart messages, 475B
- preparing outbound messages, 472B
- raising delivery errors, 19B
- receiving, 477B–479B
- sending, 477B
- server configuration, 479B
- SMTP, 471B
- testing with RSpec, 479B–481B

Action View, 293B–308B

- conditional output, 296B
- customizing validation error output, 315B
- ERb, 293B
- filename conventions, 294B
- flash messages, 300B–302B
- Haml, 293B
- instance variables, 297B–302B
- layouts, 294B–295B
- partials, see *Partials*.
- view specs, 529B–531B
- yielding content, 295B

Active Model, 561B–578B

- AttributeMethods module, 561B–562B
- Callbacks module, 563B

- Conversion module, 563B–564B
- Dirty module, 564B
- Errors class, 565B
- MassAssignmentSecurity module, 567B
- Naming module, 569B
- observers, 569B–571B
- serialization, 571B–573B
- testing compatibility of custom classes with
 - Lint::Tests, 567B
- translation, 573B
- Validations module, 574B–578B
- Active Resource 459B–471B
 - authentication, 465B
 - customizing default URLs, 463B
- Active Record
 - abstract base models, 276B
 - associations, 121B, 181B–230B
 - :counter_cache option, 190B, 195B
 - :counter_sql option, 187B
 - :dependent option, 187B, 188B, 195B
 - :finder_sql option, 187B
 - AssociationProxy class, 228B–229B
 - belongs_to. See belongs_to associationsB
 - checking inclusion of records in
 - collection, 189B
 - class hierarchy, 181B
 - destroying records, 188B
 - extensions, 226B–227B
 - foreign-key constraints, 281B
 - has_and_belongs_to_many. See has_and_belongs_to_many associations
 - has_many :through. See has_many :through associations
 - has_many. See has_many associations
 - indexing, 484B
 - many-to-many relationships, 208B–214B
 - one-to-many relationships, 183B–190B
 - one-to-one relationships, 222B–225B
 - polymorphic, 277B–281B
 - size of, 190B
 - unique sets, 190B
 - unsaved objects, 225B
- attributes, 123B–126B
 - controlling access, 140B
 - readonly, 141B
 - reloading, 131B
 - serialized, 125B
 - typecasting, 131B
 - updating, 136B–1140B
- Base class, 120B
- basic object operations, 127B–133B
- calculation methods, 265B–267B
- callbacks, 256B–265B
 - list of, 258B–259B
- cloning, 131B
- concurrency. See Database locking
- configuration, 158B
- dynamic finder methods, 132B
- dynamic scopes, 133B
- find_by_sql method, 133B–134B
- legacy naming schemes, 122B–123B
- model specs, 526B–528B
- migrations, 161B–179B
 - column type mappings, 168B–172B
 - creating, 161B–172B
 - magic timestamp columns, 172B
 - schema.rb file, 174B
 - sequencing, 162B
- observers, 10B, 268B–269B
- pattern, 119B
- query caching, 135B–136B
- querying, 146B–152B
 - exists, 152B
 - from, 150B
 - group, 150B
 - having, 150B
 - includes, 151B
 - joins, 151B

- limit, 149B
 - offset, 149B
 - order, 148B
 - readonly, 152B
 - select, 149B
 - where, 146B–148B
 - RecordInvalid exception, 187B
 - RecordNotSaved exception, 187B, 188B
 - records
 - deleting, 141B–142B
 - random ordering, 148B
 - touching, 139B
 - scopes, 251B–255B
 - session store, 429B
 - STI (Single-Table Inheritance), 269B–276B
 - translations, 386B–388, 390B
 - validations, 231B–250B
 - common options, 242B–243B
 - conditional validation, 243B–245B
 - contexts, 245B
 - custom macros, 247B–248B
 - errors, 231B–232B, 249B–250B
 - enforcing uniqueness of join models, 240B
 - reporting, 310B–312B
 - short-form, 245B–246B
 - skipping, 249B
 - testing with Shoulda, 250B
 - value objects, 281B–285B
 - Active Support, 579B–686B
 - Ajax, 409B–425B
 - changes in Rails 3, 410B
 - CSS selectors, 418B
 - HTML fragments, 421 B
 - JSON, 419B–421B
 - JSONP, 423B–424B
 - Unobtrusive JavaScript (UJS), 411B–412B
 - Array, extensions, 579B–585B
 - Asset hosts, 22B, 321B–323B
 - Asset timestamps, 323B
 - Asynchronous processing, See Background processing.
 - Atom Feeds
 - autodetection, 316B–317B
 - atom_feed method, 324B–326B
 - Authentication, 433B
 - Active Resource, 465B
 - client-side certificates, 466B
 - HTTP basic, 465B
 - HTTP digest, 466B
 - Authlogic, 434B
 - configuration, 437B–438B
- ## B
- background processing, 549B–559B
 - Base64 class, 586B
 - BasicObject class, 586B–587B
 - belongs_to associations, 191B–199B
 - building and creating related objects, 192B
 - options, 192B–199B
 - polymorphic, 197B
 - reloading, 191B
 - touch, 198B
 - with conditions, 193B
 - benchmarking, 587B
 - binary data storage, 170B
 - breadcrumbs, 400B–401B
 - Builder::XmlMarkup class, 454B
 - Bundler, 2B–7B
 - loading gems directly from Git repository, 4B–5B
- ## C
- Caching
 - :counter_cache, 190B, 195B
 - action caching, 484B–485B
 - CacheHelper module, 326B
 - controlling web caches and proxies, 497B
 - disabling in development mode, 18B
 - ETags, 498B–500B

expiration, 488B–491B
 fetch, 496B–497B
 fragment caching, 486B–488B
 page caching, 484B
 query caching, 135B–136B
 storage, 493B
 Store class, 590B–595B
 sweeping, 491B, 494B
 view caching, 483B–495B
 Callbacks. See also Active Record, callbacks
 in Active Support, 595B–597B
 CAS, 443B
 CDATA, 366B
 chars proxy, 645B–648B
 Class, extensions, 598B
 Concern module, 602B
 Concurrency. See Database Locking
 Configurable module, 603B
 const_missing method, 644B
 Controllers. See Action Controller
 convention over configuration, 119B, 122B
 Cookies
 :secure option, 432B
 integrity, 11B–12B
 reading and writing, 431B
 session store, 429B
 signing, 432B
 CRUD (Create Read Update Delete), 119B
 CSS
 linking stylesheets to template, 318B–319B
 sanitizing, 365B
 Currency
 formatting, 359B
 Money gem, 284B
D
 data migration, 173B–174B
 Databases
 connecting to multiple, 153B–154B
 foreign-key constraints, 281B

locking, 142B–146B
 considerations, 145B
 optimistic, 143B
 pessimistic, 145B
 migrations. See ActiveRecord, Migrations.
 schemas, 15B, 161B
 seeding, 175B–76B
 using directly, 154B–158B
 Date, extensions, 603B–609B
 Date input tags, 328B–331B
 DateTime, extensions, 609B
 Decent Exposure gem, 105B, 297B–298B
 decimal precision, 169B, 171B
 Delayed Job gem, 550B–553B
 Deprecation, 617B
 Devise gem, 439B
 Domain-Specific Languages, 291B
 Drag and Drop, 415B
 Duration class, 617B–618B

E

Email. See Action Mailer
 Enumerable, extensions, 619B–620B
 ETags, 498B–500B
 Excerpting text, 370B

F

Facebook, 443B
 favicon.ico file, 317B
 Files
 extensions by Active Support, 621B–622B
 reporting sizes to users, 359B
 upload field, 348B, 356B
 Firebug, 410B
 floats, 171B
 Forms
 _destroy checkbox, 345B
 _method hidden field, 64B
 accepts_nested_attributes_for method,
 344B–345B

- attributes not typecasted, 343B
- automatic view creation, 313B–314B
- button_to helper method, 391B
- custom builder classes, 347B
- dynamically adding rows of child records, 338B
- helper methods, 333B–358B
 - input, 348B–358B
- updating multiple objects at once, 337B

G

- Gemfile, 3B
- Geocoding, 260B–261B

H

- has_and_belongs_to_many associations, 208B–214B
 - bidirectional, 210B
 - custom SQL, 211B–213B
 - extra columns, 213B
 - making self-referential, 209B
- has_many :through associations, 214B–221B
 - and validations, 218B
 - join models, 215B
 - options, 219B–221B
 - usage, 216B
- has_many associations, 199B–208B
 - :class_name option, 202B
 - :conditions option, 202B
 - :include option, 204B–206B
 - callbacks, 200B–201B
- has_one associations, 222B–225B
 - options, 224B–225B
 - together with has_many, 223B
- Hash, extensions, 622B–627B
- HashWithIndifferentAccess class, 627B
- Heckle, 532B
- Helper methods, B
 - helper specs, 531B
 - writing your own 398B–407B

HTML

- escaping, 367B
- sanitizing, 364B–365B
- tags
 - a, 392B–394B
 - audio, 319B
 - label, 348B
 - form. See Forms.
 - option, 353B–355B
 - password, 349B
 - select, 350B–351B
 - script, 359B
 - submit, 349B
 - video, 320B

HTTP

- basic authentication, 465B
- foundation of REST, 55B–56B
- stateless, 427B
- status codes, 99B–101B
- verbs (GET, POST, etc.), 60B–64, 393B

I

- IMAP, 443B
- Image tags, 320B
- Initializers, 11B–14B
 - backtrace_silencers.rb, 11B
 - cookie_verification_secret.rb, 11B–12B
 - inflections.rb, 12B–13B
 - mime_types.rb, 14B
 - session_store.rb, 14B, 430B
- Inflections. See Pluralization
- Integer, extensions, 632B–633B
- Internationalization (I18n), 372B–391B
 - Active Model, 573BB
 - Active Record, 386B
 - default locale, 10B
 - exception handling, 391B
 - interpolation, 385B
 - locale files, 382B–383B
 - process, 380B–390B

setting user locales, 377B–380B
 setup, 374B–380B

J

JavaScript, 97B, 317B–318B, 358B–259B,
 409B–425B
 escaping, 358B
 including in template, 317B–318B
 link_to method enhancements, 392B–393B
 using to insert HTML into pages, 338B
 JQuery framework, 410B–411B, 418B, 421B,
 423B–424B
 JSON, 97B, 419B–421B, 633B–634B
 JSONP, 423B–424B

K

Kernel, extensions, 634B–635B

L

LDAP, 443B
 link_to helper methods, 392B–394B
 Locale files, 382B–383B
 Logging, 23B–28B
 backtrace silencing, 11B, 585B
 BufferedLogger class, 588B–590B
 colorization, 27B
 level override, 15B
 levels, 23B–24B
 log file analysis, 26B–27B
 Logger, extensions, 635B–637B
 Syslog, 28B

M

Memcache, session store, 428B
 MessageEncryptor class, 636B–637B
 MessageVerifier class, 637B–638B
 Middleware (Rack), 86B–88B
 MIME types, 13B–14B
 Module, extensions, 638B
 MongoDB, 442B
 MVC (Model-View-Controller), 85B

N

Named scopes. See Active Record, scopes.
 Nonces, 430B
 Notifications, 651B–652B
 Numbers
 delimiters, 360B
 extensions to Numeric class, 650B–653B
 conversion, 359B–361B

O

Object, extensions, 653B
 Observers, 10B
 OpenID, 443B
 OpenSSL Digests, 431B

P

params hash, 336B–338B
 Partial, 95B, 302B–307B
 passing variables to, 305B
 rendering collections, 306B
 reuse, 303B
 shared, 304B
 wrapping and generalizing, 401B–407B
 Plugins, 535B–548B
 as RubyGems, 536B
 extending Rails classes, 540B
 installation and removal, 542B–543B
 load order, 9B–10B
 plugin script, 536B
 rake tasks, 543B–544B
 testing, 545B–546B
 writing your own, 537B–548B
 Pluralization
 i18n, 385B
 Inflector class, 12B–13B
 Inflections class, 628B–632B
 pluralize helper method, 370B
 Prototype framework, 361B, 411B
 helper methods, 361B
 Prototype Legacy Helper plugin, 413B

R

Rack, 86B–88B, 90B–91B

 Rack::Sendfile middleware, 114B

 RACK_ENV variable, 1B

 routes as Rack endpoints, 41B–42B

rails.js file, 411B

Rails

 Class loader and reloading, 16B–18B,
 613B–617B

 console, 12B

 reloading, 92B

 Engines, 549B

 lib directory, 17B

 root directory, 9B

 runner, 559B

 scaffolding, 311B–314B

 settings, 1B–29B

 application.rb file, 8B–11B

 autoload paths, 9B

 boot.rb file, 8B

 cherry-picking frameworks used, 8B

 custom environments, 20B

 development mode, 15B–19B

 environment.rb file, 8B

 generator defaults, 11B

 initializers. See Initializers

 production mode, 20B–23B

 test mode, 19B–20B

RAILS_ENV variable, 1B

Railties, 546B–547B, 658B

Rake tasks (selected),

 db:migrate, 163B, 176B

 log:clear 24B

 routes, 53B

 spec, 521B

Random

 ordering of records, 148B

 SecureRandom generator class, 661B–662B

Range, extensions, 658B–659B

RecordNotFound exception, 128B

Regexp, extensions, 660B

Rendering views, 92B–101B

 another actions's template, 93B

 explicit, 93B, 94B

 implicit, 92B–93B

 inline templates, 96B

 JSON, 97B

 nothing, 97B

 options, 98B

 partials, see Partials.

 text, 96B

 XML, 97B

Request handling

 in routing, 89B

 redirecting, 101B–104B, 418B

 verification, 111B–112B

Rescuable module, 660B

Resque gem, 553B–557B

REST and RESTful design, 55B–83B

 action set, 78B–82B

 collection routes, 72B

 controller-only resources, 74B

 forms, 335B

 member routes, 70B–71B

 HTTP verbs, 60B–64B

 nested resources, 65B–69B

 routes, 31, 58B–61B

 resources and representations, 40B–41B,
 56B–57B, 76B–77B

 singular resource routes, 64B–65B

 standard controller actions, 61B–64B

REXML, 456B, 684B

Roy T. Fielding. See also REST and RESTful
 design, 55B–58B

Routing, 31B–54B

 constraining by request method, 38B–39B

 formats, 40B

 globbing, 45B–46B

- listing, 53B
 - match method, 34B–37B
 - named, 46B–50B
 - RESTful routes, 31B, 58B–61B
 - :format parameter, 76B
 - collection, 72B
 - member, 70B–71B
 - nested, 65B–69B
 - singular, 64B–65B
 - redirecting, 39B–40B
 - root routes, 44B–45B
 - routes.rb file, 33B–34B
 - scopes, 50B–53B
 - RJS, 412B–419B
 - templates, 413B
 - RSS autodetection, 316B–317B
 - RPX authentication, 443B
 - RSpactor, 531B
 - RSpec, 501B–533B
 - assertions, 510B
 - custom expectation matchers, 514B
 - generator settings, 11B
 - grouping related examples, 504B
 - let methods, 504B–506B
 - mocking and stubbing, 517B–520B
 - pending, 509B–510B
 - predicate matchers, 513B–514B
 - running specs, 520B
 - runtime options, 522B
 - shared behaviors, 517B
 - spec_helper.rb file, 515B, 522B–524B
 - subjects, 511B–513B
 - testing email, 449B–481B
 - Ruby
 - \$LOAD_PATH, 9B, 16B
 - hashes, 450B
 - macro-style methods, 121B
 - Marshal API, 425B
 - modules for reusing common behavior, 285B–289B
 - RubyGems
 - as plugins, 536B
 - Bundler, 2B–7B
 - Git repository, loading directly from, 4B
 - installing, 5B–7B
 - packaging, 7B
 - using pre-release gems, 4B
- ## S
- Scopes, see Active Record, scopes
 - Security
 - CSRF attacks, 336B
 - replay attacks, 429B
 - SQL injection, 134B
 - Session Management, 425B–432B
 - cleaning old sessions, 430B
 - storage
 - RESTful considerations, 75B
 - turning off sessions, 427B
 - Settings, 1B–29B
 - Specjour, 532B
 - Spork, 532B
 - SSL
 - OpenSSL digests, 429B
 - serving protected assets, 323B
 - X.509 certificates, 465B–466B
 - Static content, 116B
 - Streaming, 112B–116B
 - String,
 - extensions, 662B–671B
 - usage versus symbols, 130B
 - StringInquirer class, 671B
 - SOAP, 457B
 - Symbol,
 - extensions, 671B
 - usage versus strings, 130B

T

Templates. See View templates.

Thread safety, 22B–23B

Time

extensions, 673B–680B

input tags, 328B–331B

reporting distances in time, 332B–333B

storing in database, 170B

Time Zones

DateTime conversions, 612B

default, 10B

option tags helper, 354B–355B

TimeZone class, 681B

TimeWithZone class, 680B

Truncating text, 372B

U

Unicode, 364B, 385B, 646B–647B

Unobtrusive JavaScript (UJS), 411B–412B

URL

generation, 395B–398B

patterns in routing, 35B–36B

segment keys, 36B–38B

V

Validation. See Active Record, validations

Value objects, 281B–285B

View templates. See also Action View,
293B–308B

capturing block content, 326B–327B

concat method, 368B

cycling content, 369B

debugging output, 333B

encapsulating logic in helper methods,
399B

highlighting content, 370B

localization, 373B

raw output, 361B

record identification, 362B–364B

transforming text into HTML, 371B

translation. See Internationalization.

word wrap, 372B

Visual effects, 419B

W

Watchr, 532B

Web 2.0, 332B, 423B

Web architecture, 55B–56B

Whiny nils, 18B

X

XML, 445B–457B

parsing, 456B

to_xml method, 445B–454B

Active Record associations, 448B

customizing output, 446B–448B

extra elements, 452B

overriding, 453B–455B

Ruby hashes, 450B

typecasting, 457B

XML Builder, 454B–456B

XMLHttpRequestObject, 409B

XMLMini module, 684B–686B

Y

YAML, 125B–126B, 445B