A Brain-Friendly Guide

Head First SQL

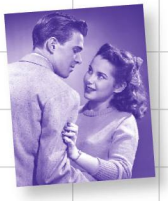

Help Greg improve his data relationships

Stop misplacing your primary and foreign keys

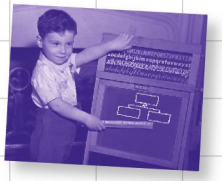

Finally be able to explain what's normal

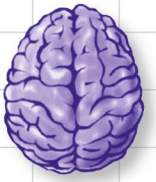

Load important SQL query concepts directly into your brain

Avoid embarrassing ALTER scenarios

Put your SQL knowledge to the test with dozens of exercises

O'REILLY®

Lynn Beighley

Head First SQL

Database/SQL

What will you learn from this book?

In today's world, data is power, but the real secret to success is having power over your data. *Head First SQL* takes you to the heart of the SQL language, from basic syntax queries using INSERT and SELECT to hardcore database manipulation with subqueries, joins, and transactions. By the time you finish reading, you will not only understand effective and efficient database design and creation, you'll be querying, normalizing, and joining data like a pro. And you'll be a true master of your data.

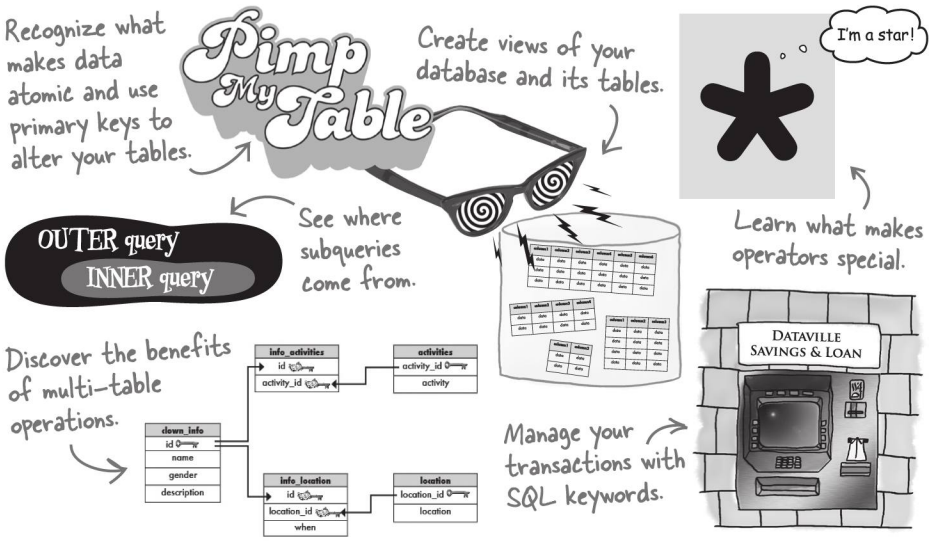

Why does this book look so different?

We think your time is too valuable to spend struggling with new concepts. Using the latest research in cognitive science and learning theory to craft a multi-sensory learning experience, *Head First SQL* uses a visually rich format designed for the way your brain works, not a text-heavy approach that puts you to sleep.

US $44.99 **CAN $53.99**

ISBN-10: 0-596-52684-9

ISBN-13: 978-0-596-52684-9

9 780596 526849 54499

O'REILLY®

www.oreilly.com

www.headfirstlabs.com

"This is not SQL made easy; this is SQL made challenging, SQL made interesting, SQL made fun. It even answers that age-old question 'How to teach non-correlated subqueries without losing the will to live?' This is the right way to learn—it's fast, it's flippant, and it looks fabulous."

—Andrew Cumming, Author of SQL Hacks, *Zoo Keeper at sqlzoo.net*

"There are books you buy, books you keep, books you keep on your desk, and thanks to O'Reilly and the Head First crew, there is the penultimate category, Head First books. They're the ones that are dog-eared, mangled, and carried everywhere. *Head First SQL* is at the top of my stack."

— Bill Sawyer, ATG Curriculum Manager, Oracle

Praise for *Head First SQL*

"There are books you buy, books you keep, books you keep on your desk, and thanks to O'Reilly and the Head First crew, there is the ultimate category, Head First books. They're the ones that are dog-eared, mangled, and carried everywhere. Head First SQL is at the top of my stack. Heck, even the PDF I have for review is tattered and torn."

— Bill Sawyer, ATG Curriculum Manager, Oracle

"This is not SQL made easy; this is SQL made challenging, SQL made interesting, SQL made fun. It even answers that age-old question 'How to teach non-correlated subqueries without losing the will to live?' This is the right way to learn—it's fast, it's flippant, and it looks fabulous."

— Andrew Cumming, Author of *SQL Hacks*, Zoo Keeper at *sqlzoo.net*

"Outrageous! I mean, SQL is a *computer* language, right? So books about SQL should be written for *computers*, shouldn't they? *Head First SQL* is *obviously* written for *human beings*! What's up with *that*?!"

— Dan Tow, Author of *SQL Tuning*

"Even many of the more advanced concepts are covered in a way that almost anyone should be able to easily follow...if your database does ever grow to the point where you need to use the more advanced commands this same book will show you how. The book not only teaches you how to code SQL, it also teaches many of the concepts of proper database design."

— Stephen Chapman, Ask Felgall (*www.felgall.com/hfsql.htm*)

"Jammed with exercises, thoughtful questions, cartoons and side comments that make you giggle while making an important point, this book provides the most enjoyable way to learn SQL than I can imagine—well, short of learning it pool side in the Bahamas anyway. It even includes crossword puzzles to help the reader quiz himself on what he's learned. If you want to have a fun time learning SQL, this is the way to do it. It gets down to hardcore SQL and keeps you thinking not just about what you're doing, but why you're doing it with its lively examples. This is gain without the pain. This is a fun way to learn.

— Sandra Henry-Stocker, *ITworld.com*

Praise for other *Head First books*

"This book's admirable clarity, humor and substantial doses of clever make it the sort of book that helps even non-programmers think well about problem-solving."

— Cory Doctorow, co-editor of Boing Boing
Author, *Down and Out in the Magic Kingdom*
and *Someone Comes to Town, Someone Leaves Town*

"If you thought Ajax was rocket science, this book is for you. Head Rush Ajax puts dynamic, compelling experiences within reach for every web developer."

— Jesse James Garrett, Adaptive Path

"I received the book yesterday and started to read it...and I couldn't stop. This is definitely très 'cool.' It is fun, but they cover a lot of ground and they are right to the point. I'm really impressed."

— Erich Gamma, IBM Distinguished Engineer, and co-author of Design Patterns

"*Head First Design Patterns* managed to mix fun, belly-laughs, insight, technical depth and great practical advice in one entertaining and thought provoking read. Whether you are new to design patterns, or have been using them for years, you are sure to get something from visiting Objectville."

— Richard Helm, co-author of *Design Patterns*

"One of the funniest and smartest books on software design I've ever read."

— Aaron LaBerge, VP Technology, ESPN.com

"I just finished reading HF OOA&D and I loved it! The thing I liked most about this book was its focus on why we do OOA&D—to write great software!"

— Kyle Brown, Distinguished Engineer, IBM

I *heart* *Head First HTML with CSS & XHTML*—it teaches you everything you need to learn in a 'fun coated' format!"

— Sally Applin, UI Designer and Fine Artist, *http://sally.com*

Praise for the *Head First* Approach

"It's fast, irreverant, fun, and engaging. Be careful—you might actually learn something!"

— Ken Arnold, former Senior Engineer at Sun Microsystems
Co-author (with James Gosling of Java),
The Java Programming Language

"I feel like a thousand pounds of books have just been lifted off of my head."

— Ward Cunningham, inventor of the Wiki
and founder of the Hillside Group

"This book is close to perfect, because of the way it combines expertise and readability. It speaks with authority and it reads beautifully."

— David Gelernter, Professor of Computer Science, Yale University

"Just the right tone for the geeked-out, casual-cool guru coder in all of us. The right reference for practical development strategies--gets my brain going without having to slog through a bunch of tired, stale professor-speak."

— Travis Kalanick, Founder of Scour and Red Swoosh Member of the
MIT TR100

"The combination of humour, pictures, asides, sidebars, and redundancy with a logical approach to introducing the basic tags and substantial examples of how to use them will hopefully have the readers hooked in such a way that they don't even realize they are learning because they are having so much fun."

— Stephen Chapman, Fellgall.com

Other related books from O'Reilly

The Art of SQL

Learning SQL

SQL in a Nutshell

SQL Cookbook

SQL Hacks

SQL Pocket Guide

The Relational Database Dictionary

Database in Depth

Other books in O'Reilly's *Head First* series

Head First PMP

Head First Object-Oriented Analysis and Design

Head Rush Ajax

Head First HTML with CSS and XHTML

Head First Design Patterns

Head First Servlets & JSP™

Head First Java™

Head First EJB™

Head First JavaScript (2007)

Head First C# (2007)

Head First SQL

Lynn Beighley

O'REILLY®

Beijing • Cambridge • Köln • Sebastopol • Tokyo

Head First SQL

by Lynn Beighley

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly Media books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*safari.oreilly.com*). For more information, contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

Series Creators: Kathy Sierra, Bert Bates

Series Editor: Brett D. McLaughlin

Editor: Catherine Nolan

Design Editor: Louise Barr

Cover Designers: Louise Barr, Karen Montgomery

Production Editor: Sanders Kleinfeld

Indexer: Julie Hawks

Page Viewer: Andrew Fader

Printing History:

August 2007: First Edition.

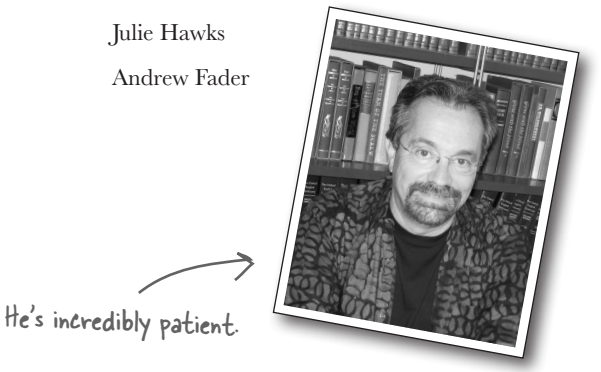

No clowns, doughnuts, or Girl Sprouts were harmed in the making of this book. Just my car, but it's been fixed.

RepKover™ This book uses RepKover™, a durable and flexible lay-flat binding.

ISBN: 978-0-596-52684-9

[M] [2011-02-25]

To our world, awash in data.
And to you, who want to master it.

Author of Head First SQL

Lynn Beighley

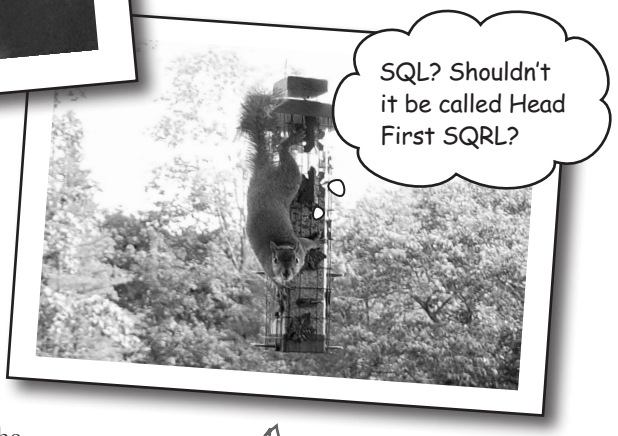

The view from Lynn's window.

Lynn is a fiction writer stuck in a technical book writer's body. Upon discovering that technical book writing actually paid real money, she learned to accept and enjoy it.

After going back to school to get a Masters in computer science, she worked for the acronyms NRL and LANL. Then she discovered Flash, and wrote her first bestseller.

A victim of bad timing, she moved to Silicon Valley just before the great crash. She spent several years working for Yahoo! and writing other books and training courses. Finally giving in to her creative writing bent, she moved to the New York area to get an MFA in creative writing.

Her Head First–style thesis was delivered to a packed room of professors and fellow students. It was extremely well received, and she finished her degree, finished *Head First SQL*, and can't wait to begin her next book.

Lynn loves traveling, cooking, and making up elaborate background stories about complete strangers. She's a little scared of clowns.

Table of Contents (Summary)

Table of Contents (the real thing)

Intro

Your brain on SQL. Here *you* are trying to *learn* something, while here your *brain* is doing you a favor by making sure the learning doesn't *stick*. Your brain's thinking, "Better leave room for more important things, like which wild animals to avoid and whether naked snowboarding is a bad idea." So how *do* you trick your brain into thinking that your life depends on knowing SQL?

data and tables

1 A place for everything

Don't you just hate losing things? Whether it's your car keys, that 25% off coupon for Urban Outfitters, or your application's data, there's nothing worse than not being able to **keep up with what you need**... when you need it. And when it comes to your applications, there's no better place to store your important information than in a **table**. So turn the page, come on in, and take a walk through the world of **relational databases**.

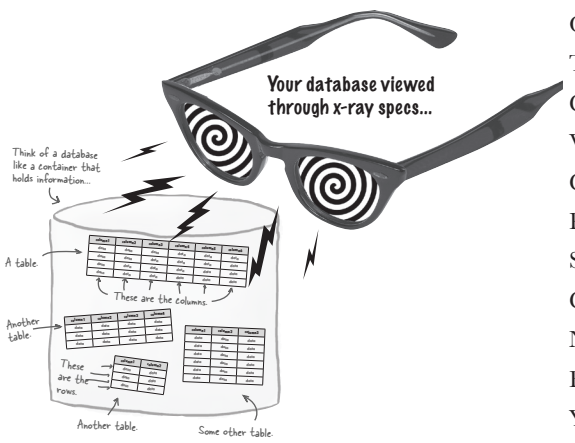

2 the SELECT statement

Gifted data retrieval

Is it really better to give than retrieve? When it comes to databases, chances are you'll need to **retrieve your data** as often than you'll need to insert it. That's where this chapter comes in: you'll meet the powerful **SELECT** statement and learn how to **gain access to that important information** you've been putting in your tables. You'll even learn how to use **WHERE**, **AND**, and **OR** to selectively get to your data and even avoid displaying the data that you *don't* need.

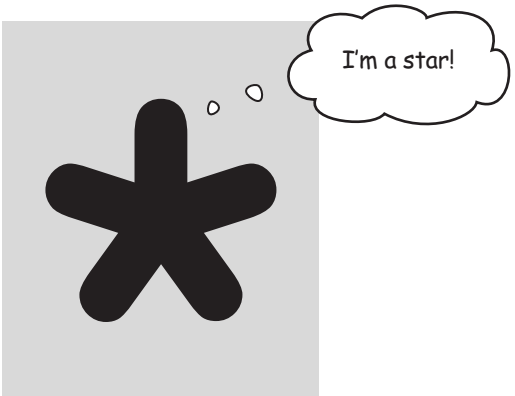

3 DELETE and UPDATE

A change will do you good

Keep changing your mind? Now it's OK! With the commands you're about to learn—**DELETE** and **UPDATE**—you're no longer stuck with a decision you made six months ago, when you first inserted that data about mullets coming back into style soon. With UPDATE, you **can change data**, and DELETE lets you **get rid of data** that you don't need anymore. But we're not just giving you the tools; in this chapter, you'll learn how to be selective with your new powers and avoid dumping data that you really do need.

smart table design

4 Why be normal?

You've been creating tables without giving much thought to them. And that's fine, they work. You can `SELECT`, `INSERT`, `DELETE`, and `UPDATE` with them. But as you **get more data**, you start seeing **things you wish you'd done** to make your `WHERE` clauses simpler. What you need is to make your tables more *normal*.

5 ALTER
Rewriting the Past

Ever wished you could correct the mistakes of your past?

Well, now is your chance. By using the **ALTER command**, you can apply all the lessons you've been learning to tables you designed days, months, even years ago. Even better, you can do it without affecting your data. By the time you're through here, you'll know what **normal** really means, and you'll be able to apply it to all your tables, past and present.

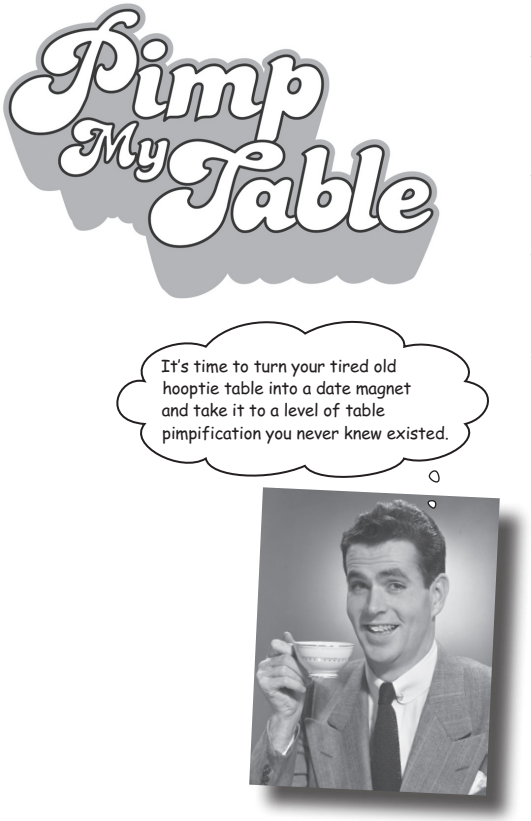

6 advanced SELECT

Seeing your data with new eyes

It's time to add a little finesse to your toolbox. You already know how to SELECT data and use WHERE clauses. But sometimes you need more **precision** than SELECT and WHERE provide. In this chapter, you'll learn about how to **order and group** your data, as well as how to **perform math operations** on your results.

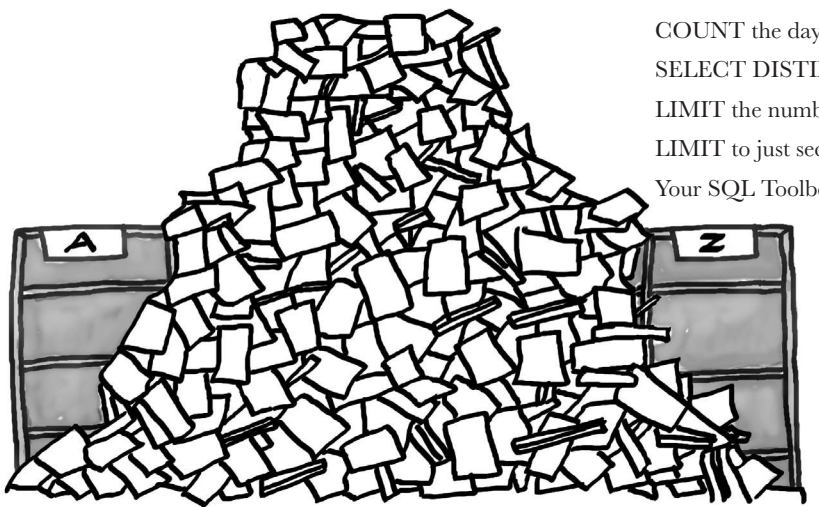

multi-table database design

7 Outgrowing your table

Sometimes your single table isn't big enough anymore.

Your data has become more complex, and that **one table** you've been using just **isn't cutting it**. Your single table is full of redundant data, wasting space and slowing down your queries. You've gone as far as you can go with a single table. It's a big world out there, and sometimes you need **more than one table** to contain your data, control it, and ultimately, be the master of your own database.

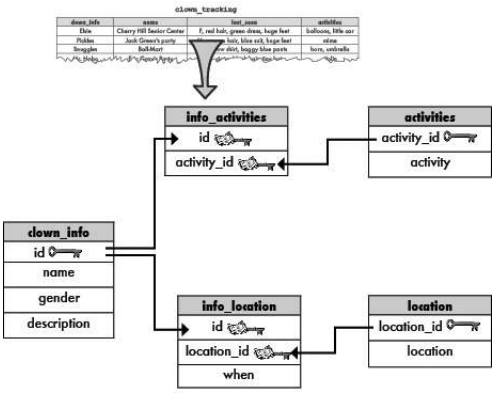

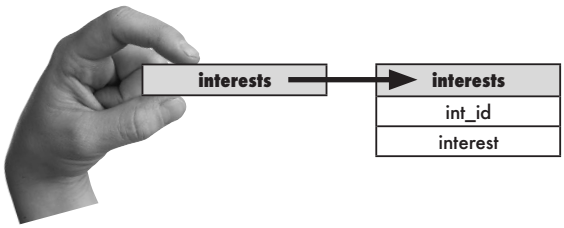

8 joins and multi-table operations

Can't we all just get along?

Welcome to a multi-table world. It's great to have **more than one table** in your database, but you'll need to learn some ***new tools and techniques*** to work with them. With multiple tables comes confusion, so you'll need **aliases** to keep your tables straight. And **joins** help you connect your tables, so that you can get at all the data you've spread out. Get ready, it's time to **take control of your database** again.

subqueries

9 Queries within queries

Yes, Jack, I'd like a two-part question, please. Joins are great, but sometimes you need to ***ask your database more than one question***. Or ***take the result of one query and use it as the input to another query***. That's where **subqueries** come in. They'll help you **avoid duplicate data**, **make your queries more dynamic**, and even get you in to all those high-end concert afterparties. (Well, not really, but two out of three ain't bad!)