

Search Engine Optimization SECRETS

DO WHAT YOU NEVER THOUGHT POSSIBLE WITH SEO



DANNY DOVER

with insights from ERIK DAFFORN

- + Insider Tips
Expert advice to help you optimize performance
- + Valuable Insights
Eye-opening descriptions of poorly documented or undocumented features and procedures
- + Unbeatable Advice
Real-world workarounds and little-known tips and techniques

Search Engine
Optimization

SECRETS

Search Engine Optimization **SECRETS**

DO WHAT YOU NEVER THOUGHT POSSIBLE WITH SEO

Danny Dover

with insights from Erik Dafforn



Wiley Publishing, Inc.

EXECUTIVE EDITOR: Carol Long
ACQUISITIONS EDITOR: Mary James
SENIOR PROJECT EDITOR: Kevin Kent
PROJECT EDITOR: Kristin Vorce
TECHNICAL EDITOR: Tim Buck
PRODUCTION EDITOR: Rebecca Anderson
COPY EDITOR: Kim Cofer
EDITORIAL DIRECTOR: Robyn B. Siesky
EDITORIAL MANAGER: Mary Beth Wakefield
MARKETING MANAGER: Ashley Zurcher
PRODUCTION MANAGER: Tim Tate
VICE PRESIDENT AND EXECUTIVE GROUP PUBLISHER: Richard Swadley
VICE PRESIDENT AND EXECUTIVE PUBLISHER: Barry Pruett
ASSOCIATE PUBLISHER: Jim Minatel
PROJECT COORDINATOR, COVER: Katie Crocker
COMPOSITOR: Craig Woods, Happenstance Type-O-Rama
PROOFREADER: James Saturnio, Word One New York
INDEXER: Robert Swanson
COVER DESIGNER: Ryan Sneed
COVER IMAGE: © Chad Baker / Lifesize / Getty Images

Search Engine Optimization Secrets

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

Copyright © 2011 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-55418-0
ISBN: 978-1-118-07829-7 (ebk)
ISBN: 978-1-118-07831-0 (ebk)
ISBN: 978-1-118-07830-3 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

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

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

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

Library of Congress Control Number: 2010929309

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

This book is dedicated to my family
(That's you, Mom, Dad, Jessica, and Josh!) for their
support and encouragement. I love all of you!

It is also dedicated to my fantastic friend Ian Lauth,
(not you Kevin Tower :-p) for his patience and support.
Thanks for putting up with me buddy!

Last but not least, I am dedicating this to all of my
brilliant co-workers at SEOmoz. Without all of you,
this would have been an unpublished disaster!

I don't know what I did to get lucky enough to have all of
you in my life but I appreciate my time with you every day.

—Danny Dover

To my wife and children, who love me even
though I never finish working when I say I will.

—Erik Dafforn

About the Authors

Danny Dover is a passionate SEO and influential writer. During his tenure at SE0moz.org (where he was the Lead SEO), he did SEO consulting for many of the world's most popular companies including Facebook, Microsoft, and Comcast. His expertise has been cited in *Time*, *PCWorld*, *Smashing Magazine*, and the *Seattle Post-Intelligencer* and has been translated into Japanese, French, Spanish, Italian, Chinese, German, and Hungarian.

Danny has spoken at numerous SEO conferences (spanning three continents) and his written posts and articles have been read over a million times and accessed online in more than 175 different countries.

Erik Dafforn is Executive Vice President and Director of Organic SEO for Intrapromote, LLC, a Cleveland-based Search and Social Marketing firm. At Intrapromote's blog and ClickZ.com, he's written over 200 articles on SEO strategy and techniques, many of which focus on architecture's effects on the crawling and indexing processes. Erik lives in Indianapolis with his wife and three children.

About the Technical Editor

Tim Buck worked for 15 years as IT Manager for a small software development company. Being the sole source of IT support there, he was responsible for server management, desktop support, web development, and software testing, and he wore many other hats as well. As a result, he learned a little about everything, including the basics of getting his company's website listed in Google's search engine results.

Now Tim works as a web application developer in state government; in this role, he continues to learn a little about everything, supporting legacy applications as well as developing new ones.

Acknowledgments

I would like to acknowledge the extraordinary efforts of Kevin Kent (Project Editor), Mary James (Acquisitions Editor), Carol Long (Executive Editor), and Jenny Watson (who originally found me for this book) for their work on this enormous and fun project. Your guidance and leadership made it possible for me to complete this book and I sincerely appreciate your patience and support.

I would also like to acknowledge the SEO community as a whole for creating the invigorating environment that made this book possible. Whether I met you in person, online, or not at all, you have been my driving force and an unconditional source of encouragement and important constructive criticism.

Thank you!

—Danny Dover

I would like to acknowledge the help and encouragement of several people who made my contribution to this book possible.

Thanks to Danny Dover for his hard work conceiving an excellent collection of content not typically found in SEO books. Contributing to the project has been an honor and a challenge.

Special thanks go to John Lustina and Doug Ausbury, co-founders of Intrapromote, LLC, for their encouragement during the writing stage; and to James Gunn, who was instrumental long ago in helping me understand fundamental SEO concepts and who continues to be a source of great insight and knowledge today.

Finally, I want to acknowledge the expertise and professionalism of the Wiley acquisitions, editorial, and production staff, including such excellent editors as Kevin Kent, Mary Beth Wakefield, and Mary James. They are an excellent team.

—Erik Dafforn

Contents at a Glance

Read This First ▶▶	1
Chapter 1 ▶ Understanding Search Engine Optimization	5
Chapter 2 ▶ Relearning How You See the Web	15
Chapter 3 ▶ Picking the Right SEO Tools	43
Chapter 4 ▶ Finding SEO Problems	85
Chapter 5 ▶ Solving SEO Problems	107
Chapter 6 ▶ SEO Best Practices	129
Chapter 7 ▶ The SEO Consulting Process	149
Chapter 8 ▶ Comprehensive Site Audit (Informational Website)	179
Chapter 9 ▶ Comprehensive Site Audit (E-Commerce Website)	259
Chapter 10 ▶ Understanding the SEO Industry	333
Chapter 11 ▶ Search Engine Verticals	355
Chapter 12 ▶ Optimizing for Alternative Search Engines	371
Chapter 13 ▶ Test, Test, Test	383
Chapter 14 ▶ SEO Resources	395
Appendix A ▶ Attending SEO Conferences	405
Index ▶	417

Contents

Read This First	1
Chapter 1 ▶ Understanding Search Engine Optimization	5
The Secrets of Popularity	6
The Secrets of Relevancy	11
Summary	14
Chapter 2 ▶ Relearning How You See the Web	15
The 1,000-Foot View—Understanding the Neighborhood	16
The 100-Foot View—The Website	21
The 10-Foot View—The Webpage	27
The 1-Foot View—Individual Content Pieces	40
Summary	42
Chapter 3 ▶ Picking the Right SEO Tools	43
View Source	44
Useful Search Engine Queries	51
Search Engine–Provided Tools	56
Google Webmaster Tools	59
Relevancy Determining Tools	68
SEO Toolbars	73
HTTP Header Analyzer	76
Firefox User Agent Switcher	78
Firefox Rendering Modifier	81
Summary	83
Chapter 4 ▶ Finding SEO Problems	85
The 15-Minute SEO Audit	86
The 5-Minute Brand Reputation Audit	97

Identifying Search Engine Penalties	103
Summary	106
Chapter 5 ▶ Solving SEO Problems	107
First Things First	108
Fixing Popularity Problems	109
Fixing Relevancy Problems	114
Fixing Penalties	117
Link Building Techniques	122
Summary	127
Chapter 6 ▶ SEO Best Practices	129
Page-Level Information Hierarchy	130
Domain-Level Information Hierarchy	130
URL	131
Title Tag	134
Meta Description	135
HTML Headings (H1 – H6)	136
Images	138
Nofollow	139
Rel Canonical Link Element	140
Meta Keywords	142
JavaScript and Flash	143
301 Redirects	143
Blocking Pages from Search Engines	144
Traffic and Rankings	145
Parameter-Driven URLs	146
Footer Links	147
Summary	148
Chapter 7 ▶ The SEO Consulting Process	149
Answering Hard Questions from Clients	150
Preparing for the First Meeting	159

Paperwork: Setting Expectations	165
Deliverables	169
Establishing Price Points	170
What to Give Away for Free	175
Subsequent Meetings	176
Summary	178
Chapter 8 ▶ Comprehensive Site Audit (Informational Website)	179
How to Read This Chapter	180
Sample Website	181
Sample Report	182
After Completing the Report	199
Summary	200
MusicArtistDatabase.com Sample SEO Site Audit	201
Chapter 9 ▶ Comprehensive Site Audit (E-Commerce Website)	259
How to Read This Chapter	260
Sample Website	260
Sample Report	261
Summary	277
SellTheWidget.com Sample SEO Site Review	279
Chapter 10 ▶ Understanding the SEO Industry	333
A Brief History of SEO	334
Who Are Internet Marketers and Where Can I Find Them?	335
The SEO Pyramid and Wearing Multiple Hats	340
SEO Leaders	343
The People and Technology behind Google and Bing	346
Long-Term Perspective in SEO	351
Summary	353
Chapter 11 ▶ Search Engine Verticals	355
Universal Search	356
Local Search	357

	Image Search	362
	Video Search	365
	Instant Search	368
	Summary	370
Chapter 12	▶ Optimizing for Alternative Search Engines.....	371
	Amazon	372
	YouTube	374
	Facebook	375
	Delicious	376
	Flickr	378
	Twitter	380
	Summary	382
Chapter 13	▶ Test, Test, Test.....	383
	Setting Up a Testing Platform	384
	Running a Test	387
	Recording Results	389
	The Importance of Sharing Knowledge	391
	Summary	393
Chapter 14	▶ SEO Resources.....	395
	SEO Cheat Sheet Part 1: On-Page Optimization	396
	SEO Cheat Sheet Part 2: Canonicalization Errors	397
	SEO Cheat Sheet Part 3: Meta Robots and Robots.txt	398
	SEO Cheat Sheet Part 4: Sitemaps	400
	SEO Cheat Sheet Part 5: User Agents	402
	Switching Domains Checklist	403
	SEO Quick Hit List	404
	Summary	404
Appendix A	▶ Attending SEO Conferences.....	405
	Picking the Right Conference	406
	What to Expect	410

Preparing to Go	412
What to Bring	413
What to Do	414
After the Conference	416
Summary	416
Index	417

Read This First

Why would someone like myself want to publish my SEO secrets for the world to read? Doesn't this destroy my competitive advantage? Won't I surely go broke and starve on the street? Won't my friends mock me and my family disown me?

For two reasons, the answer is probably not.

- ▶ The first reason is the size of the market. The Internet is incredibly large and growing at an astounding rate. The market for SEO is following a similar path. There is absolutely no way I could work for all of the websites that need SEO consulting. As such, I am happy to pass the work on to others and teach them how to succeed. It is no money out of my pocket, and it makes me feel like I am contributing to a greater good. I learned most of what I know about SEO from others and, as such, feel obligated to spread the knowledge.
- ▶ The second reason has to do with SEOMoz, the company I used to work for. SEOMoz provides tools to help SEOs do their jobs. As such, it is to my advantage to promote and train other SEOs. Just like Google benefits from getting more people online, I benefit from teaching others how to do SEO. You may choose to use SEOMoz's competitors' services or you may not. That is completely up to you, and I will do my best to show you all the available options.

WHO THIS BOOK IS FOR

This book is for the SEO who already knows the basics of SEO and wants to take this knowledge to the next level so that they can make more money. In the SEO industry, the best way I have found to do this is to do SEO consulting.

This book is written as a guide to becoming an SEO consultant or for those who want to use the strategies of professional SEO consultants. It clearly lays out the processes and perspectives I have used at SEOMoz when I did consulting for some of the most well-known websites on the Internet. It is intended for those who love the Internet and strive to influence how it operates.

WHY THIS BOOK IS BETTER THAN OTHER SEO BOOKS

Modern SEO is complicated, fast moving, and rife with misconceptions. This makes it extremely difficult to learn. When I began researching for this book, I read all of the major SEO books that were available. I quickly found that they were full of theory and lacked actionable steps to really help the reader master the subject.

I wrote this book with the goal of building the bridge between theory and action by bringing together all of the best sources of information I have found and putting them in a format that makes it easy to understand and, more importantly, do SEO like a professional. This emphasis on action follows the steps I originally used to learn SEO. I believe this focus on process followed by explanation is unique among SEO books on the market, and I believe it will make the difference that allows you to out rank your competition.

HOW I LEARNED THE SECRETS SHARED IN THIS BOOK

The brutal truth is that I do not work at Google or Microsoft and I have never read a single line of code that powers the search engine algorithms. Surprisingly, as an SEO professional, I am not unique.

So what gives me the authority to write a book about SEO? The answer is simple. I get results. I have dedicated years of my life to studying the search engines and have learned how to influence search engine result pages. I use my skills almost every day to help people improve their rankings and drive traffic to their sites. To me, there is no better feeling than helping people achieve their online dreams.

This book is the next step for me. Instead of helping others in a one-to-one fashion, this book will enable me to help others in a one-to-many fashion. That is where you come in. My hope is that after reading this book, you will choose to use your skills to help others (but be sure to look out for yourself first). Either way I support you.

WEBSITE SUPPORTING THE BOOK

You will find additional supporting material at the accompanying online resource at www.dannydover.com/search-engine-optimization-secrets/. This resource includes:

- ▶ Beginner's Guide to SEO
- ▶ A Comprehensive SEO Audit Report (Informational Website)
- ▶ A Comprehensive SEO Audit Report (E-commerce Website)
- ▶ A Center for Learning SEO
- ▶ Web Developer's SEO Cheat Sheet
- ▶ Internet Marketing Handbook
- ▶ 15 Minute SEO Audit Checklist
- ▶ Updates to this book
- ▶ Resources on how to learn more

FEATURES AND ICONS USED IN THIS BOOK

The following features and icons are used to help draw your attention to some of the most important or useful information in the book, some of the most valuable tips, insights, and advice.

▶ Watch for margin notes like this one that highlight some key piece of information or that discuss some valuable technique or approach.

SIDEBARS

Sidebars like this one feature additional information about topics related to the nearby text.

TIP The Tip icon indicates a helpful trick or technique.

NOTE The Note icon points out or expands on items of importance or interest.

CROSSREF The Cross-Reference icon points to chapters where additional information can be found.

WARNING The Warning icon warns you about possible negative side effects or precautions you should take before making a change.

Enough talk; it is now time to get started. Thank you, and best of luck with your Internet endeavors.

CHAPTER 1

Understanding Search Engine Optimization

IN THIS CHAPTER

- ▶ Learning how search engines see websites
- ▶ Taking a look at popularity in SEO
- ▶ Considering the role of relevancy in SEO

At Google, search engineers talk about “80-20” problems.

They are describing situations where the last 20 percent of the problem is 80 percent of the work. Learning SEO is one of these problems. Eighty percent of the knowledge SEOs need is available online for free. Unfortunately, the remaining 20 percent takes the majority of the time and energy to find and understand. My goal with this book is to solve this problem by making the last 20 percent as easy to get as the first 80 percent. Though I don't think I will be able to cover the entire 20 percent (some of it comes from years of practice), I am going to write as much actionable advanced material as humanly possible.

This book is for those who already know the basics of SEO and are looking to take their skills to the next level. Before diving in, try reading the following list:

- ▶ robots.txt
- ▶ sitemap
- ▶ nofollow
- ▶ 301 redirect
- ▶ canonicalization

If you are not sure what any of the items in this list are, you should go over to the nearest computer and read the article “The Beginner’s Guide to SEO” at

<http://www.seomoz.org/article/beginners-guide-to-search-engine-optimization>

This free article can teach you everything you need to know to use this book to its fullest. Done with that? Great, now we can begin.

THE SECRETS OF POPULARITY

Once upon a time there were two nerds at Stanford working on their PhDs. (Now that I think about it, there were probably a lot more than two nerds at Stanford.) Two of the nerds at Stanford were not satisfied with the current options for searching online, so they attempted to develop a better way.

Being long-time academics, they eventually decided to take the way academic papers were organized and apply that to webpages. A quick and fairly objective way to judge the quality of an academic paper is to see how many times other academic papers have cited it. This concept was easy to replicate online because the original purpose of the Internet was to share academic resources between universities. The citations manifested themselves as hyperlinks once they went online. One of the nerds came up with an algorithm for calculating these values on a global scale, and they both lived happily ever after.

Of course, these two nerds were Larry Page and Sergey Brin, the founders of Google, and the algorithm that Larry invented that day was what eventually became PageRank. Long story short, Google ended up becoming a big deal and now the two founders rent an airstrip from NASA so they have somewhere to land their private jets. (Think I am kidding? See <http://searchengineland.com/your-guide-to-the-google-jet-12161>.)

RELEVANCE, SPEED, AND SCALABILITY

Hypothetically, the most relevant search engine would have a team of experts on every subject in the entire world—a staff large enough to read, study, and evaluate every document published on the web so they could return the most accurate results for each query submitted by users.

The fastest search engine, on the other hand, would crawl a new URL the very second it's published and introduce it into the general index immediately, available to appear in query results only seconds after it goes live.

The challenge for Google and all other engines is to find the balance between those two scenarios: To combine rapid crawling and indexing with a relevance algorithm that can be instantly applied to new content. In other words, they're trying to build *scalable relevance*. With very few exceptions, Google is uninterested in hand-removing (or hand-promoting) specific content. Instead, its model is built around identifying characteristics in web content that indicate the content is especially relevant or irrelevant, so that content all across the web with those same characteristics can be similarly promoted or demoted.

This book frequently discusses the benefits of content created with the user in mind. To some hardcore SEOs, Google's "think about the user" mantra is corny; they'd much prefer to know a secret line of code or server technique that bypasses the intent of creating engaging content.

While it may be corny, Google's focus on creating relevant, user-focused content really is the key to its algorithm of scalable relevance. Google is constantly trying to find ways to reward content that truly answers users' questions and ways to minimize or filter out content built for content's sake. While this book discusses techniques for making your content visible and accessible to engines, remember that means talking about content constructed with users in mind, designed to be innovative, helpful, and to serve the query intent of human users. It might be corny, but it's effective.

That fateful day, the Google Guys capitalized on the mysterious power of links. Although a webmaster can easily manipulate everything (word choice, keyword placement, internal links, and so on) on his or her own website, it is much more difficult to influence inbound links. This natural link profile acts as an extremely good metric for identifying legitimately popular pages.

NOTE Google's PageRank was actually named after its creator, Larry Page. Originally, the algorithm was named BackRub after its emphasis on backlinks. Later, its name was changed to PageRank because of its connections to Larry Page's last name and the ability for the algorithm to rank pages.

Larry Page's original paper on PageRank, "The Anatomy of a Large-Scale Hypertextual Web Search Engine," is still available online. If you are interested in reading it, it is available on Stanford's website at <http://infolab.stanford.edu/~backrub/google.html>. It is highly technical, and I have used it on more than one occasion as a sleep aid. It's worth noting that the original PageRank as described in this paper is only a tiny part of Google's modern-day search algorithm.

Now wait a second— isn't this supposed to be a book for advanced SEOs? Then why am I explaining to you the value of links? Relax, there is a method to my madness. Before I am able to explain the more advanced secrets, I need to make sure we are on the same page.

As modern search engines evolved, they started to take into account the link profile of both a given page and its domain. They found out that the relationship between these two indicators was itself a very useful metric for ranking webpages.

Domain and Page Popularity

There are hundreds of factors that help engines decide how to rank a page. And in general, those hundreds of factors can be broken into two categories—relevance and popularity (or "authority"). For the purposes of this demonstration you will need to completely ignore relevancy for a second. (Kind of like the search engine Ask.com.) Further, within the category of popularity, there are two primary types—domain popularity and page popularity. Modern search engines rank pages by a combination of these two kinds of popularity metrics. These metrics are measurements of link profiles. To rank number one for a given query you need to have the highest amount of total popularity on the Internet. (Again, bear with me as we ignore relevancy for this section.)

This is very clear if you start looking for patterns in search result pages. Have you ever noticed that popular domains like wikipedia.org tend to rank for everything? This is because they have an enormous amount of domain popularity. But what about those competitors who outrank me for a specific term with a practically unknown domain? This happens when they have an excess of page popularity. See Figure 1-1.

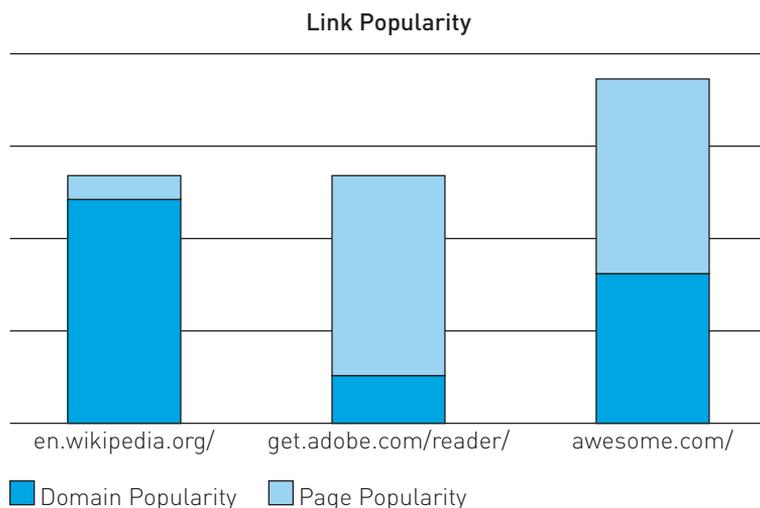


FIGURE 1-1: Graph showing different combinations of relevancy and popularity metrics that can be used to achieve high rankings

Although `en.wikipedia.org` has a lot of domain popularity and `get.adobe.com/reader/` has a lot of page popularity, `www.awesome.com` ranks higher because it has a higher *total* amount of popularity. This fact and relevancy metrics (discussed later in this chapter) are the essence of Search Engine Optimization. (Shoot! I unveiled it in the first chapter, now what am I going to write about?)

POPULARITY TOP TEN LISTS

The top 10 most linked-to domains on the Internet (at the time of writing) are:

- ▶ Google.com
- ▶ Adobe.com
- ▶ Yahoo.com
- ▶ Blogspot.com
- ▶ Wikipedia.org
- ▶ YouTube.com
- ▶ W3.org
- ▶ Myspace.com

continued

(continued)

▶ Wordpress.com

▶ Microsoft.com

The top 10 most linked-to pages on the Internet (at the time of writing) are:

▶ <http://wordpress.org/>

▶ <http://www.google.com/>

▶ <http://www.adobe.com/products/acrobat/readstep2.html>

▶ <http://www.miibeian.gov.cn/>

▶ <http://validator.w3.org/check/referer>

▶ <http://www.statcounter.com/>

▶ <http://jigsaw.w3.org/css-validator/check/referer>

▶ <http://www.phpbb.com/>

▶ <http://www.yahoo.com/>

▶ <http://del.icio.us/post>

Source: SEOmoz's Linkscape—Index of the World Wide Web

▶ Not only that, but at any given time, the TbPR (Toolbar PageRank) value you see may be up to 60–90 days older or more, and it's a single-digit representation of what's probably very a long decimal value.

▶ Google makes scraping (automatically requesting and distributing) its PageRank metric difficult. To get around the limitations, you need to write a program that requests the metric from Google and identifies itself as the Google Toolbar.

Before I summarize I would like to nip the PageRank discussion in the bud. Google releases its PageRank metric through a browser toolbar. This is not the droid you are looking for. That green bar represents only a very small part of the overall search algorithm.

Just because a page has a PageRank of 5 does not mean it will outrank all pages with a PageRank of 4. Keep in mind that major search engines do not want you to reverse engineer their algorithms. As such, publicly releasing a definitive metric for ranking would be idiotic from a business perspective. If there is one thing that Google is not, it's idiotic.

In my opinion, hyperlinks are the most important factor when it comes to ranking web pages. This is the result of them being difficult to manipulate. Modern search engines look at link profiles from many different perspectives and use those relationships to determine rank. The takeaway for you is that time spent earning links is time well spent. In the same way that a rising tide raises all ships, popular domains raise all pages. Likewise, popular pages raise the given domain metrics.

In the next section I want you to take a look into the pesky missing puzzle piece of this chapter: relevancy. I am going to discuss how it interacts with popularity, and I may or may not tell you another fairy tale.

THE SECRETS OF RELEVANCY

In the previous section, I discussed how popular pages (as judged by links) rank higher. By this logic, you might expect that the Internet's most popular pages would rank for everything. To a certain extent they do (think Wikipedia!), but the reason they don't dominate the rankings for *every* search result page is that search engines put a lot of emphasis on determining relevancy.

Text Is the Currency of the Internet

Relevancy is the measurement of the theoretical distance between two corresponding items with regards to relationship. Luckily for Google and Microsoft, modern-day computers are quite good at calculating this measurement for text.

By my estimations, Google owns and operates well over a million servers. The electricity to power these servers is likely one of Google's larger operating expenses. This energy limitation has helped shape modern search engines by putting text analysis at the forefront of search. Quite simply, it takes less computing power and is much simpler programmatically to determine relevancy between a text query and a text document than it is between a text query and an image or video file. This is the reason why text results are so much more prominent in search results than videos and images.

As of this writing, the most recent time that Google publicly released the size of its indices was in 2006. At that time it released the numbers shown in Table 1-1.

TABLE 1-1: Size of Google Indices

DATA	SIZE IN TERABYTES
Crawl Index	800
Google Analytics	200
Google Base	2
Google Earth	70
Orkut	9
Personalized Search	4

► This is especially true until Google finds better ways to interpret and grade non-textual media

So what does this emphasis on textual content mean for SEOs? To me, it indicates that my time is better spent optimizing text than images or videos. This strategy will likely have to change in the future as computers get more powerful and energy efficient, but for right now **text should be every SEO's primary focus.**

But Why Content?

The most basic structure a functional website could take would be a blank page with a URL. For example purposes, pretend your blank page is on the fake domain `www.WhatIsJessicaSimpsonThinking.com`. (Get it? It is a blank page.) Unfortunately for the search engines, clues like top-level domains (.com, .org, and so on), domain owners (WHOIS records), code validation, and copyright dates are poor signals for determining relevancy. This means your page with the dumb domain name needs some content before it is able to rank in search engines.

The **search engines must use their analysis of content as their primary indication of relevancy** for determining rankings for a given search query. For SEOs, this means the content on a given page is essential for manipulating—that is, earning—rankings. In the old days of AltaVista and other search engines, SEOs would just need to write “Jessica Simpson” hundreds of times on the site to make it rank #1 for that query. What could be more relevant for the query “Jessica Simpson” than a page that says Jessica Simpson 100 times? (Clever SEOs will realize the answer is a page that says “Jessica Simpson” 101 times.) This metric, called *keyword density*, was quickly manipulated, and the search engines of the time diluted the power of this metric on rankings until it became almost useless. Similar dilution has happened to the keywords meta tag, some kinds of internal links, and H1 tags.

Hey, Ben Stein, thanks for the history lesson, but how does this apply to modern search engines? The funny thing is that modern-day search engines still work essentially the same way they did back in the time of keyword density. The big difference is that they are now much more sophisticated. Instead of simply counting the number of times a word or phrase is on a webpage, **they use natural language processing algorithms and other signals on a page to determine relevancy.** For example, it is now fairly trivial for search engines to determine that a piece of content is about Jessica Simpson if it mentions related phrases like “Nick Lachey” (her ex-husband), “Ashlee Simpson” (her sister), and “Chicken of the Sea” (she is infamous for thinking the tuna brand “Chicken of the Sea” was made from chicken). The engines can do this for a multitude of languages and with astonishing accuracy.

Don't believe me? Try going to Google right now and searching `related:www.jessicasimpson.com`. If your results are like mine, you will see websites about her movies, songs, and sister. Computers are amazing things.

► Despite being more sophisticated, modern-day search engines still work essentially the same way they did in the past—by analyzing content on the page.

In addition to the words on a page, search engines use signals like image meta information (alt attribute), link profile and site architecture, and information hierarchy to determine how relevant a given page that mentions “Jessica” is to a search query for “The Simpsons.”

Link Relevancy

As search engines matured, they started identifying more metrics for determining rankings. One that stood out among the rest was link relevancy.

The difference between link relevancy and link popularity (discussed in the previous section) is that link relevancy does not take into account the power of the link. Instead, it is a natural phenomenon that works when people link out to other content.

Let me give you an example of how it works. Say I own a blog where I write about whiteboard markers. (Yes, I did just look around my office for an example to use, and yes, there are actually people who blog about whiteboard markers. I checked.) Ever inclined to learn more about my passion for these magical writing utensils, I spend part of my day reading online what other people have to say about whiteboard markers.

On my hypothetical online reading journey, I find an article about the psychological effects of marker color choice. Excited, I go back to my website to blog about the article so (both of) my friends can read about it. Now here is the critical takeaway. When I write the blog post and link to the article, *I get to choose the anchor text*. I could choose something like “click here,” but more likely I choose something that it is relevant to the article. In this case I choose “psychological effects of marker color choice.” Someone else who links to the article might use the link anchor text “marker color choice and the effect on the brain.”

This human-powered information is essential to modern-day search engines. These descriptions are relatively unbiased and produced by real people. This metric, in combination with complicated natural language processing, makes up the lion’s share of relevancy indicators online.

Other important relevancy indicators are link sources and information hierarchy. For example, the search engines can also use the fact that I linked to the color choice article from a blog about whiteboard markers to supplement their understanding of relevancy. Similarly, they can use the fact that the original article was located at the URL `www.example.com/vision/color/` to determine the high-level positioning and relevancy of the content. As you read later in this book (Chapter 2 specifically), these secrets are essential for SEOs to do their job.

Beyond specific anchor text, proximal text—the certain number of characters preceding and following the link itself—have some value. Something that’s logical,

► People have a tendency to link to content using the anchor text of either the domain name or the title of the page. Use this to your advantage by including keywords you want to rank for in these two elements.

but annoying is when people use a verb as anchor text, such as “Frank said . . .” or “Jennifer wrote . . .”, using “said” or “wrote” as the anchor text pointing back to the post. In a situation like that, engines have figured out how to apply the context of the surrounding copy to the link.

Tying Together Popularity and Relevancy

So far in this chapter I have discussed both popularity and relevancy. These two concepts make up the bulk of Search Engine Optimization theory. They have been present since the beginning of search engines and undoubtedly will be important in the future. The way they are determined and the relationship between them changes, but they are both fundamental to determining search results.

This fact is critical to SEOs. We have very little control over how the major search engines operate, yet somehow we are supposed to keep our jobs. Luckily, these immutable laws of popularity and relevance govern search engines and provide us with some job security.

SUMMARY

In this chapter, I explained the concepts of popularity and relevancy in relation to modern search engines. This information, along with your prior SEO experience, will make up the foundation for all of the SEO secrets and knowledge that you learn throughout the rest of the book. You no doubt have some questions. I’ll start answering many of your questions in the next chapter, but you will likely form many more. Welcome to the mindset of a Professional SEO. Prepare to be questioning and Googling things for the rest of your life.

► Popularity and relevancy are the two concepts that make up the bulk of Search Engine Optimization theory.

CHAPTER 2

Relearning How You See the Web

IN THIS CHAPTER

- ▶ **Analyzing how a website fits in its “web neighborhood”**
- ▶ **Viewing websites like an SEO**
- ▶ **Assessing good site architecture and webpages from an SEO perspective**
- ▶ **Assessing website content like an SEO**

When people surf the Internet, they generally view each domain as its own island of information. This works perfectly well for the average surfer but is a big mistake for beginner SEOs. Websites, whether they like it or not, are interconnected. This is a key perspective shift that is essential for understanding SEO.

Take Facebook, for example. It started out as a “walled garden” with all of its content hidden behind a login. It thought it could be different and remain completely independent. This worked for a while, and Facebook gained a lot of popularity. Eventually, an ex-Googler and his friend became fed up with the locked-down communication silo of Facebook and started a wide open website called Twitter. Twitter grew even faster than Facebook and challenged it as the media darling. Twitter was smart and made its content readily available to both developers (through APIs) and search engines (through indexable content).

Facebook responded with Facebook Connect (which enables people to log in to Facebook through other websites) and opened its chat protocol so its users could communicate outside of the Facebook domain. It also made a limited amount of information about users visible to search engines. Facebook is now accepting its place in the Internet community and is benefiting from its decision to embrace other websites. The fact that it misjudged early on was that websites are best when they are interconnected. Being able to see this connection is one of the skills that separates SEO professionals from SEO fakes.

In this chapter you learn the steps that the SEO professionals at SEOmoz go through either before meeting with a client or at the first meeting (depending on the contract). When you view a given site in the way you are about to learn in this chapter, **you need to take detailed notes**. You are likely going to notice a lot about the website that can use improvement, and you need to capture this information before details distract you.

► I highly recommend writing down everything you notice in a section of a notebook identified with the domain name and date of viewing.

KEEP YOUR NOTES SIMPLE

The purpose of the notebook is simplicity and the ability to go back frequently and review your notes. If actual physical writing isn't your thing, consider a low-tech text editor on your computer, such as Windows Notepad or the Mac's TextEdit.

Bare-bones solutions like a notebook or text editor help you avoid the distraction of the presentation itself and focus on the important issues—the characteristics of the web site that you're evaluating.

If you think it will be helpful and you have Internet access readily available, I recommend bringing up a website you are familiar with while reading through this chapter. If you choose to do this, be sure to take a lot of notes in your notebook so you can review them later.

THE 1,000-FOOT VIEW—UNDERSTANDING THE NEIGHBORHOOD

Before I do any work on a website I try to get an idea of where it fits into the grand scheme of things on the World Wide Web. The easiest way to do this is to run searches for some of the competitive terms in the website's niche. If you imagine the Internet as one

giant city, you can picture domains as buildings. The first step I take before working on a client's website is figuring out in which neighborhood its building (domain) resides.

This search result page is similar to seeing a map of the given Internet neighborhood. You usually can quickly identify the neighborhood anchors (due to their link popularity) and specialists in the top 10 (due to their relevancy). You can also start to get an idea of the maturity of the result based on the presence of spam or low-quality websites. Take a look at Figures 2-1 and 2-2.

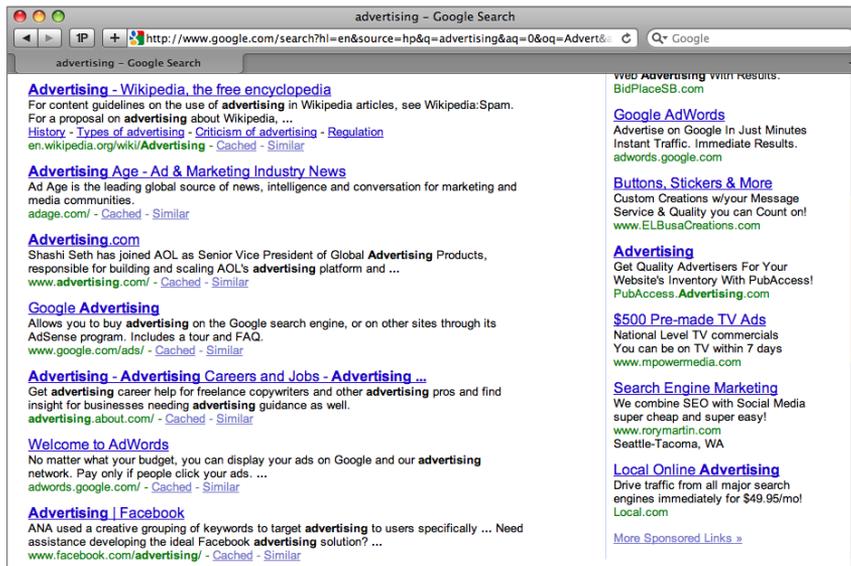


FIGURE 2-1: Google search result for "advertising"

Notice the difference in the maturity (quality) of the search results. In the second set of results (Figure 2-2), you see some of the same big names again (Wikipedia, for example, appears in both searches) but this time they are mixed with some sites that appear spammier (iab.net, freewebdirectory.us).

During client meetings, when I look at the search engine result page for a competitive term like *advertising*, I am not looking for websites to visit but rather trying to get a general idea of the maturity of the Internet neighborhood. I am very vocal when I am doing this and have been known to question out loud, "How did that website get there?" A couple times, the client momentarily thought I was talking about his website and had a quick moment of panic. In reality, I am commenting on a spam site I see rising up the results.

Also, take note that regardless of whether or not you are logged into a Google account, the search engine will automatically customize your search results based on links you click most. This can be misleading because it will make your favorite websites rank higher for you than they do for the rest of the population.

► To turn this off, append "&pws=0" to the end of the Google URL.

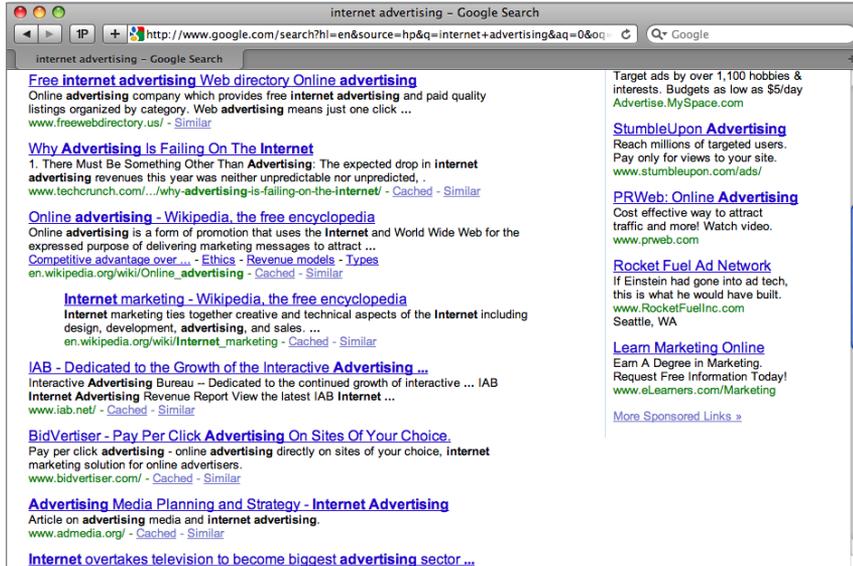


FIGURE 2-2: Google search result for “Internet advertising”

Along with looking at the results themselves, I look at the other data present on the page. The amount of advertisements on the search result gives a rough idea of how competitive it is. For example, a search for buy *viagra* will return a full page height worth of ads, whereas a search for women that look like Drew Carey won't likely return any. This is because more people are searching for the blue pill than are searching for large, bald women with nerd glasses.

In addition to the ads, I also look for signs of temporal algorithms. *Temporal algorithms* are ranking equations that take into account the element of time with regards to relevancy. These tend to manifest themselves as news results and blog posts.

TAKING ADVANTAGE OF TEMPORAL ALGORITHMS

You can use the temporal algorithms to your advantage. I accidentally did this once with great success. I wrote a blog post about Michael Jackson's death and its effect on the search engines a day after he died. As a result of temporal algorithms my post ranked in the top 10 for the query “Michael Jackson” for a short period following his death. Because of this high ranking, tens of thousands of people read my article. I thought it was because I was so awesome, but after digging into my analytics I realized it was because of unplanned use of the temporal algorithms. If you are a blogger, this tactic of quickly writing about news events can be a great traffic booster.