

Greg Nudelman

with Contributing Editor Pabini Gabriel-Petit

Foreword by internationally recognized information architecture expert and bestselling author Peter Morville

DESIGNING SEARCH

UX Strategies for eCommerce Success



OPTIMIZING ECOMMERCE SEARCH RESULTS PAGES

Search is a wicked problem of terrific consequence... the defining element of user experience. It changes the way we find everything... It shapes how we learn and what we believe.

—PETER MORVILLE

- CHAPTER 1** Starting from Zero: Winning Strategies for No Search Results Pages
- CHAPTER 2** How Shoppers Search
- CHAPTER 3** Choosing the Right Search Results Page Layout
- CHAPTER 4** Balancing Pogosticking and Page Relevance
- CHAPTER 5** Making \$10,000 a Pixel: Optimizing Thumbnail Images in Search Results
- CHAPTER 6** Best Practices for Ads in Search Results

**Please enjoy this sample chapter. If you find it useful,
we'd love a tweet! Just click the button on the right.**

Thanks!

Greg Nudelman, Author



STARTING FROM ZERO:

Winning Strategies for No Search Results Pages

The No Search Results Page: Your Key to Competitive Search Advantage **PAGE 4**

Create a Robust Partial Match Strategy **PAGE 11**

Employ Multiple Content Strategies **PAGE 15**

Case Study: Hotmail No Search Results Page **PAGE 17**

How Can Your Site-Search Analytics Help You, by Louis Rosenfeld **PAGE 19**

Search, more than any other activity, is a living, evolving process of discovery—a conversation between a customer and the Web site. Unfortunately, this conversation is often fraught with miscommunication, and so it is critical for you to keep this conversation going even when the customer has initiated a search that yielded no results.

Unfortunately, the typical product team has no coherent strategy for cases when there are no search results, generally treating such occurrences as a “user error.” Most teams spend the bulk of their design phase working on the search results pages for a successful search. Then in the harried moments prior to launch, the engineers hurriedly slap something together for the no search results page. Such an approach is detrimental to the success of a customer’s search experience.

In my experience, the effort and ingenuity a product team invests in the no search results page is indicative of its overall dedication to customer success. Ignoring this special kind of search results page virtually guarantees a mediocre search experience and contributes to obscurity in the ecommerce marketplace. On the other hand, if a team thinks creatively about the case when there are no search results and focuses on customer needs, it can turn a temporary snag in communication into an opportunity for deeper connection and a source of tremendous competitive advantage.

THE NO SEARCH RESULTS PAGE: YOUR KEY TO COMPETITIVE SEARCH ADVANTAGE

To see how original thinking about the no search results page can revolutionize an entire industry, consider the history of Google. Although a relative late-comer to an already saturated search market, Google blew away all competition through an unapologetic dedication to customer success. One critical innovation was the Google *Did you mean...* feature, which gave the process of discovery a safety net and made exploration more fun. This feature was the result of deliberate and original thinking about how to help people correct the misspellings that were a common cause of the appearance of the no search results page. Controlled vocabulary substitution redefined the way Google does search, and today, the *Did you mean...* feature shown in Figure 1.1 is a virtual necessity for a successful search implementation.

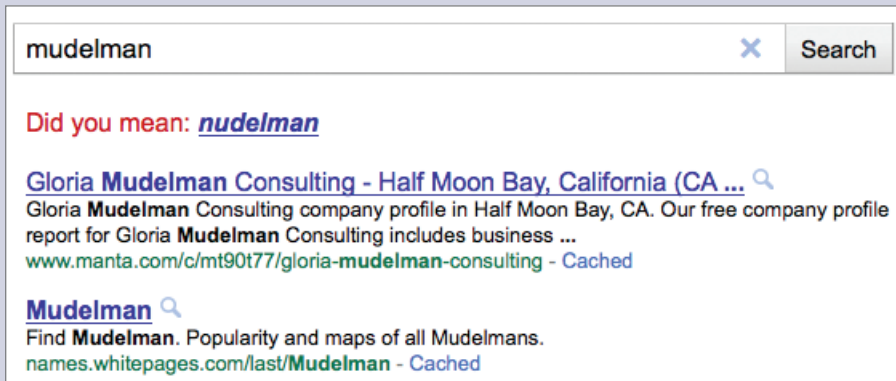


FIGURE 1.1

The Google *Did you mean...* feature

If you want to create a killer search app, begin with the no search results page. Starting from zero forces your team to address the most difficult design questions up front and honestly assess your engineering budget and capabilities. More important, your team can define the entire search user interface design problem more precisely—in a compelling and possibly even original way.

No Search Results Strategy: Not a Zero-Sum Game

No simple set of rules exist that guarantees a successful implementation of a no search results page. However, the following four broad design principles provide a useful starting point:

1. Don't be afraid to say *I did not understand*. Clearly indicate there are no search results, so your customer can recover.
2. Focus on providing a way out. Make sure every control on the search results page does something productive to help resolve the no search results condition.
3. Create a robust partial match strategy. Over-constraining is the most frequent mistake people make when searching ecommerce sites. Having a robust partial match strategy is critical.
4. Employ multiple content strategies. Draw from multiple sources to provide the most relevant content first to aid recovery from the no search results condition, while staying true to your customer's original intent.

Now look at some examples of how you can combine these design principles creatively to develop both compelling customer experiences and opportunities for business success.

Don't Be Afraid to Say *I Did Not Understand*

Clear communication of system state forms the foundation of all human-computer interaction. In his *Alertbox* blog post “Ten Usability Heuristics,” Jacob Nielsen put “Visibility of system status” at #1. If customers cannot understand that a search system did not find what they asked for, they can’t take the appropriate corrective action. Yet many search applications do not display the no search results condition clearly or accurately.

Figure 1.2 shows an example of a search results page from a leading financial information site, Morningstar.com. Can you tell whether the system found what the customer searched for?

FIGURE 1.2

Morningstar.com does not communicate zero search results effectively

The screenshot shows the Morningstar.com website. At the top, there is a navigation bar with links like 'Join', 'Home', 'Portfolio', 'Stocks', 'Options', 'Funds', 'ETFs', 'Hedge Funds', 'Markets', and 'Tools'. Below this, there is a search bar with the text 'Opphmr Small & Mid Cap Value Fund' and a 'VIEW RESULTS' button. The search results are displayed in a table with columns for 'Date', 'All', and '1-0 of 0'. The table is empty, indicating no results were found. To the right of the search bar, there is a sidebar with links for 'Membership', 'My Portfolio', 'Email', 'Discuss', 'My Profile', 'Glossary', and 'Help'. At the bottom right, there is a 'tiny*prints' advertisement for holiday cards and party invitations.

On the Morningstar search results page, there is no explicit message that states **No results found**. The only zero on the page appears in the black bar in the form of **1-o of o**, which does not make much sense. Compare the Morningstar page with the Google page shown in Figure 1.3.

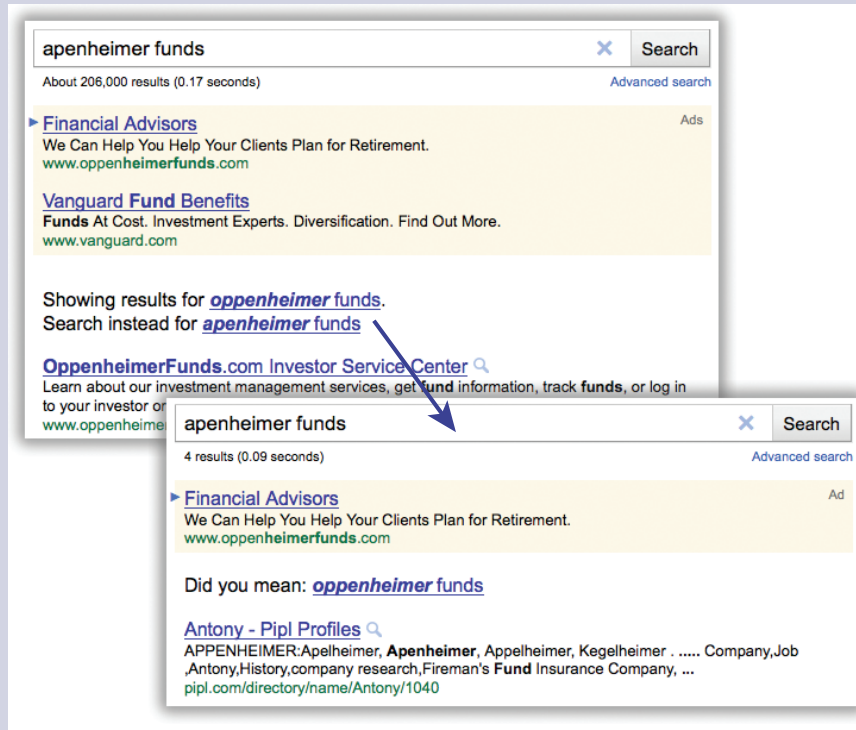


FIGURE 1.3

Google clearly communicates a zero search results condition and often takes appropriate corrective action automatically

The status message on the Google search results page is clear and straightforward: **Showing results for apenheimer funds. Search instead for apenheimer funds**, making it obvious that the system assumed a misspelling and made the most likely substitution automatically. In case the person actually did want to search for **apenheimer**, Google also makes the original query available. Clicking **Search instead for apenheimer funds** yields the result shown in Figure 1.3 in the lower right.



If you look carefully, you noticed that this Google system behavior is a little different than the **Did you mean:** implementation shown in Figure 1.1. In the case of the search for **Mudleman**, the system merely suggested the substitution

but still ran the original query. On the other hand, in the case of **apenheimer**, the system performed the substitution automatically. The words **Did you mean:** appear in red font in Figure 1.1, whereas they are less prominent in black ink in the lower right in Figure 1.3. This system behavior demonstrates another key principle: The best design for your application is the one that works best for a particular query, your audience, and your business model. There are no hard-and-fast rules for implementing any of the strategies that handle no search results conditions. Even the industry's best sites work tirelessly to make continuous, small improvements. No search results pages are critical to your site's success, so it pays to keep a particularly close eye on any design changes to make sure that they are indeed providing a better experience for your customers. That is why it is important to continuously test and evaluate the performance of your no search results pages using A/B testing (see sidebar).

A/B TESTING AND MULTIVARIATE TESTING

To ensure you create successful design solutions that meet both business and customer goals, your team should conduct frequent, quantitative A/B testing of your no search results page and other search results pages. Follow up with qualitative lab and field testing to help you make sense of your A/B testing results and suggest ideas for future improvements.

The central idea behind A/B testing is to have two different user interface designs running on your site at the same time, while collecting key performance metrics (KPMs) that enable you to measure desired customer behavior.

For example, say you want to introduce some improvements to your current no search results page, which you can call variant A. To determine whether your redesigned version of the page, variant B, offers any improvement, you can deploy variant B and send a small percentage of site visitors—for example, 1% to 10%—to that server and observe the metrics for variant B:

- Did those visitors buy more stuff?
- Did they stay on your site longer?
- Did they add more things to their favorites or wish lists?
- Did they come back to your site sooner to make additional purchases—showing loyalty to your site?

Of course, these are ecommerce metrics, but you can set up your own relevant KPMs and figure out whether variant A or B is more successful.

Most successful companies use multivariate testing, which can be thought of in simple terms as numerous A/B tests performed on one page at the same time. Whereas A/B tests are usually performed to determine the better of two content variations, multivariate testing can theoretically test the effectiveness of limitless combinations. In practice the number of combinations is limited by the amount of time it will take to get a statistically valid sample of visitors and computational power. However, it can be substantial, with some of the larger multivariate tests including 40 or more combinations.

Multivariate testing is the primary method by which Netflix decides whether any new user interface design is an improvement and should go live on its site. Netflix regularly displays new user interfaces to a small percentage of its customers for a few hours or days at a time. If customers are more successful using a new user interface, variant B, according to the KPMs, that new variant goes live on the site for all customers. If not, the new design is rejected. By doing this type of testing, Netflix can remain confident that its site is continuously improving because the key metrics tell that it is. This improvement strategy has certainly worked well for Netflix. Since 2005, the company has consistently ranked #1 in customer satisfaction across the entire ecommerce industry!

Focus on Providing a Way Out

After the system indicates that the no search results condition occurred, it must now help the customer recover. Whenever you display a no search results page, always provide a helpful way forward to get your customer back on track as quickly as possible. Virtually every control on the page should be focused on doing *something* to help the customer recover from the no search results condition, and any extraneous controls for filtering and sorting search results should be removed.

On the Morningstar.com no search results page, the most prominent controls on the page are the links on the gray bar, shown in Figure 1.4, which include **Analyst Reports & Data**, **Morningstar Articles**, and **Tools**. At first glance, it seems that clicking these links would enable customers to browse content relevant to their query. Unfortunately, the links are actually *filtering* controls that serve only to further constrain a search that already has no results. Instead of providing a way out, filter links on a no search results page actually suck the customer deeper into the quagmire of the no search results state, now requiring the customer to click the Back button multiple times to get back to some useful content.



Morningstar is not the only site that is guilty of providing unproductive and confusing links that do not give customers a way out. Can you spot some useless filtering controls on the Endless.com no search results page shown in Figure 1.5?

FIGURE 1.4

Misleading links on the Morningstar no search results page

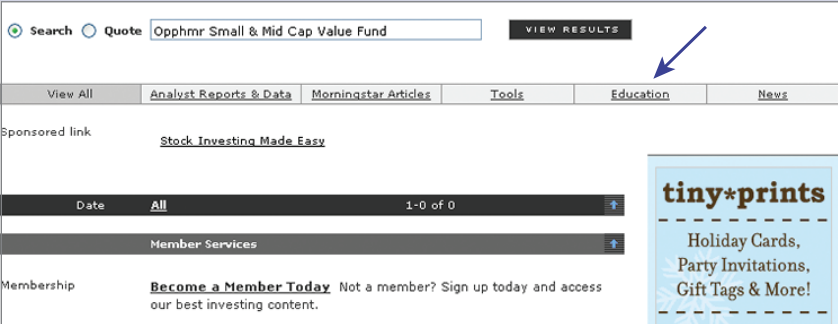
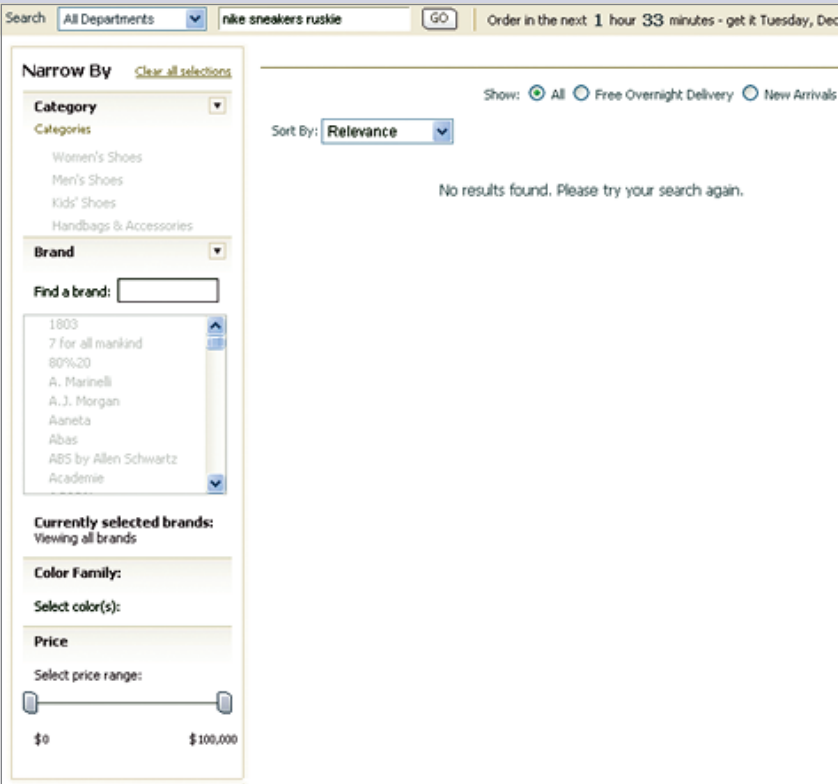


FIGURE 1.5

Endless.com no search results page



If you answered *The entire page except the search box*, you are correct. Further constraining already nonexistent search results to those with Free Overnight

Delivery, changing the sort order, narrowing by category, or using a fancy price slider would all yield an endless stream of no results. Such system behavior is confusing to customers who are already frustrated by getting no search results.

Instead of providing controls that further constrain a customer's query on your no search results page, make sure every control gives the customer a productive way out. Compare these endless nonproductive results to the page in Figure 1.6, which shows the result of the same query on Amazon.com.




FIGURE 1.6

Amazon.com shows customers the way to find Nike sneakers

Amazon.com's no search results page succeeds in a big way, because of its very prominent and useful links and content that helps customers find their way to Nike sneakers and get away from the no search results page.

CREATE A ROBUST PARTIAL MATCH STRATEGY

Many of the user research participants I've observed over the years have leaned toward over-constraining their queries by entering too many search keywords.

Note—After misspelling, over-constraining is probably the most common cause of no search results pages on ecommerce sites.



Because over-constraining search results is so common, partial match strategy deserves special consideration for most ecommerce sites.

Partial match is a strategy that shows some relevant content that matches a viable part of the customer's query, while also indicating which keywords were omitted from the search results. Showing partial matches is important because what people find often changes what they seek by giving them additional clues about the target content. Effective partial match strategy gives your customers the best chance to recover and succeed in finding what they are looking for on your site.

Adding more information to explain what you mean is also a common human response when you are feeling misunderstood. Thus, a natural human response to a no search results page is to add more keywords to the original search query. Unfortunately, for most ecommerce search engines, this strategy works poorly, because additional keywords only constrain the query further without fixing the original reason for getting zero search results.



Note—The problem of over-constrained searches is especially prevalent when people are searching for an item of particular interest to them, so you will not always be able to observe this behavior in a usability test with predefined search tasks and made-up items for which participants are to search.

This makes field studies one of the most instructive environments in which to observe how people interact with your no search results pages.

For example, during one field study I conducted for one of the top ecommerce sites, a particularly representative participant mistakenly searched for the book titled *Harry Potter and the Sleepy Hollows* and, as a consequence, received no search results. This person then attempted to recover from the no search results condition by adding more keywords—a misspelled author name *J.K. Rawlin*, thus, inadvertently multiplying the confusion. Unfortunately, the site I was studying did not, at the time, provide any partial matches to show which of the keywords was the problem. When this participant came up empty after several search attempts involving incorrect or misspelled keywords, she concluded that the site must not carry any *Harry Potter* books and moved on to another site to make her purchase, never coming back to the original site that failed her. It is interesting to note that this behavior is typical.



Note—In most studies I've observed, even after participants have figured out which part of the query they've gotten wrong, not a single person has ever come back to the site that originally failed to deliver partial matches that would in some way guide them to the product they were seeking.

Several participants have commented that the lack of guidance as to which keyword was the problem “made them feel stupid” and adversely affected their perception of the site. Thus, an effective partial match recovery strategy is critical for ecommerce no search results pages.

Few sites do a better job at partial match than Amazon.com. As Figure 1.7 shows, the Amazon.com no search results page shows not just one, but two partial match results that use some of the customer’s original keywords, while at the same time clearly indicating which keywords were omitted from the partial match query by using effective strike-through font.

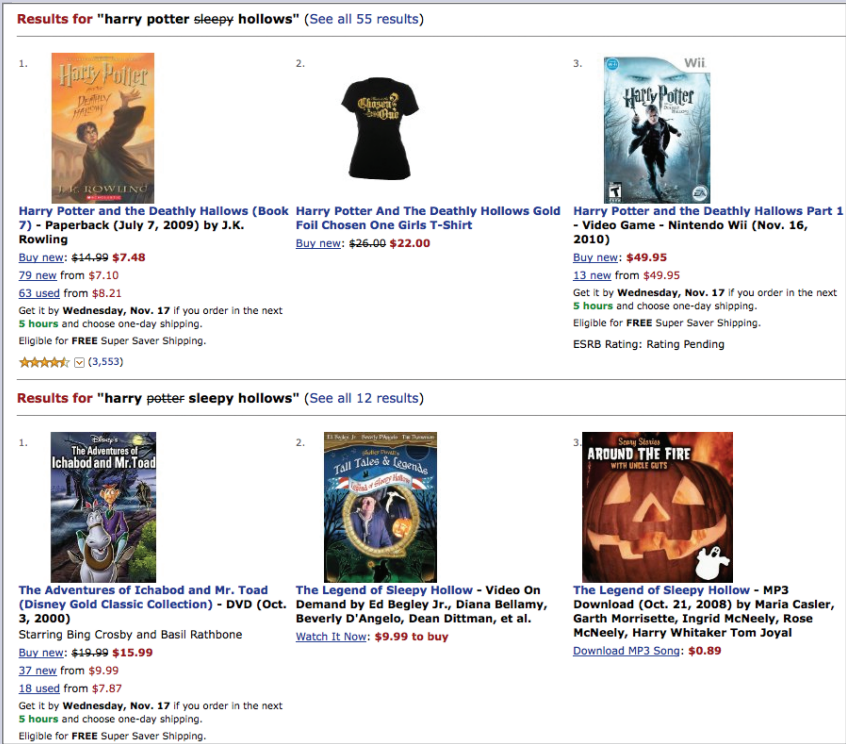


FIGURE 1.7

Amazon.com uses an effective partial match strategy

Not every company can afford to create a partial match strategy as robust as that of Amazon.com, but even a small step in the right direction can lead to a better customer experience. Often, over-constrained searches involve picking the wrong category in a filter set, so an easy place to start developing your partial match strategy might be with omitting a category (if the customer has selected one) from a query that led to a no search results page.



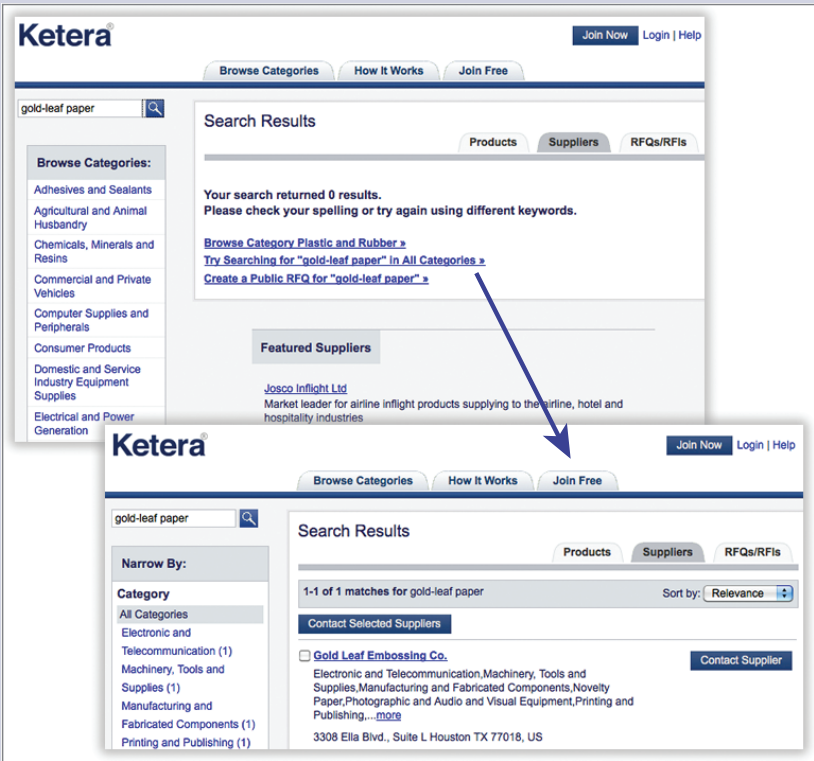
Note—Picking a wrong category need not be the result of any confusion about the category structure. It is very common for customers to forget to change the category when they switch gears and start searching for a different product or when they start their search via an external link that catapults them deep into your search results with a category filter applied.

Thus, a basic keyword-only search of **All Categories**, with no other constraints, should form a basis of your partial match strategy.

For example, the top of Figure 1.8 shows Ketera Network’s no search results page that resulted from picking the wrong category “Plastics and Rubbers” while looking for manufacturers of “gold-leaf paper.”

FIGURE 1.8

Ketera Network’s rudimentary no search results page



As a rudimentary recovery mechanism, the page provides a link that enables the customer to rerun their original keyword query in All Categories. Clicking this partial match search link yields the search results page containing “Gold Leaf Embossing Co” (bottom of Figure 1.8).

Although this approach works, it forces the customer to take action (click a link) without indicating in any way that a link would yield a useful result. A better experience would be to run the keyword-only query automatically on the back end, and, should it yield any search results, to display these results as partial matches within the body of the no search results page using a “More Like This” design pattern described in Chapter 12. Of course, the resulting page needs to indicate to customers clearly how the query was changed to create these results so they may respond accordingly.

EMPLOY MULTIPLE CONTENT STRATEGIES

Good content is a key to helping people recover from a no search results condition. In providing content for no search results pages, you need to keep in mind the customer’s original goal and stick as closely as possible to the spirit of the original query. Typically, in addition to the partial match discussed in the previous section, successful sites employ multiple content strategies. Normally, what is judged to be the most relevant content is provided first, followed by less relevant content shown lower on the page.

In addition to an “All Categories” partial match link, the Ketera Networks no search results page in Figure 1.8 contains a variety of simple but useful content designed to help the customer recover from a no search results condition. The page includes a link that enables customers to browse their original category, “Plastics and Rubbers,” and all of the 28 other categories of suppliers in the column on the left. Below links directly relevant to the customer’s query, the page displays a list of featured suppliers. If customers still can’t find what they are looking for, the page enables them to create their own Request for Information (RFQ) so the right services can find *them* instead. The Ketera Network no search results page is a simple, albeit rudimentary example of employing multiple content strategies to deal with a no search results condition.



Following is a list of some useful ideas for no search results content strategies, in order of their relevance to the customer’s original query:

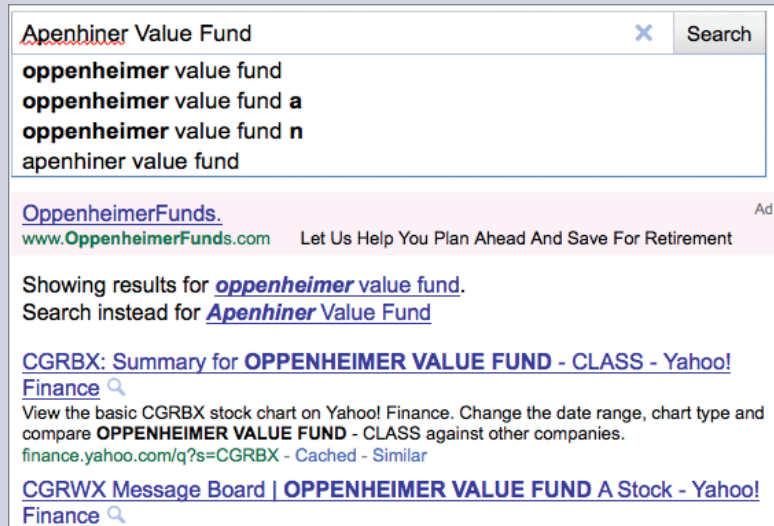
1. Spelling correction and substituting a customer’s original keywords with different keywords from a controlled vocabulary
2. Removing some of a customer’s original keywords, or making partial matches
3. Matching only categories or aspects, without the keywords
4. Top searches, featured results, or most popular results
5. Third-party resources and ads

You already looked at a good example of controlled vocabulary keyword substitution: the Google *Did you mean...* feature shown in Figure 1.1. To provide relevant content, Google draws upon its enormous list of the indexed keywords for which people have previously searched, which forms its controlled vocabulary, to suggest the closest matching available keyword.

Successful content strategies for handling cases when there are no search results are not limited to the five just mentioned. Google's brilliant auto-suggest feature, another industry-defining innovation shown in Figure 1.9, is an excellent example of a successful marriage of two content strategies: making partial keyword matches and using a controlled vocabulary for keyword substitution.

FIGURE 1.9

Google auto-suggest prevents a no search results condition from occurring



When automatically suggesting search keywords, Google chooses the top keywords from its controlled vocabulary of the most popular keywords. By matching the beginning of a string the customer begins to type with popular search keywords, Google ensures a successful search, forestalling the no search results condition before it ever occurs. In doing so, Google has devised another tremendously useful and industry-changing search tool. The point: Don't stick with just one specific type of search assistance. Instead combine design features and content creatively to meet the goals of your partner in conversation on the other side of the human-computer interface.

CASE STUDY: HOTMAIL NO SEARCH RESULTS PAGE

Now take a look at the Hotmail no search results page, following a scenario in which I'll try to look up Luke Wroblewski's email address by searching for any email messages I've recently sent or received from him. But I've inadvertently misspelled Luke's last name, so my search query is *Luke Wroblwski*. This situation is, unfortunately, all too common. (I personally experience it almost daily with people's creative spellings of my own last name.) Indeed, it is likely that millions of people experience the Hotmail no search results page on a daily basis.

Unfortunately, the Hotmail no search results page, shown in Figure 1.10, completely misses the mark, violating *all* the design principles discussed so far and thus, committing the three deadly sins of no search results pages, as follows:



1. There is no indication of the no search results condition. Instead, the central message seems to be **No message is selected**, which is technically true, but hardly useful in this situation, because there are no messages to select.
2. Instead of focusing on providing a way out, the page presents several completely useless and confusing controls—like the **Sort by** check box with no options and the **Change Your Reading Pane Settings** link.
3. There is absolutely no content that helps me reach my goal. The **Return to inbox** link and the various filter links on the left do nothing to help me connect with Luke.

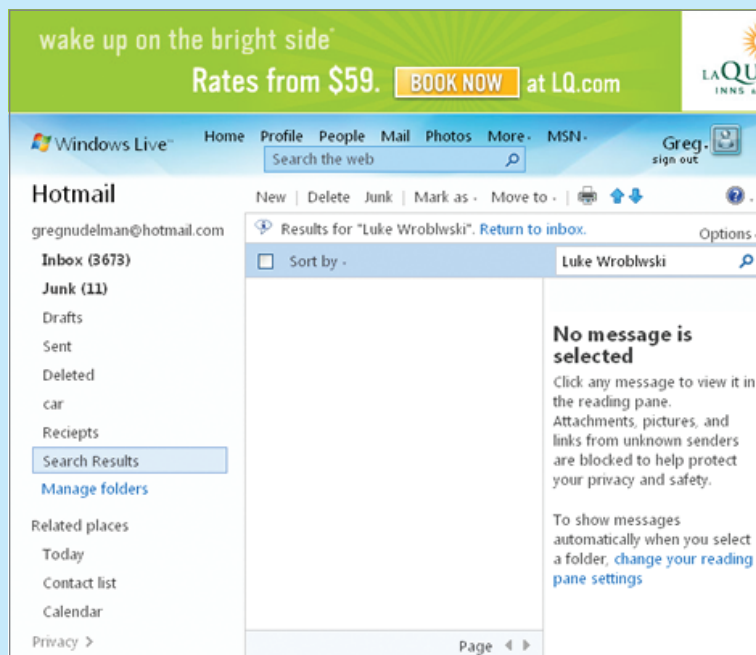


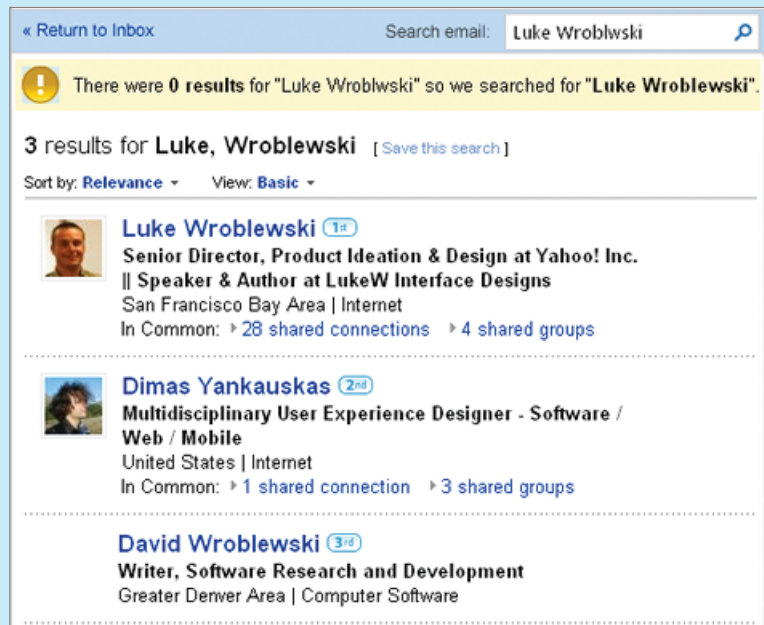
FIGURE 1.10

The Hotmail no search results page commits all three deadly sins

What can you do to improve this page? Obviously, removing some of the confusing controls and improving the no-results messaging would help. Further, you could think about the search problem creatively and try applying a social networking model to email. For example, if Luke and I have ever corresponded in the past, we could remain connected in the system, which would automatically remember Luke's email address. If we've never connected before, the system could helpfully offer to introduce me to Luke for a fee. Of course, we also need robust controlled-vocabulary matching, so the system always finds my existing connections first—no matter how badly I've misspelled their names. In Figure 1.11, my simple wireframe shows how a redesigned Hotmail no search results page might look. It employs the three design principles described earlier and uses sample content borrowed from LinkedIn.com.

FIGURE 1.11

Employing the three design principles to improve the Hotmail no search results page



By neglecting to give its no search results page the attention it deserves, Hotmail has quite possibly lost a massive opportunity to launch a social networking revolution. However, its loss resulted in a huge opportunity for LinkedIn, to the tune of millions of dollars.

No search results conditions provide a precious opportunity to go the extra mile to connect your customers to the content they are looking for. You can greatly deepen the customers' connection with your brand, create legions of loyal followers, and blow away your competition, but only if you pay attention to what your customers are trying to do and address their problems with creativity and the care they deserve.

HOW CAN YOUR SITE-SEARCH ANALYTICS HELP YOU?

Not sure whether people searching on your Web site are getting what they need from your search system? There's good news: Your organization is probably logging searchers' queries. Find that data and use it. It's a gold mine for user research and insights, and you *already own it*. Unlike other types of Web analytics, this is *semantically rich data*—it records users' needs, not just their actions. And most important, users can express their information needs in *their own words*.

You can benefit from this data in many ways, including the following:

- You can identify gaps in your offerings. Do people keep searching for a product you don't currently offer? Maybe you should reconsider your product mix.
- You can determine ways to improve your site's user interface design for search and search results, leading to better conversion rates. At minimum, you know how wide that search box should be.
- You can discover the points at which navigation and search fail and remedy the problems you identify.
- You can learn the tone and flavor of your users' language and make sure your content—titles, metadata, and copy—matches up well with their language.
- You can devise a variety of search-related metrics you may not have considered before. In turn, you can integrate these into your organization's Key Performance Indicators (KPIs), enabling you to do a better job of assessing how well your organization meets its goals. So, your managers should like site-search analytics.

Ironically, as useful as all this sounds, many organizations don't bother to analyze their users' search queries. They may just be ignorant of the potential of site-search analytics, or perhaps their query data is buried in some corner of their organization that seems too hard to reach. But it's likely that the biggest barrier to the adoption of site-search analytics is simply that it sounds like it would be hard to do.

That's a misconception: Even free analytics applications such as Google Analytics now provide at least minimal site-search analytics features. Spending an hour a month reviewing even the most basic reports can yield some great insights that can immediately help your organization—and ideally, pave the


way to going further with site-search analytics. Following are two types of reports that are great starting points:

- ■ **Most frequent queries**—Study this at least once a month to get a sense of what your users' most important information needs are—and how they change over time. Test those queries out yourself: Do they retrieve relevant results?
- ■ **Most frequent queries with zero results**—Assuming no results means failure, which kinds of queries go wrong the most? Is the content there, but mistitled or rife with jargon? Or do you need to create new content?

In all cases, start with the most frequent queries—the *short head* to the *long tail* you may have heard of—because frequent queries *really* occur frequently, much more often than you might expect. If you sorted all your site's queries by frequency and stacked them from most frequent to least frequent, you'd get a distribution that looks like that shown in Figure 1.12, which depicts the *Zipf Distribution*, in which a few frequent queries—the ones in the *short head*—account for a huge volume of all search activity.

FIGURE 1.12

Typical distribution of site queries by frequency



Named after economist George Kingsley Zipf, the *Zipf Distribution* curve describes just about every Web site's search activity and is at the core of site-search analytics. This graph makes a simple, but powerful point: *A few unique queries go a really long way.* So, if you start by analyzing short-head queries, your efforts can also go a really long way.

This point is so important I want to illustrate it in yet another way: through a *Zipf Distribution* rendered in text. Table 1.1 shows that the most frequent query on the Michigan State University Web site, *campus map*, accounts for 1.4% of all searches during the time period they studied. To reach 10% of *all* search activity on their Web site, look at only the 14 most frequent queries—not a big number. (In this table, #14 is *housing*.) To get to 20% of all search activity takes only 42 queries; 50%, 500 queries; and so on.

RANK	CUMUL. %	COUNT	QUERY
1	1.40	7,218	campus map
14	10.53	2,464	housing
42	20.18	1,351	webenroll
98	30.01	650	computer center
221	40.05	295	msu union
500	50.02	124	hotels
7,877	80.00	7	department of surgery

TABLE 1.1

The Zipf Distribution
of Search Activity.
Rendered in Text

Depending on your available time, your efforts to analyze search queries can scale quite nicely. Have only an hour this month? Analyze your top-ten queries; they'll account for a good chunk of all search activity. Have twice as much time next month? Keep going, picking off the most common queries, one by one. As long as you let the Zipf distribution drive your priorities, your work can have a huge impact.

Searchers on your Web site want to have a conversation with your organization. If people do not find what they need, they may go to your competitors' sites. If they had no alternatives and are stuck with you, their poor experience would make them unhappy and damage your brand. Through their search queries, people are asking you—in their own words—for products, information, help, and ideas. Site-search analytics are one of the best ways to understand what your customers are looking for and ensure you give them the answers they need.

—LOUIS ROSENFELD

REFERENCES

Gill, Judith. “Effective Internet Search: Search Help.” *SearchHelpCenter*, March 2004. Retrieved January 31, 2009.

Mondosoft. “Web Site Usability Metrics: Search Behavior—Search Trends.” May 2004. Retrieved January 31, 2009.

Nielsen, Jakob. “Search: Visible and Simple.” *Alertbox*, May 13, 2001. Retrieved January 31, 2009.

Nielsen, Jakob. “Eyetracking Research.” *Alertbox*, April 17, 2006. Retrieved January 31, 2009.

Nudelman, Greg. “The Timestamp-Based Caching Framework.” *JavaWorld*, January 3, 2005. Retrieved January 31, 2009.

Nudelman, Greg. “Improve the Usability of Search-Results Pages.” *JavaWorld*, January 23, 2006. Retrieved January 31, 2009.

Spool, Jared M. “Producing Great Search Results: Harder than It Looks.” *UIE*, July 09, 2008. Retrieved January 31, 2009.

