Recommender Systems in Commercial Use

Susan E. Aldrich

■ *Commercial recommender systems are deployed by marketing teams to increase revenue or personalize user experience. Marketers evaluate recommender systems not on their algorithms but on how well the vendor's expertise and interfaces will support achieving business goals. Driven by a business model that pays based on recommendation success, vendors guide clients through continuous optimization of recommendations. While recommender technology is mature, the solutions and market are still young. As a result, solutions are not fully integrated with other business systems and technology platforms. While the market is retail-focused today, interest and vendor offerings are rapidly expanding to other areas. Retail clients will drive social, location, and mobile enhancements.*

Recommender systems are a mature technology that is ubiquitous in commercial websites. They are deployed to support the end user who is selecting content from a website and to support the marketer who is personalizing customer interactions with websites, emails, mobile phones, kiosks, and even dialogue with customer-service personnel.

This article presents primary research from an industry analyst and aims to analyze and convey the state of commercially available recommender systems. The data was collected from vendors of such solutions and their clients. The research did not involve tests or benchmarks. The research into recommendation solutions began with preparing an evaluation framework that synthesizes the requirements and evaluation criteria identified through customer interviews and analysis of existing solutions. It describes roughly 150 requirements in seven categories: guidance and advice; recommendation structure; managing recommendations; integration; operations; vendor's development and maintenance; and product and company viability. Using the evaluation framework, leading recommendation solutions were analyzed, compared and ranked. Findings are summarized in Aldrich (2011).

This article describes the business models, components, and tasks of today's commercial recommender solutions; describes how these systems are deployed in practice; analyzes how recommendation solutions are evaluated by businesses; presents the current recommendation solution landscape; identifies the shortcomings of current solutions from a commercial perspective; and ends with some ideas of what the future might hold for recommendation solutions in commercial environments.

In the following pages, I refer to the company providing the recommendation technology as the vendor, the company implementing this technology on its website as the client, and the user interacting with the website to acquire a product or obtain a service as the customer or simply user.

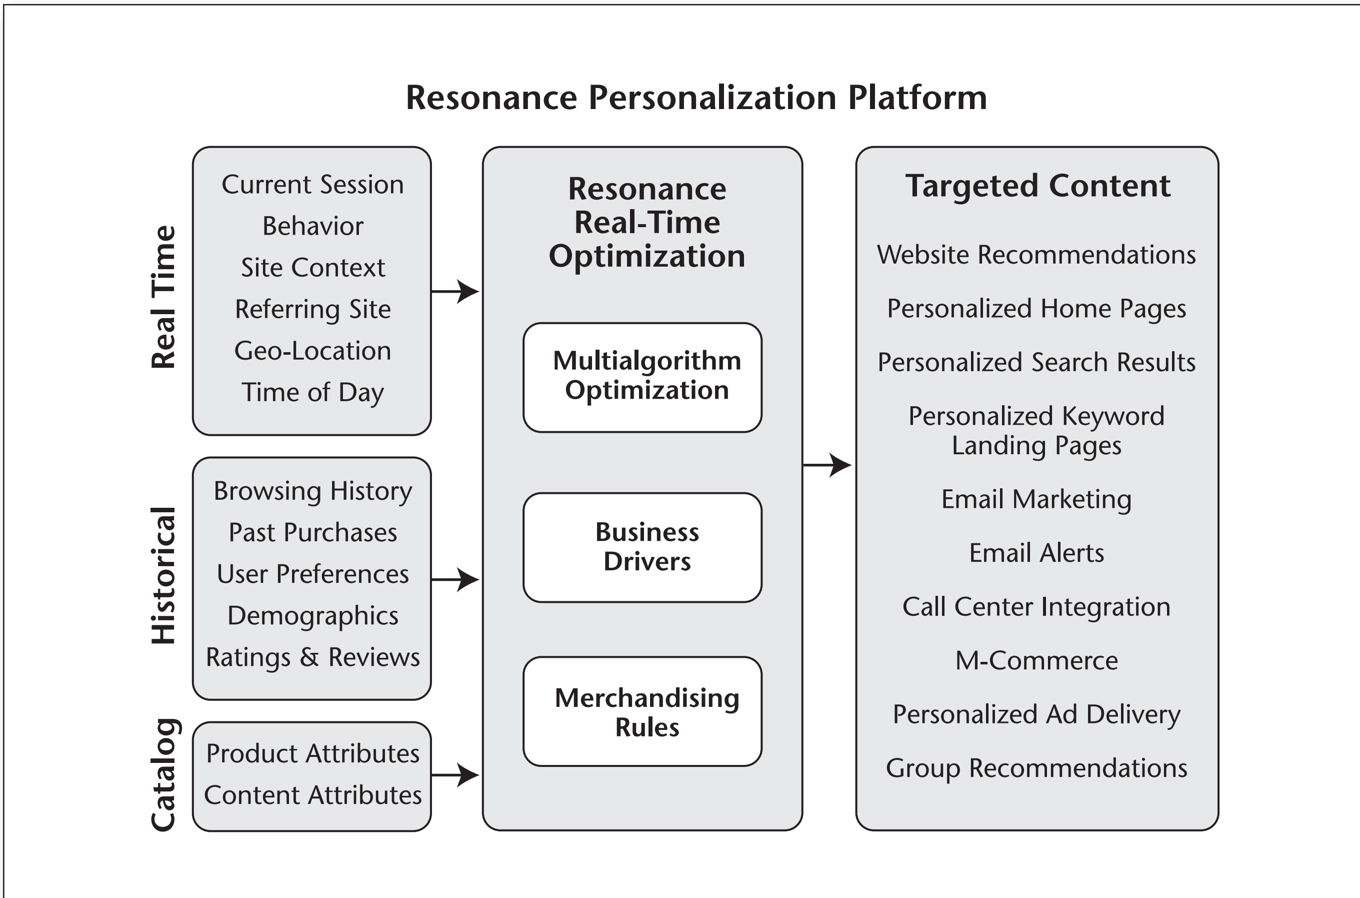

Figure 1. Example of a Commercial Recommender System.

This diagram represents Certona's Resonance recommendation platform, which is typical of commercial recommender systems in its high-level architecture. The left column lists the recommendation engine's data sources, which include current session context and actions; historic data; and product catalog data. The middle column lists variables used by the recommendation engine: algorithms, business drivers, and rules. The right column lists the types of content that might be recommended. Adapted from a diagram © 2010 Certona Corporation. Reprinted with permission.

Commercial Recommendation Solutions

Commercially available recommendation solutions all have similar compositions, as shown in figure 1. They also all have (1) proprietary recommendation engines and their data, hosted by the vendor; (2) services that receive and process bulk client data (such as customer profiles or order history, and store or website information such as the catalog); (3) services that receive the visitor context and requests, and return recommendations; (4) user interfaces for business people (usually marketing departments) to analyze and optimize results, add business rules, and control the recommendation engine; (5) rules engines, to allow clients to set constraints; (6) *A/B* or multivariate testing to verify how campaigns, business rules, and placement affect results; (7) templates that describe and organize recommendation strategies (for example, a product page template and a category page template); (8) reporting tools, with data exportable to comma-separated value (CSV) files; and, finally, (9) client care, provided by the vendor, to teach and guide the client in optimum use of the recommendation engine.

Commercial Recommendation Solutions in Action

All the recommendation solutions I've reviewed operate, at a high level, in a similar way. First they gather information about all users' interactions with all content, using data sent by the "listener" installed on each web page. They analyze and model the current user's activity. The listener data is collected and analyzed by the recommendation engine. Second, they obtain the user's information (online), usually through a cookie, or sometimes through user log in. Third, they apply any of a variety of algorithms to select the content to be recommended to a user, based on the client's specifications. Fourth they return the recommendations. Fifth and finally, they display the content.

Recommendation Solution Business Model

Recommender systems offered commercially today are generally offered as a service (software as a service, or SaaS). Each vendor manages the large amount of data involved (web visitor behavior, content catalog, customer profiles), while the vendors' clients manage the venues for recommendation display — typically web and email pages. Vendors typically price the service based on its success. In an e-commerce website, the retailer will pay a percentage of the incremental revenue derived from recommendations. In other types of sites, the equation is not as simple.

Because vendors rely on their clients' success for revenue (and for some portion of revenue growth), vendors typically provide excellent client care. This takes the form of initial deployment strategy, training, and then subsequent quarterly optimization reviews.

How Recommendation Systems Are Deployed in Practice

The initial deployment of the recommender system is generally swift (days or weeks). It will be focused on high value recommendation venues and low-hanging fruit. These are typically product detail pages, landing pages, and category pages for an e-commerce site; or article and topic pages for media sites.

Within two to three weeks of initial deployment, the client's marketing team will be working with the vendor's client-care team to improve the effectiveness of the recommendations. As the client becomes comfortable and knowledgeable, recommendations are extended into more web pages and then across more channels, such as email and mobile. Ultimately, recommendations will be served across the customer life cycle — not only during exploration and buying, but during postsales communications, repair and replenishment, and retargeting to bring the customer back.

The setup process for a commercial recommender system follows five steps.

First, establish a strategy for recommendations, based on business goals. For e-commerce, the goals are typically revenue and customer experience.

Second, determine where and in what way to apply recommendations. For e-commerce, the most important pages will be landing, category, and product pages, as well as follow-up communication emails. Vendors guide clients in selecting recommendation types (for example, people who viewed this bought that); establishing rules to constrain the recommendation engine (for example, prioritize red items before Valentine's day); specifying where on the page what type and number of recommendations will appear; and specifying the default for recommendations if the selected algorithm is unable to recommend items.

Third, provide information about content items, typically through a data feed or crawl. This information is typically updated daily or weekly by FTP.

Fourth, add a snippet of listener code (provided by the vendor) on each web page, email, and so on. This step is done by the client's IT staff, offline.

Finally, define, through templates, how many recommendations will be displayed; the web page coding defines where on the page the recommendations will appear. Vendors supply examples of the web page coding.

How These Solutions Are Evaluated by Clients

In the long-term use of a recommendation solution, algorithms may prove to be the biggest enabler or limiter — but it is not the algorithms that drive vendor selection. Business people want to hear how the solution will affect revenue and costs and how quickly they can achieve results. So they don't really press vendors who don't want to share details of their algorithms, exposing their intellectual property.

Moreover, it would be very difficult to determine how results would vary, based on the way vendors describe their recommendation structures and types; and studying the algorithms is not a great predictor of real-life results either. In fact, most of the time, the vendor's top technical talent and the client's most knowledgeable product manager won't be able to predict what recommendations the engine will come up with in a given situation, nor will they be able to predict how the website's visitors will respond to the recommendations. Will consumers buy more if shown "people who viewed this bought that?" compared to "people who put that in their cart also bought this?" Will they buy items with higher margins if shown viewed / bought compared to bought / bought? If the algorithms used to create the recommendations are changed slightly, how will that affect what consumers do? The only accurate answer to these questions is achieved by live testing with actual consumers.

Live testing is precisely how clients evaluate the effectiveness of a recommendation solution; they conduct their own tests once they have installed the solution, using the multivariate or *A*/*B* testing capabilities that all the solutions provide — and they generally discover that the recommendation solution produces great results, independent of vendor. Why?

First, clients have been using a very poor system — manual recommendations or no recommendations. It is very seldom that they replace an existing recommender engine with a new one or compare two systems to each other.

Second, the initial test is sure to employ behavioral recommendations, which are the most successful type of recommendation in commercial sites. Behavioral recommendation algorithms compare the current visitor's website behavior to the behavior of the thousands or millions who have preceded him or her, and based on past successful outcomes (for example, buying from a retailer or reading an article at a magazine site), select content (for example, products or articles) for the visitor to consider.

Behavior that is fodder for the algorithms includes the site that the visitor arrives from, search terms that brought the visitor, search terms the visitor has entered while on the site, the visitor's navigation steps, mouse clicks, elapsed time spent on each page, and mouse movements (Aldrich 2010a). Examples of behavioral recommendations include, "people who navigated our site the way you have, usually bought this," "based on the contents of your cart, other users think you will be interested in this," "people like you gave high ratings to this," "those who came here from Bing searching for cherry usually liked this item."[1]

Behavioral-based recommendations are roughly three times as effective as manually selected recommendations, according to a 2007 Avail Intelligence study.[2] This is generally in line with the experiences reported by many recommendation solution users. Each of the vendors on our list reports pretty spectacular success, including a 300 percent revenue increase,[3] a 150 percent higher conversion rate (Aldrich 2010b), and a 60 percent higher average order value.[4]

The decision criteria for selecting a recommendation solution, then, is generally less focused on the algorithms and more focused on vendors' geographic coverage, industry expertise (evidenced by target markets and clientele), and the mechanisms provided for optimizing recommendations.

But after initial deployment, and after a period of optimization, better algorithms become a competitive weapon for vendors and for their clients. As clients perfect their use of recommendations, they will see smaller and smaller improvements — unless research produces big breakthroughs in what recommender systems can accomplish. At some point, perhaps two years into their recommendation journey, clients will become anxious for recommendations that are better — new types, perhaps, and certainly better algorithms for the types already deployed. To keep their clients, or to steal clients from other vendors, vendors will have to improve their recommendations — an ongoing, continuous improvement process that is already underway. Vendors will find more heuristics to include in the recommendation analysis, and test better algorithms, and find ways to apply the best algorithms to each problem.

There is certainly an upper limit to the impact recommendations can have on average order value — shoppers have budgets. But there isn't an upper limit to the improvements in the customer experience from having a personalized selection of content throughout the dialogue, selection that is performed by the recommendation engine. Retailers' initial goals for recommender systems are revenue and conversion, but other industries deploy recommender systems to create a personalized experience. Those marketing teams are working to create a completely unique experience for every visitor, based on what that visitor cares about at the moment of his/her visit.

Recommendation Types and Structures

Ultimately, the flexibility of a recommendation solution rests on its algorithms, the vendor's willingness to prepare custom algorithms for the client, and the power of the management interface.

Custom algorithms are prepared by all solution providers, for clients whose requirements are not met by the vendor's standard algorithms. Most often, the customization involves incorporating additional item or customer attributes into the algorithm. For example, some sites use customer-profile information, stored in the customer database and associated with the visitor because the visitor has logged in or carries a cookie. Customer-profile data might include customer segment, purchase history, and preferences. Another method of achieving custom algorithms is by creating rules to constrain the basic algorithms. Most recommendations are created from a mix of behavioral data, content (item) attributes, and customer-profile data.

The solutions we evaluated varied in terms of the recommendation types their technology supports. Recommendation types can be described in consumer terms ("customers who viewed this bought that," "related items") or in terms of technical capabilities, such as whether the algorithms can associate people with items. Behavioral data such as mouse movements, time spent, clicks, search terms typed, and navigation path are also used to model a user's behavior, predict his/her intent, and calculate recommendations. Item associations typically make use of metadata (attributes) as well as behavioral data. See table 1 for a summary of vendor support for recommendation types and structures.

Recommendation Types and Structures	Avail	Baynote	Certona	Adobe	Rich Relevance
Recommendation Types					
Rules driven (choice of recommendations is entirely dependent upon a rule, for example, top-selling items in this category; no algorithm involved)	X	X	X	X	X
Rules modified (an algorithm selects recommendations, but rules constrain the recommendations that are shown; for example, always show recommendations from same category as item being viewed)	X	X	X	X	X
Recommendations for Facebook, Twitter, other social media (a visitor is matched with people with similar behavior, who are recommended as friends)		X			
Recommendations based on ratings or reviews submitted by user (user's ratings, which indicate likes and dislikes, are criteria for selecting recommendations)	X	X	X	X	X
Ad-word based (the ad the user has clicked on is factored into the recommendation; for example, "People like you who arrived by clicking on a Google ad for ... usually ended up buying one of these...")	X	X	X		X
Internet search term based (what the user has searched for is factored into the recommendation)	X	X	X		X
Syndicated recommendations (recommending items from another retailer's catalog)	X		X		X
Collaborative recommendations (a merchant recommends products that match the customer's interest at other retailers, based on cross-matching activities on both sites. Example: "People who bought a digital camera, also bought Adobe Photoshop at Software.com")	X		X		X
Recommendation Structures					
Item(s)-item(s) associations (one or more items is associated with one or more items)	X	X	X	X	X
Many items to item(s) (many items are associated with many items)	X	X	X		X
Person-person associations (a person is associated with one or more persons)	X	X	X		X
Item(s)-person(s) associations (one or more items is associated with one or more persons)	X	X	X		X
Heuristics based on behavioral data (search terms, navigation path, mouse clicks, time spent on pages, and other behaviors are used as input to algorithms that select recommendations based on outcomes of past visitors who exhibited similar behaviors)	X	X	X		X

Table 1. Recommendation Types and Structures.

Each solution supports a variety of recommendation types. This is due in part to the recommendation structures that are supported. The terminology presented in the table reflects the common usage in a commercial setting. "Collaborative recommendations," for example, in the commercial realm has a specific meaning that differs from its academic meaning.

Optimization and Marketer's Control of Recommendations

Clients using recommendation solutions don't establish recommendations on pages and then forget about them. They continually analyze and monitor the effectiveness of recommendations in meeting business goals, such as revenue, conversion, time on site, impressions. The three steps involved are as follows (Certona Corporation 2010b):

Step one: Track, analyze, and optimize the effectiveness of the recommendations. The business user interface provides reporting on the success of recommendations by page, category, and other filters.

Step two: Clients test the effectiveness of different recommendation strategies (combinations of algorithms and rules) by starting an *A/B* or multivariate test. Vendors provide industry and technical support for this activity.

Step three: At the conclusion of the test (based on time or traffic), the business user reviews the report that analyzes the results and selects the winning approach to implement — and possibly begins another test.

Vendors provide a management console, which is the user interface for continually improving the efficacy of recommendations on their web pages, emails, mobile interactions, kiosks, and so on. For maximum efficiency and effectiveness, marketers should not have to go through a middleman — a technical person — in order to add recommendations to a page, change the type of recommendation displayed, add a rule to constrain the recommendation engine's choices, and kick off a multivariate test to see which approaches are the most successful.

The more places recommendations are deployed, and the more strategies and rules are created, the more onerous the optimization task becomes. This is why automated optimization of recommendation strategies is a critical requirement for all but the simplest deployments. At this juncture, only Adobe,[5] Certona,[6] and RichRelevance[7] automate this task.

The optimization task starts with understanding what has happened, and it helps to be able to filter the data in many ways, such as by time, geography, customer segment, product category, and campaign. The solutions are currently too young to provide all the reports a business user can imagine, and they are too young to have yet invested in tools to enable the business user to create his or her own reports. Fortunately, all of the solutions allow export of report data to an Excel file; and all of the vendors will create specialized reports on request from a client.

Once the business person understands what has happened, he or she will want to try to improve results in the future, by selecting a different algorithm, displaying a different number of recommendations or displaying them differently, or changing the rules that constrain the algorithm. Then, the business person will start a multivariate test to determine the effectiveness of the new approach; wait some days; review the reported results and understand the meaning; and make yet another round of changes to improve results.

Present and Future for Recommendations

The recommendation market is currently dominated by retail, an eager audience for a solution with clear cost justification. Media sites, those that depend on advertising for revenue, have been the next wave of targets for recommender systems, followed by travel, B2B, telecoms, and advertising. Firms in many more industries are expressing interest in the personalization and revenue opportunities that recommender systems offer. Marketers want to apply the personalization capabilities of recommender systems to potentially all interactions with their customers. Keeping in touch with customers throughout the customer life cycle virtually requires a cross-channel approach. All of the vendors we've been tracking are making sure that they can present a cross-channel and cross-life-cycle story, either through built-in capabilities, add-on modules, or partnerships.

This widespread interest in recommendations leads to a dynamic vendor community. The barriers to entering this market are not high. The technology is not hugely complex, and clients are willing to engage with unfamiliar brands. The expense for the vendor of setting up the service can be eased by using hosting services such as Amazon's. So it is no surprise that there are dozens of recommendation solutions available, with new suppliers appearing frequently. Established vendors of marketing, search, and e-commerce solutions are adding recommendations to their portfolios; they are appearing in e-commerce platforms, site search solutions, and merchandising solutions. These vendors are not buying established recommendation solution providers; they are building their own or buying precommercial startups.

Shortcomings of Current Solutions

The solutions I've reviewed have shortcomings that are mostly due to the relative youth of the market. When any new technology arrives, it begins its life as an isolated tool; as it matures, it becomes integrated into business systems and technology platforms. Recommender systems are still in their isolation phase. Recommendation solutions are isolated from other business systems, not yet integrated with the marketing functions that complement them, such as advertising and campaigns.

The recommender systems are also isolated technically. They produce valuable insights that could create, for example, behavioral segmentation of customers or a new view of product categories. But this insight is typically digested by the recommender system, and never surfaces for use elsewhere. It would be useful if this insight could be exported to analytics platforms (such as a data warehouse).

Another failing of youth: the vendors don't have the ecosystem in place to provide all that their clients would like to consume. A developer ecosystem would help solve the integration shortfall and also provide additional applications. Currently, applications are very retail and web oriented, and a broader library of applications could address other industries and platforms. A library of e-learning modules is needed so that clients' staff (especially new hires) can learn the art and science of recommendations and how to use the interfaces provided by their vendor.

Future for Recommendations Market

The recommendations market will be growing rapidly for the next several years as it expands outside retail. But retail is the lead adopter, and retailers will be pushing the vendors to add new capabilities: social networking, mobile, and location based.

For retailers — and for shoppers — the ideal recommendation is, “people like you liked this,” or, even better, “your friends like this.” Recommendation solution providers are incorporating social networks into recommendations. Most already have some ability to base recommendations on a user’s ratings and reviews. Many are now experimenting with making recommendations based on what a user’s social network likes. In some categories, this is the only type of recommendation that matters — tell me what TV show my friends are watching, so we can talk about it at lunch tomorrow.

The customer’s location can be highly relevant to selecting the right content to show him/her: what tires to buy, what sports articles to read. I expect to see vendors adding algorithms this year that leverage location information, given the demand from their clients for this support.

Knowing the consumer’s location makes it possible to deliver offers to a mobile phone that can be acted on immediately, such as by stepping into the nearby store. Vendors already support recommendations for mobile and tablet devices; given the increasing consumer willingness to buy from mobile devices, it is likely this is the year that retailers will be using recommendation technology to improve search, browse, and buy on mobile phones and tablets.

But retailers won’t be driving all the requirements. Marketing executives will demand solutions tailored to their industries: content-oriented sites will be early on this list, such as magazines, investment services, and so on. I also expect that solutions tailored to a specific business goal — personalized experience — will be one of the next types of solutions to appear. Personalization has been the holy grail of marketing since the Internet hit the scene. It seemed briefly possible in the late 1990s, when it was the hottest buzzword on the web, but the buzz receded along with everyone’s hopes as it became clear that the technology wasn’t effective. The desire never died, though, and today’s commercial recommender systems solve the most intractable problem of the early personalization systems: mass personalization with decent performance and minimal manual effort. Recommendation technology is currently the most effective at delivering personalized experience to a large audience, and businesses recognize that personalized experiences deliver higher value than generic experiences.

Another arena that is ripe for development is the technical tool kit, targeted at IT. There are two reasons for this. First, the application of recommendations is potentially very broad, and IT will eventually step in and take ownership. Second, not all applications are well suited to the SaaS model. I have seen several deployments of recommendations into a customer-experience platform that is not web based. This environment requires a technical tool kit, not a business solution. Recommendations, as I observed earlier, have wide application, and IT will be the implementation partner that deploys recommendations into existing applications and systems.

Notes

1. See www.Avail.net (verified as of April 5, 2011).

2. See www.Avail.net (verified as of April 5, 2011).

3. See www.Baynote.com (verified as of April 5, 2011).

4. See Stalwart Communications, www.stalwartcom.com/clients_certona.php (verified as of April 13, 2011).

5. See Adobe Recommendations, www.Omniture.com/en/products/conversion/recommendations (verified as of April 5, 2011).

6. See www.Certona.com (verified as of April 5, 2011).

7. See www.RichRelevance.com (verified as of April 5, 2011).

References

Aldrich, S. E. 2011. *Comparative Rating of Five Recommendation Solutions*. Boston, MA: Patricia Seybold Group.

Aldrich, S. E. 2010a. *Baynote Collective Intelligence Platform and Recommendation Applications*. Boston, MA: Patricia Seybold Group.

Aldrich, S. E. 2010b. *Side-by-Side Comparison of Six Recommendations Solutions*. Boston, MA: Patricia Seybold Group.

Avail Intelligence 2007. *Avail Intelligence Customer Research, Business Case X*. Malmo, Sweden: Avail Intelligence.

Certona Corporation 2010a. *Resonance Personalization Platform White Paper*, 3. DelMar, CA: Certona Corporation.

Certona Corporation 2010b. *Resonance Insight User Guide*, 22–28. DelMar, CA: Certona Corporation.

Susan Aldrich is the research director for recommendations and site search for Patricia Seybold Group, a research and advisory firm focused on business and technology strategies to improve customer experience; and a principal of Greenhill Analysis, an information management and e-commerce research and consulting firm. Aldrich is recognized as a leading authority in the commercial applications of technologies for recommendations, site search, navigation, and discovery. She has consulted with Autodesk, Cisco, Comcast, Dell, HP, Philips, Symantec, Texas Instruments, Vodafone, and Xilinx, among others.