Speak<geek>
Tech Brief

2

Inside the RichRelevance Core Platform: dynamic ensemble learning

richrelevance®

Inside the RichRelevance Core Platform: dynamic ensemble learning

Even in a lackluster economy, online retailers can dramatically improve conversion rates by offering better, more personalized service and recommendations. A JupiterResearch study found that though online shoppers were planning to spend less in 2009, 61% could be persuaded to make purchases if they were given helpful product content that included recommendations.[1]

To deliver automated recommendations, most online retailers are leveraging recommendation technology that was developed in the early ’90’s and first popularized by Amazon.com a few years later. These systems are still evolving, promising higher conversion rates by predicting users’ behavior with increasing accuracy. However, not all recommendations systems are created equal. In this SpeakGeek paper we take you under the hood of the RichRelevance enRICH™ Platform, the core technology that powers the RichRelevance line of personalization products, to show how it delivers not only some of the most powerful purchase recommendations available, but also some of the most helpful, personalized guidance during every stage of the customer lifecycle.

A Solid Foundation

The enRICH personalization platform is built on an ensemble of more than 40 independent algorithms, each of which makes use of different kinds of user behavior and catalog data. The enRICH platform decides, in real time, which of these algorithms is best suited to meet a particular customer’s needs in a specific place at a specific time. In contrast, most recommendation systems leverage one highly complex algorithm for use across the entire customer base. Before we discuss how the enRICH platform operates on its ensemble of algorithms, let’s take a quick look at some of the most common single-algorithm approaches.

Collaborative Filtering. This approach offers recommendations to users based on the activity of users with similar histories. Most Internet shoppers have seen this approach at work in recommendations delivered in the form of, “People who bought X also bought Y.” Amazon popularized several variants of this approach in the late ’90’s.

1 JupiterResearch. Online Retail: Driving Relevant Experiences, February 23, 2009.

Recent Behavior Pattern Matching. This approach analyzes real-time user behavior in order to choose products to recommend. The premise of this approach is that the manner in which people browse through a site contains clues as to the products they are most likely to eventually purchase. The system then gets these products in front of them quickly.

Visual Similarity. This strategy uses the similarities of digital images to recommend similar products. Such an approach has the potential for quickly generating new lists of products that cut across traditional categories, such as a list of wooden products across Kitchen and Living Room categories.

Neural Network Modeling. In contrast with the three approaches above, which are built to use specific properties of specific kinds of data, such as page views, purchase history, and so on, neural networks are a more general approach designed to learn patterns in any kind of data. The downside of this approach is that it is difficult, if not impossible, to explain why any given recommendation was made by such systems.

A Recommendations System Should:

- Employ a variety of recommendation strategies
- Choose the right strategy, at the right time, in the right context
- Be employed sitewide from landing page to confirmation page, as well as beyond the merchant's site through advertising and email correspondence
- Continually improve by incorporating dynamic user and product data
- Rebuild its models many times per day to capture new trends and incorporate inventory data

The RichRelevance Approach

At RichRelevance, we believe that one single approach is not sufficiently powerful to make the best possible recommendations to all users in all contexts. People shop for washing machines in a fundamentally different way than they shop for DVDs — which means we have to recommend products differently. For example, we would encourage a person shopping for DVDs to buy two or three DVDs instead of one. That cross-sell technique would never work with washing machines, since people only need one.

Beyond using different algorithms in different places with different kinds of products, it is also critical that we use the right algorithm at the right time — and what is right can change in a matter of hours. For example, in a music store, recommendations for titles from similar artists often do well. But in the hours after news of Michael Jackson's death spread around the globe, a very simple algorithm that recommended the top selling albums of the past few hours (all of which were by Michael Jackson) converted more customers than any other. A few days later, this was no longer the case, and it was time to revert to what were normally the best algorithms for each context.

Additionally, not every one of the algorithms in the ensemble is capable of making good recommendations in all cases. For example, an algorithm that relies on past purchases has nothing to offer a brand new unrecognized user visiting a site for the first time. Our system allows individual algorithms to opt out of the ensemble when the appropriate data to support them is not available. The ensemble then relies on just the eligible algorithms.

Finally, because our ensemble chooses among individual algorithms that are kept separate we are always able to present transparent messaging about why specific products are being recommended, thereby driving up engagement dramatically. Users feel that we actually know and understand them, rather than assuming that we are dumping inventory on them. This is something that other types of ensembles, and opaque approaches like neural networks, cannot do.

With these concerns in mind, RichRelevance based the enRICH platform on multiple recommendation strategies, ranging from simple categorical top sellers, to collaborative filtering algorithms, to search and navigation based recommendations, to replenishment algorithms, to accessory algorithms, to highly specialized algorithms for individual merchant partners. In using multiple algorithms, the enRICH platform employs the basic principles of ensemble learning, a machine learning approach that combines many individual algorithms into a powerful overall algorithm.

*

The RichRelevance enRICH Platform is built on the principles of Online Dynamic Ensemble Learning, a patented recommendation system based on the tenets of Ensemble Learning, a proven, machine learning approach.

Ensemble learning is powerful, because it is free of the constraint of a single algorithm. However, prior to RichRelevance's innovations, most ensemble approaches were hampered by the fact that they used a collection of similarly structured algorithms to analyze structurally similar sets of data. RichRelvance, by contrast uses a wide variety of different underlying algorithms, each of which has the potential to look at very different data sets than the others. For example, one algorithm may look at the affinity between an individual user's clicks and purchases while another may look at affinity between his/her purchases and other users' purchases, while yet another may be based simply on past behavior.

Another weakness — which RichRelevance was also able to solve — of some early applications of the ensemble approach which is that they combine results from many sub-algorithms into one final result by weighting the scores of each sub algorithm and then adding them up. When groups of algorithms are combined in this way, the rationale for the recommendation can no longer be succinctly described to the user. It would be unreasonable to try to say something like, "30% of the reason you are seeing this product is because you bought a related product last week, 45% is because of the keyword you just searched for, and 25% is because of the last three pages you looked at." Instead, these systems have to rely on messaging that's more generic and therefore less effective. RichRelevance, by contrast is able to message recommendations clearly and provide the best products to recommend through the most effective algorithm strategy to a particular customer at a particular time.

Next-Generation Ensemble Learning

The RichRelevance enRICH platform innovates on traditional ensemble learning approaches: rather than trying to blend separate results together in an overarching result, the enRICH platform selects a single algorithm for each recommendation placement, based on its performance. Utilizing this data (updated in real-time), the 40+ individual algorithms constantly compete for dominance. The winning algorithm produces the highest level of user responsiveness, thus allowing the platform to always select the highest-performing strategy for each unique placement. On many pages, there are several recommendation placements, which allows us to display a combination of high performing messages based on high-performing algorithms.

We call this approach Online Dynamic Ensemble Learning. Because only one algorithm is active for a recommendation at a given time, messaging is presented to the user in clear, transparent terms that inspire trust, the necessary ingredient in encouraging purchase behavior. Also, because the platform gathers a wide variety of data from its ensemble of algorithms, and because the algorithms operate independently, the platform enables a variety of RichRelevance applications beyond on-site product recommendations. For example, ClickSee leverages user data in an interface that shows users up to 100 relevant products per click during a search; RichReach delivers recommendations via ads to your most valuable customers even if they leave your site; and RichMail leverages data from the enRICH platform to deliver highly personalized email.

How the enRICH Platform Works

As mentioned above, most recommendation systems deliver the same recommendation type (typically purchase-related recommendations) on limited numbers of pages. The enRICH platform, in contrast, offers recommendations that provide navigational assistance across every stage of the customer's process — offering many recommendation types that directly accommodate the customer's current behavior and location.

To illustrate this process, let's follow a hypothetical user on a short trip through a site.

1. A customer arrives at the site's home page. Because the customer has made purchases from the site before, we can apply algorithms that make recommendations based on past purchase behavior. (Other more generic algorithms, like top sellers, are also available. But on the home page, we know that personalized algorithms generally do better, so we use those. If the user had no purchase history, some of those personalized algorithms would be ineligible, and we might fall back on more generic algorithms.)

2. The platform looks at a sample of other users viewing the home page, and narrows down the list of algorithms to those that have proven successful, as measured by the correlation of each algorithm with the merchant's chosen criteria, such as the number of clicks or purchases.

*

Customers Gain Powerful Results from the enRICH Platform

Burton Snowboards
burton.com
73% year-over-year sales growth

D.M.insite
dminsite.com
15% increase in sales-per-visitor

Wine.com
wine.com
26% increase in revenue

3. Next, the customer does a search, which takes them to a search results page. Now several algorithms that rely on recent search keywords are also eligible to appear. They compete with all of the other eligible algorithms for placement on the search results page.

4. Next, the user clicks through to a product page. Algorithms like collaborative filters or accessory strategies that are based on viewing a specific product are now eligible and compete for placement with all of the others.

5. Finally, the user puts the item in their cart and proceeds to Checkout. At this point, cart-based algorithms are also eligible.

On all of the pages above, as the customer moves from page to page to page, we choose the most effective recommendation algorithms, allow each algorithm to recommend the most appropriate products for the user based on who they are and where they are, and then label the products recommended by each algorithm with a clear and direct message.

Learning from Experience

Over time, the system learns to make better recommendations in two ways. First, the individual algorithms are exposed to more and more of the underlying data they rely on, whether it be clicks, cart adds, page views, searches, etc., which allows them to get progressively better at recommending products. Second, the ensemble as a whole learns more about which algorithms work best in each context. This learning and adjustment occurs continuously, and the platform is rebuilt several times a day to accommodate such developments. This constant optimization allows the ensemble to react to both relatively short-term changes, like the previous Michael Jackson example, and longer term macro trends in the marketplace.

The Bottom Line

The enRICH platform's unique architecture enables us to provide merchants with immediate, tangible feedback. After deploying the enRICH platform, retail customers report improvements across a range of KPIs, including increased conversion, revenue, and repeat visits. And because the platform self-optimizes, these results can only improve with age. Visit www.RichRelevance.com to learn more.

richrelevance®

www.richrelevance.com

415.956.1947
info@richrelevance.com

275 Battery St. Suite 1150
San Francisco, CA 94111