



RESOLVE DIGITAL

Building the Smart Shop

How machine learning can help improve customer experience, increase your conversion rates, and boost sales across channels.

Overview of Machine Learning

Learning from Amazon and Netflix

What if you could understand your customer's behavior based on her past purchases and browsing patterns? What if your site could predict which products she's predisposed to buy?

We know that if you show your customers the right content at the right time they're more likely to take the action you want: Click on that "add to cart" button and complete a purchase.

Amazon.com and Netflix have been doing it for years. It's called "predictive analytics" and is one of the many benefits machine learning can provide your eCommerce business.

The fact is every visitor is different. Machine Learning (ML) algorithms continuously adapt to each customer's unique behavior to provide relevant, *personalized* product recommendations.

This tailored approach improves each customer's experience. Up-selling and cross-selling results in higher conversion rates and larger average order sizes.

Can Machine Learning help *my* site?

You don't have to be a Fortune 500 company to benefit from machine learning.

To make the investment worthwhile, you do need a certain number of products and volume of visitors. The more history you have, the better. Ideally each product will have been purchased and/or rated at least ten times.

ML is especially effective on sites with at least a few hundred products, where subtle differences between them may not be immediately obvious.

It comes down to this: can visitors easily discover products they want or do they need help finding new merchandise?

Cheese.com is one site that could greatly benefit from a machine learning recommendation engine. With over 700 cheeses from around the world they offer a huge selection. But not everybody knows the difference between *Double Gloucester* and *Double Worcester*. Even their most cultured customers may not be familiar with *Abbaye du Mont des Cats*.

A recommendation engine sorts through all the options to deliver specific product suggestions for each customer based on their purchase history and preferences.

The human touch is still important.

Machine Learning by itself is not enough. Human intervention and supervision is required.

To achieve the best results you should set clear and measurable business objectives. You need to know what's most important to your customers and how they make decisions.

Armed with that knowledge, we develop logic and set filters to achieve specific goals for your site. What works well for one business may not translate to another.

Examples of Machine Learning in action

- Promote the highest margin products to increase revenue
- Send personalised recommendations via email marketing
- Use recommendations to tailor search results
- Up-sell through the check-out process
- Integrate with your CRM to provide real-time suggestions for your sales team

Case Study: United Cellars

United Cellars is an award-winning online wine retailer based in Australia. Their site is built on the Spree Commerce platform and supported by a well-trained telemarketing team.

We've been working with United for over two years on different aspects of web development. Recently we've begun to focus more attention on conversion optimization. The site currently has over:

- 16,000 product views
- 60,000 orders
- 3,000 reviews

With this valuable data we knew we could make quality product recommendations using Machine Learning algorithms provided by Prediction.IO - an open source prediction server.

Getting started with a recommendation engine

Because we have a long history with United we were already familiar with their business goals and customer behaviors. While everybody appreciates a good deal, price is not the main motivator for United's best customers. They're looking for wine they haven't tried before that's similar *in style* to vintages they've already purchased.

With this in mind, we set up a recommendation engine and A/B testing so we could compare the results to the baseline. We ran the test until we had statistically significant data.

We spent one week getting everything set up and running smoothly.

- Prediction.IO server setup on Amazon Web Services
- Push ratings, orders, and product views data to Prediction.IO server
- Connect Prediction.IO server to Ruby Prediction.IO SDK running in the store
- Retrieve recommended products
- Display recommended products on the site
- Set up A/B test

Results of the A/B test

We ran the first test for 15 days. 50% of the traffic saw the Prediction.IO recommendations; 50% viewed the original site. The test demonstrated

- 45% longer average session
- 22% increase in conversion rate
- 37% increase in average order value
- 71% increase in revenue

The significant revenue increase is due to a *combination* of higher conversion rate and bigger average order value. The results were so compelling it was a no-brainer to start showing the recommendations to *all* visitors.

More optimization opportunities

In addition to on-site product recommendations, Prediction.IO unlocks many more possibilities. To build on their initial success, United is considering

- A tool for the telemarketing staff to offer customers personalized suggestions over the phone
- Email marketing campaigns that send individualized recommendations to each recipient
- Applying the recommendation engine to customize search results

Learn more

We hope this glimpse into the potential benefits of Machine Learning will inspire you to think about your site in a different way. If you're really excited about all the possibilities, Stanford University offers a free course on ML available through [Coursera](#).

But if you'd rather focus on running your business and serving your customers, contact us about implementing a machine learning strategy on your site.



Barry Harrison

Managing Director

barry@resolvedigital.com

Skype: barryhar

David Jones

Technical Director

david@resolvedigital.com

Skype: mrdgjones

