MACHINE LEARNING AT 4R

What It Is and How It Can Help Retailers

Analytics in Retail over the past couple of decades has grown from better predictions to better decisions. This evolution has been both motivated and enabled by the confluence of three factors:

1 *The ability to ingest and process formidable amounts of data*

2 *The availability of increasingly sophisticated algorithms*

3 *The cut-throat pursuit of profitability and market share*

These are not unrelated factors, though. Seeking that 1% or 2% of incremental profit has sparked the need to better understand the information contained in retail transaction data and use it to facilitate better decision making in functional areas ranging from planning to assortment, to inventory management throughout the supply chain.

In the Retail industry necessity is truly the mother of invention, where increased profitability has been the reward for the more technology savvy retailer which have adopted the right science in the right places, at the right time.

But what exactly is this science that we all hear so much about?

First, a bit of historical background: few people might know that Machine Learning has been around for almost a century and the current state of the discipline reflects the advances in data processing and algorithmic sophistication over time. About 80 years ago a British statistician by the name of Ronald Fisher (unrelated to Dr. Marshall Fisher, founder of 4R Systems!) applied a -back then- revolutionary classification technique to identify which of the three types of iris flower each observation in a set of 50 belongs to. While this instance of early Machine Learning application went on to become a famous case study for generations of students and practitioners, its importance lies in ushering a new way of thinking aimed at applied a collection of techniques to achieving tangible benefit in a variety of scientific and industrial settings. These techniques are today referred as Machine Learning.

Fast forward almost a century. Today ML is ubiquitous in all fields where massive amounts of data contain potentially valuable information. ML per se is not a single trick or algorithm. It is instead a variety of techniques that fundamentally do different flavors of one thing: predict the value of a dependent variable (the "response") based on the values of one, two, or millions of independent variables (the "predictors").

These dozens or more different ML techniques calculate this prediction for different needs and frameworks used for training the prediction model.

The response variable can be numeric (your FICO score is calculated using a number of behavioral financial

attributes and reflects your credit worthiness), or categorical (character recognition, where an image is mapped to a set of known characters, i.e., classification).

The training environment can be supervised, when for every observation used in training, i.e. calibrating, the model we know the true value of the response variable, such as character recognition. Unsupervised is when we don't know the true value of -or there is no "true" value for- the response variable, such as when we try to segment a population based on demographic data and purchasing habits so that a retailer can fine-tune their promotional activity to better reach their target market.

So, what can ML do for a retailer? The areas of application are numerous: targeted marketing, customer recommendations, dynamic pricing, streamlined logistics, etc. At 4R Systems we use ML in a variety of applications ranging from inventory to assortment to markdown optimizations. We employ ML in a very similar way as it is used in the context of autonomous driving (AD). Contrary to popular belief, it is not a "learning" algorithm that decides -for example- to stop at a red light and go with green. More than 90% of all ML in AD is used to read and interpret the surrounding environment and make better predictions about a large number of environment variables. From the simple question of classifying a light as red/yellow/green to more complex questions such as "is that object on the edge of the sidewalk a pedestrian about to jay-walk or just the shadow projected by a lamppost?" In other words, ML is not used to make the decision per se, it is used to gather better information that will feed into a decision engine that makes the decisions. The decision engine itself is ML-free or ML-light at best. And it makes sense that it be this way: we don't want a learning algorithm to figure out that you stop at red and go at green. The learning process would be very costly!

Similarly, ML is used at 4R Systems to "read" the retail environment and make more accurate predictions, which will in turn feed to a decision engine. Unlike AD, though, the decision engine is not based on hard-coded rules but on a different set of algorithms known as stochastic optimization. Stochastic because predicting demand is not about "guessing" correctly what product the customer will buy tomorrow in store XYZ, it is instead about describing as accurately as possible the probabilistic nature of the underlying demand process, which is inclusive of seasonal profile(s), trend, demand modifier (promotions, demand shifts, weather events, etc.) plus a random component, i.e., unexplained fluctuations.The optimization part comes into play when seeking profit-optimal decisions given the environmental variables mentioned above.

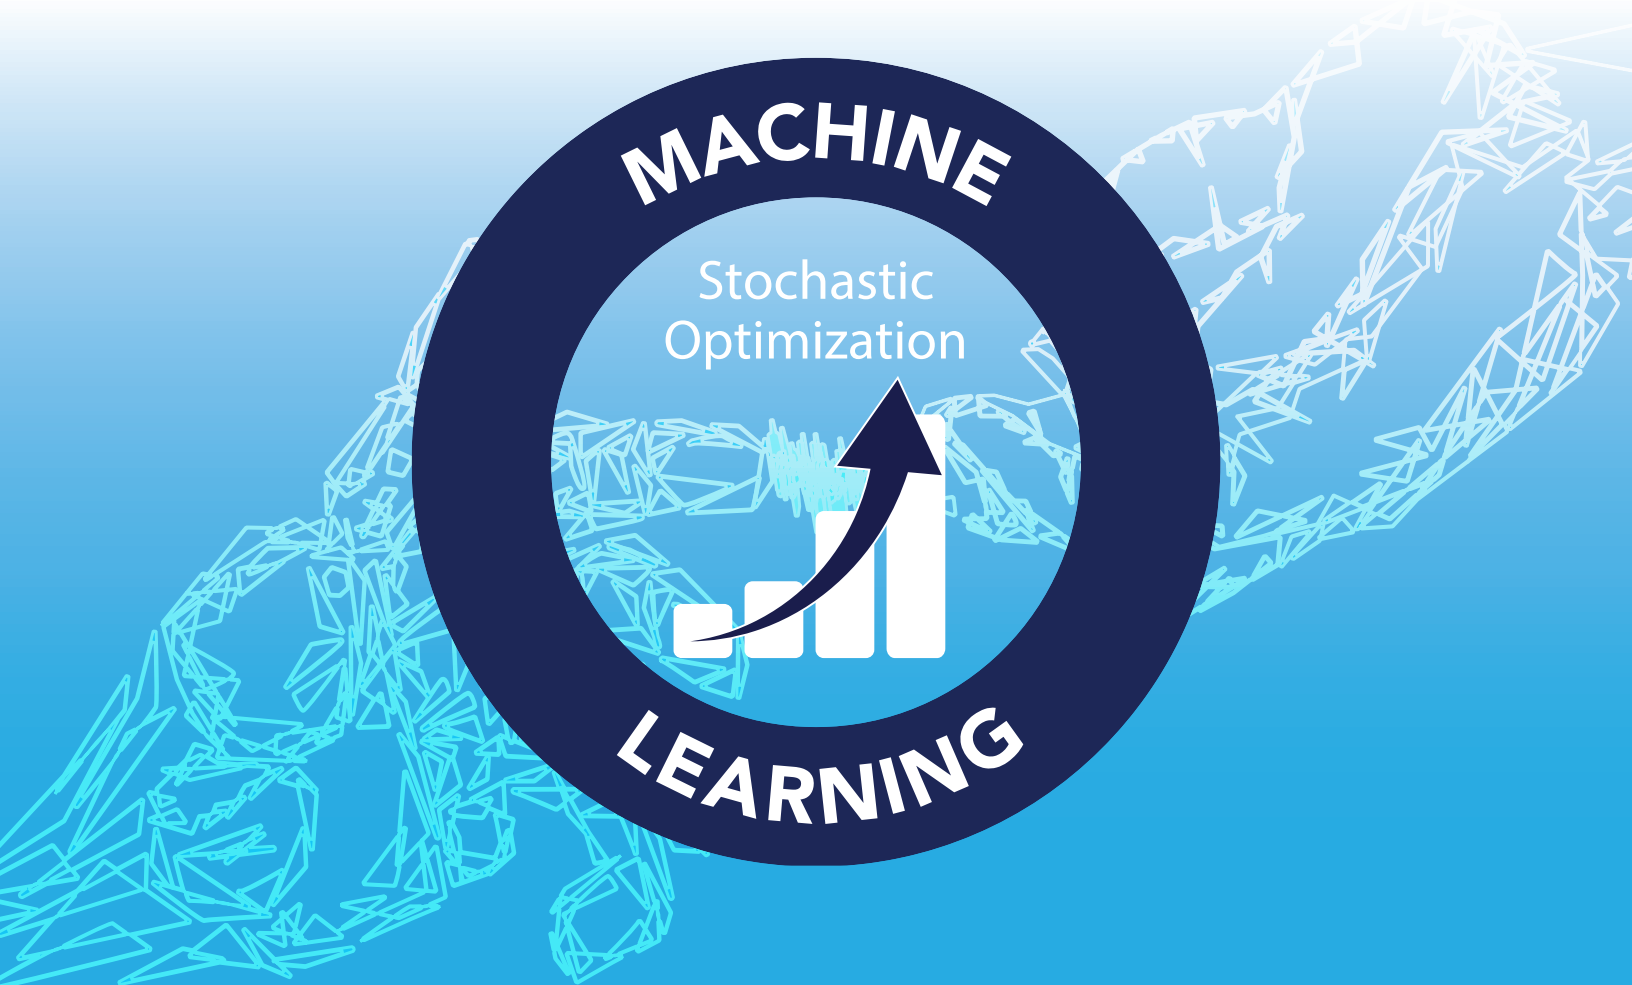

Examples of specific areas where ML is used in the suite of 4R Systems' solutions include:

- Unsupervised Clustering for creation of seasonal profiles
- Neural Networks for predicting demand response to TPR promotional activity
- Classification for outlier filtering
- Non-linear multiple regression for BOGOs and other structured promotions
- Pattern Recognition to identify out specific demand patterns such as pillows selling in pairs, tires selling in 1,2 and 4 only, etc.

At 4R Systems Machine Learning is not a buzzword that we throw around to build a veneer of credibility. It is a methodically researched and carefully applied set of techniques that greatly improve the accuracy of the predictions used to make optimal inventory decisions. Has someone else built a better mousetrap? We can't tell for sure, but ours works pretty well and has demonstrably generated millions of dollars of incremental benefit for our clients over the course of many years. And the process keeps getting better as we continuously find more and better ways to apply and integrate ML across our set of solutions.

4R Systems
801 Cassatt Road, Suite 202
Berwyn, PA 19312 USA

+1 610-644-1234

info@4rsystems.com

US 01/10/19