WHITE PAPER

APPRISS® RETAIL

Machine Learning and the Evolution of Exception Reporting

David Speights, PhD,
Vishal Patel, PhD,
Alan Gillette, PhD,
Sruthi Shyamsunder, MS,
and Dr. Adi Raz,
Appriss Retail.

In collaboration with Chris Hoffman,
Formerly Director of Corporate Investigations
& Loss Prevention Systems, The Bon-Ton

Introduction

Recent years have seen an explosion of consumer products that incorporate artificial intelligence (AI). Self-driving cars, Alexa, Siri, chatbots, Google search,and smart traffic lights are examples of how artificial intelligence is impacting the world. While artificial intelligence has gained mainstream attention recently, AI applications have been "hidden" for many years in a variety of business applications such as: credit scores, credit card fraud detection, direct marketing, identity theft protection, oil exploration, and more.

Retail loss prevention has also seen applications of AI in technologies such as facial recognition and return fraud authorization. In this white paper, we extend the loss prevention professionals toolkit and show a machine learning algorithm (a specific type of AI) developed to help retailers detect fraudulent and abusive employee activity.

The ultimate objective for developing machine learning models is not to replicate a single retailer's successful investigations. Rather, it is to combine many learnings and obtain collective intelligence to curb employee deviance. Appriss Retail serves more than 300 retailers worldwide, and this experience provides a solid foundation for such development.

"In our testing, we saw an increase in fraud cases identified and a decrease in time required to track down the fraudulent activity within many of the cases."

—Chris Hoffman

APPRISS® RETAIL

From Rules and Queries to Predictive Models and Machine Learning

Today, most Loss Prevention departments use traditional exception based reporting (EBR) tools to monitor their employees. Exception based reporting is usually viewed as a query run on point-of-sale (POS) data that flags suspicious activities. EBR tools used by retailers are similar to the tools used by the credit card industry in the 1980s to detect fraud. In the early 1990s, the credit card industry mostly abandoned the exception based reporting approach in favor of methods of detection based on predictive modeling and machine learning (see timeline below). A similar trend was seen with healthcare fraud about 15 years ago when the industry's tools converted from rules and queries to predictive modeling. At Appriss Retail, loss prevention solutions are evolving in a manner comparable to credit card fraud detection.

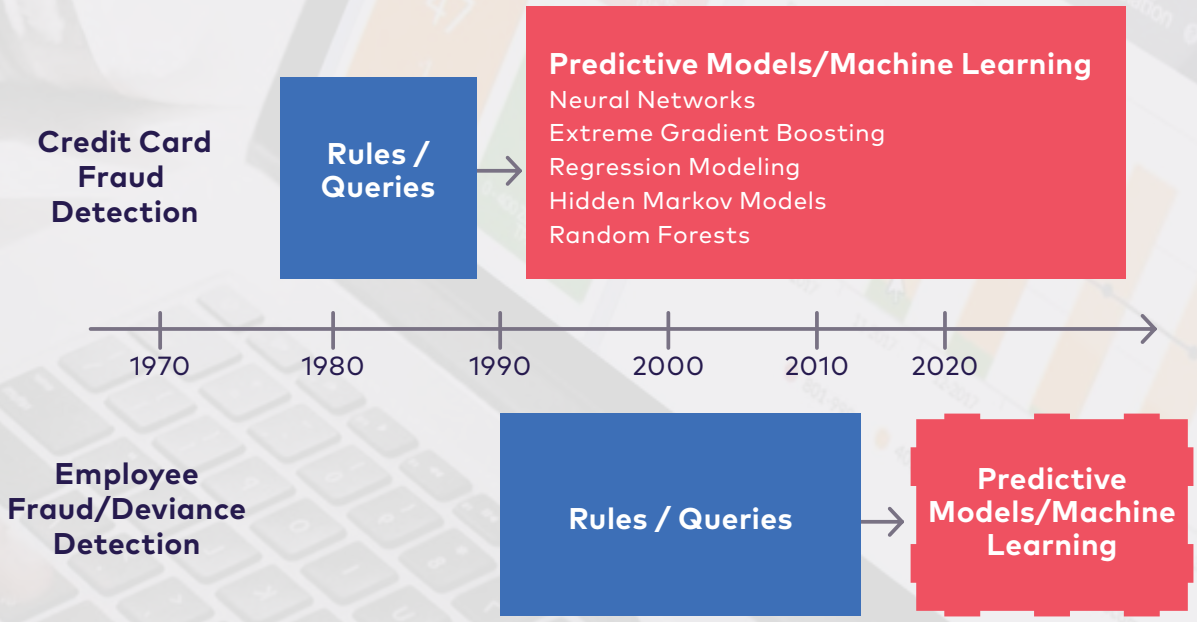

"Appriss Retail's machine learning-based solution has shown a significant increase in the quality of fraud cases we can identify and the efficiency with which we are able to track down the source of fraud."

—Chris Hoffman

A traditional exception based reporting process follows the model seen below. A query is run to look for anomalous behavior and its output is provided to an analyst in the form of a list of transactions or employees to investigate. Then, the analyst or investigator reviews the exceptions presented and determines, sometimes based on looking at additional information, if the transaction or employee in question warrants further review. Some of the exceptions will result in an action from the retailer against the employee such as investigation or termination of employment, which we label as "success" below.

The goal of machine learning is to apply a complex mathematical model to the data and learn, using thousands of variables, the most useful elements in an employee's transactions for identifying cases that are most likely to result in action by the retailer. With machine learning, the top ranked exceptions will have a much higher rate of success than those that are generated from traditional EBR (see the following diagram). Over time, an approach with a higher success rate will make loss prevention departments far more efficient and will decrease employee fraud rates.

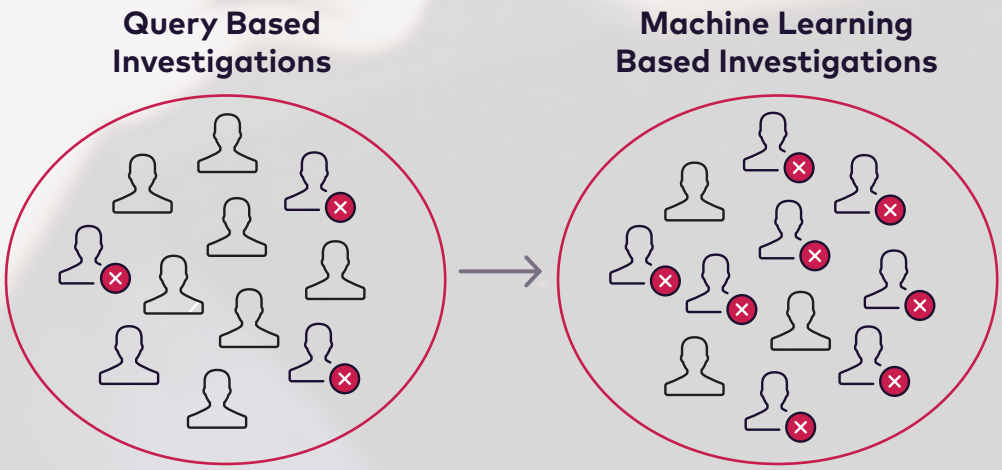

APPRISS® RETAIL

Overview of Artificial Intelligence and Machine Learning

Artificial intelligence refers to intelligence exhibited by machines. It can cover a wide range of tasks such as authorizing a loan application or automatically driving a car. The diagram illustrates some of the components of AI that are used frequently. Note that AI is the broad category that contains the other topics.

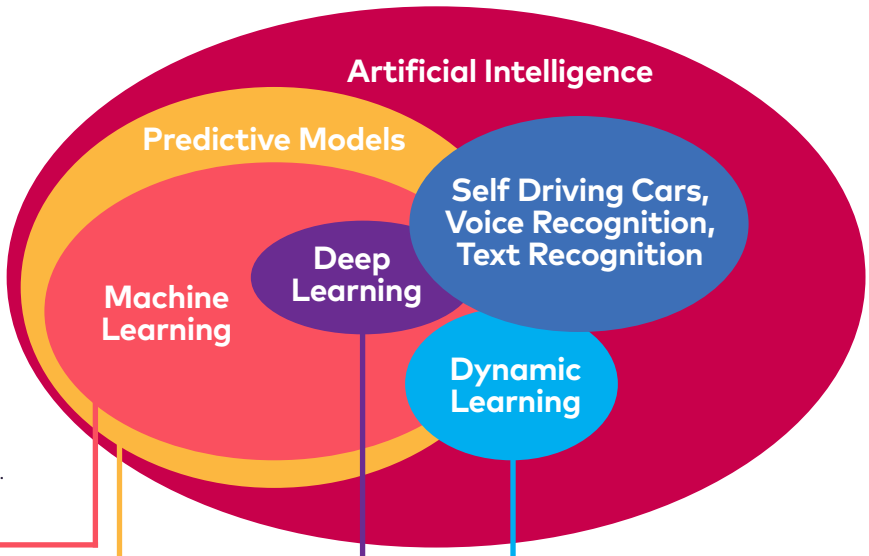

Predictive Models and Machine Learning:

Predictive models and machine learning are shown in the diagram as almost completely overlapping circles and are very close in definition and purpose. The primary difference between the two is machine learning originated in computer science; predictive modeling was developed in the field of statistics. In the past decade or more, the two fields have essentially merged and now share a highly overlapping definition. Both are concerned with using data to predict or model outcomes, and both are tools used in artificial intelligence.

Deep Learning:

To model relationships, a mathematical equation is defined that has variables that are known and parameters that are unknown. The unknown parameters are estimated using a mathematical optimization process to best fit the data. Typically, variables used in machine learning models are defined by a person as candidate variables for a prediction, and are derived from a raw set of data. For example, the raw data may be something like the point-of-sale data for a retailer and an example of a derived variable would be the number of line voids in the past three months. Deep learning is a process that uses highly complex model structures to learn patterns in the raw data without any human intervention, thus eliminating the need for derived variables. The difficulty in deep learning is that the more complex models are often more difficult to fit accurately.

Dynamic Learning:

There are two very different types of model categories discussed in this section, static models and dynamic models. A static model is fit to a static pattern in data. The predictions from a static model can change as the data it is presented with changes, but if the model is presented with the same exact data, it will produce the same exact answer. By contrast, a dynamic model's prediction can change from day to day even if it is presented with the same information. Dynamic learning is a process by which models are re-fit or updated dynamically as new data presents itself. This approach is useful when a system is continually bombarded with additional information and the relationships in the data are dynamic and the model itself needs to change to adapt.

Collective Intelligence

Machine learning allows developers to aggregate deidentified and anonymized data from many retailers to detect fraud trends and create a more powerful solution. (Data from multiple retailers is not comingled.) Each retailer is unique in the exception-based approaches and investigation techniques that result in employee investigations, terminations, or prosecutions by the retailers. Machine learning collectively learns the most successful strategies. The table below highlights the general structure across a set of retailers.

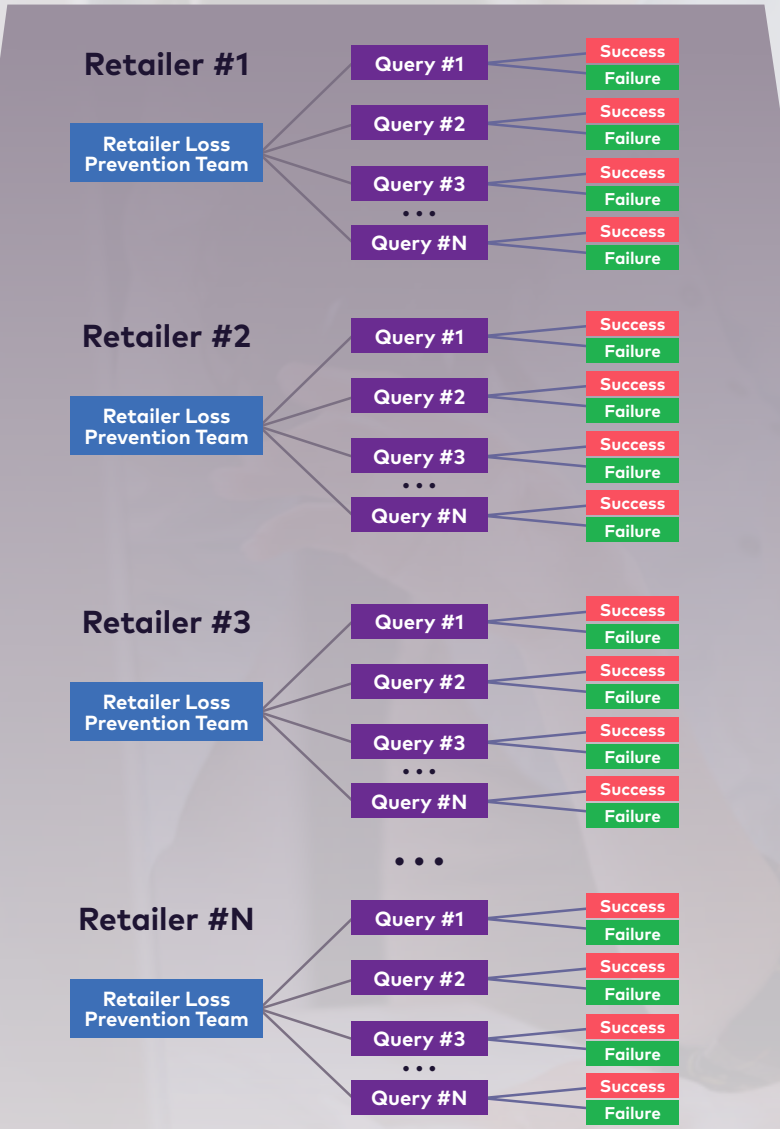

Combining retailer strategies leads to better detection

- Variables which result from shared questions
- Solutions built on many successful strategies

APPRISS® RETAIL

Using Dynamic Learning and Feedback

Beyond the collective intelligence described in the prior section, a final objective of an AI-based exception based reporting system is a feedback loop from the investigations back into the solution. During an investigation, the investigator will perform many actions that can be captured in a system and later used to improve system accuracy. For example, the investigator will click on related transactions and individuals, will open a case, will abandon a case or a transaction, and so on. The captured usage activity can be converted into information that machine learning can utilize to improve. This type of feedback loop ensures that as new fraud schemes are detected at one retailer, that pattern is captured in a broader model that propagates across all retailers.

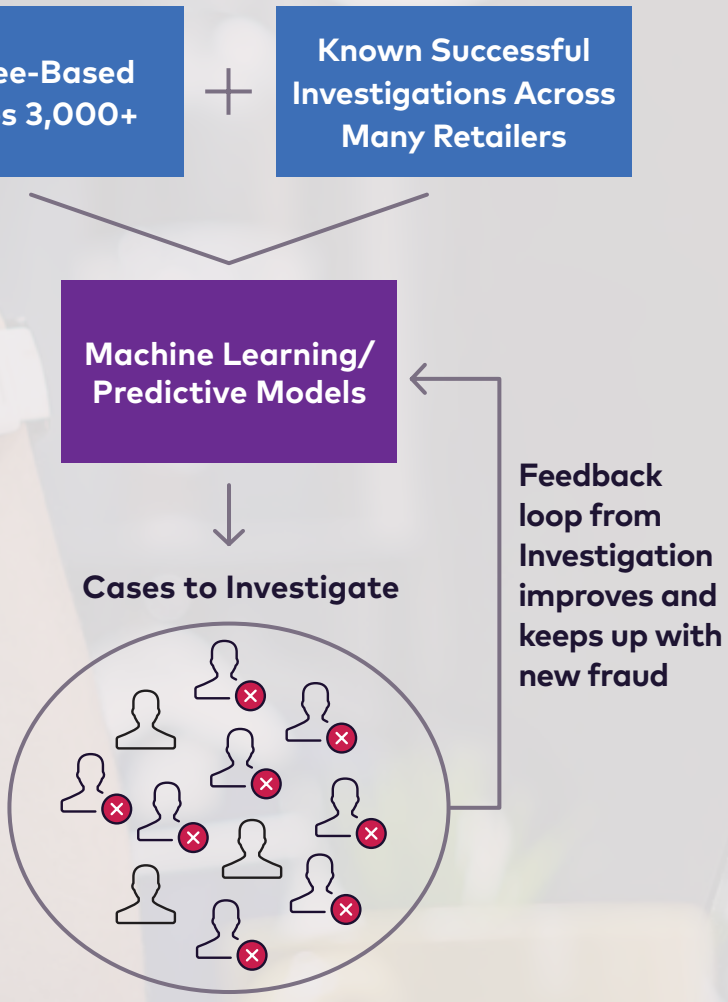

Summary

The ultimate objective for developing machine learning solutions is to combine learnings and obtain collective intelligence from data that has been deidentified and anonymized and not comingled to curb suspicious employee behavior.

Americas +1 949 262 5100
Europe/Middle East/Africa +44 (0)20 7430 0715
Asia/Pacific +1 949 262 5100

APPRISS® RETAIL
apprissretail.com