


AYASDI

Enterprise AI for Financial Services

WHITE PAPER



Introduction

Ayasdi has built an AI platform that powers the design, development, and deployment of enterprise-scale, intelligent applications. Our approach, our underlying technology, and our products are expressly crafted to deliver against a financial institution's requirements in this area – with the goal of delivering extraordinary business value.

This paper details our approach to creating intelligent applications and our suite of solutions in the financial services area.

Our Areas of Focus

While our underlying technology can be applied broadly, our focus in financial services can be categorized under three areas:

- Financial Crimes
- Customer Intelligence
- Regulatory Risk

In each of these areas we seek to create applications that solve discrete business problems through acceleration, performance, operational efficiency or risk mitigation. Our results, where publically reportable, are often orders of magnitude better than the existing state of the art.

Our Application Orientation and Our Philosophy of Intelligence

We believe that all applications need to become intelligent. To that end, we have built an AI platform that supports the creation of intelligent applications.

In some cases, where we see large, repeatable problems (e.g. anti-money laundering, regulatory risk modeling) we will build those applications. More often we will work with our clients to enable them to build intelligent applications. Additionally, we have several partners, such as Accenture and Deloitte, that build these applications on behalf of their clients.

Our platform supports application development because it achieves the broadest penetration of intelligence in the organization – empowering business leaders and subject matter experts with extraordinary capabilities that deliver equally impressive business results.

We believe that for an application to be truly intelligent it needs to provide specific capabilities. We call these capabilities the five pillars of enterprise intelligence and they have particular applicability in the financial services domain.

While some of these capabilities may seem self-evident, that is because they are taken as a single item. For an application to be truly intelligent, all of the elements must work in conjunction with each other.

The five pillars are:

1. Discover
2. Predict
3. Justify
4. Act
5. Learn

Let's take each of these concepts in turn and discuss their importance in the financial services context.

Discover

Discovery is the ability of an intelligent system to learn from data without being presented with an explicit target. It relies on the use of unsupervised and semi-supervised machine learning techniques (such as segmentation, dimensionality reduction, anomaly detection, etc.), as well as more supervised techniques where there is an outcome or there are several outcomes of interest.

This is important for financial services institutions because, in complex datasets, it is nearly impossible to ask the "right" questions. To discover what value lies within the data one must understand the relationships that are inherent and important in the data. That requires a principled approach to hypothesis generation.

Ayasdi's unique technology, topological data analysis (TDA), is exceptional at surfacing the hidden relationships that exist in the data and identifying those relationships that are meaningful without having to ask specific questions of the data. Ayasdi's output, topological models, are able to represent complex phenomena, and are therefore able to surface weaker signals as well as the stronger signals. This permits the detection of emergent phenomena.

As a result, enterprises can now discover answers to questions they didn't even know to ask. Given the size, complexity and inter-related nature of financial data, this capability is incredibly powerful and leads to orders of magnitude improvements in performance – as evidenced by our work with HSBC and Citi.

Still, Discovery is only one element of what makes an application truly intelligent.

Predict

Once the data set is understood through intelligent discovery, supervised approaches are applied to predict what will happen in the future. These types of problems include classification, regression and ranking.

For this pillar, Ayasdi uses a standard set of supervised machine learning algorithms including random forests, gradient boosting, and linear/sparse learners. The discovery capabilities of our technology are highly useful in that they: 1) generate relevant features for use in prediction tasks and 2) find local patches of data where supervised algorithms may struggle.

This is immensely important.

Traditional machine learning techniques depend on the appropriate selection of a discovery algorithm followed by the appropriate selection of a prediction algorithm. This process means that data scientists, not business executives, make these determinations and drive this process. This approach creates significant resource constraints because it is inefficient and iterative. It also often generates biased outcomes.

The automated use of unsupervised learning techniques alongside supervised learning techniques engages data scientists and subject matter experts in a constructive, collaborative approach that creates better outcomes in less time and with less friction. Furthermore, this approach creates an environment of justifiability.

Justify

We believe that intelligent applications need to support interaction with humans in a way which makes outcomes recognizable and believable to them. For example, when one builds a predictive model, it is important to have an explanation of how the model is doing what it is doing, i.e. what the features in the model are doing - in terms that are familiar to the users of the model.

This kind of familiarity is critical for generating trust. Trust, however, is by itself insufficient in highly regulated industries like financial services. Justification becomes a requirement and extends well beyond what is commonly referred to as transparency.

Transparency is the concept that one can identify what algorithm is being used, when it is used, and with what parameters. Transparency, however, does not tell you what the algorithm is doing, or, more importantly, why. This detailed knowledge of an algorithm's operation is called justification.

The concept of justification is far more robust than transparency and is required to move AI into production. Like transparency, justification identifies the algorithm that was used and the parameters that were applied, but it also provides the additional ingredient of intuition – the ability to see what the machine is thinking: “when x, because y.”

Justification tell us, for every atomic operation, here is the reason(s). For every classification, prediction, regression, event, anomaly, or hotspot, we can identify matching examples in the data as proof. These are presented in human understandable output and represent the variables or the ingredients of the model.

Justification is perhaps the most important attribute to consider when evaluating AI in the context of financial services and is the only way to build the trust necessary to move beyond models and into production environments.

Act

Insights without action have limited utility in an organization.

To maximize the impact of Discovery, Prediction and Justification, an organization needs to turn that effort into something that can be activated by subject matter experts. This requires packaging of these steps into an application that is “live” in the business process, that is seeing new data, and that is automatically executing the loop of Discovery, Prediction, and Justification on a frequency that makes sense for that business process.

For some processes, this frequency may be quarterly, for others daily. This frequency can even be measured in seconds. Too many solutions provide an answer at a point of time; an intelligent system is one that is always learning through the framework outlined here.

This is an area where financial services organization are inherently advantaged – they build and deploy models in the normal course of their business. What is different in this context is that financial services organizations will upgrade those models with intelligent applications that encode existing knowledge while improving upon their performance.

It should be noted that the process of successfully operationalizing an intelligent application within an enterprise requires change and evolution. These may include change in the organization, an acceptance that the application will evolve over time, and a willingness to make downstream changes – automated or otherwise.

Learn

Finally, intelligent systems are designed to detect and react as the data evolves. An intelligent system is one that is always learning, is live in the workflow, and is constantly improving. In the modern data world, an application that is not getting more intelligent is getting dumber.

Financial Crimes Intelligence

With the rise in global terrorism, international crime, and fraud, financial institutions are seeing a commensurate increase in regulatory pressure and an exponential increase in the fines levied against institutions that fail to implement effective crime-deterrence systems. One of the key challenges in financial crimes is that it is broad – encompassing a number of different areas that require different approaches. Finding a technology solution that can meet a firm’s needs across these dimensions is difficult, but certainly not impossible. Financial institutions around the world have applied Ayasdi’s technology to problems as diverse as anti-money laundering, fraud, and cybercrime.

AML

What makes AML such a difficult problem to solve is that it involves complex data, detailed workflows, and significant human involvement. The result is that the cost of compliance is increasing 50% year-over-year. The dimensionality of the data and the costs to support a robust AML process are further amplified because these are large, geographically diverse financial institutions.

The major issue is the volume of investigations; most do not result in a Suspicious Activity Report (SAR). As many as 95% of the investigations do not result in a SAR, which means that 95% of the effort of a large investigations team (anywhere from 500-5,000+ people) is not required.

At the heart of the problem is the balance between signal and noise. Too much noise in the form of false positives increases costs for the bank since they need to investigate the suspicious activity reports. Too little signal means that the bank is exposed to regulators because they might miss criminals.

The challenge is to find the lever points in the AML process by looking at the end-to-end process from KYC to the transaction monitoring systems and through to alert investigations. By identifying the lever points, an enterprise can deliver massive efficiency improvements with minimal impact on the existing AML infrastructure.

Ayasdi's Intelligent AML application targets three such lever points:

- **Intelligent Segmentation:** segment customers, transactions, and alerts for more optimized scenario threshold tuning.
- **Intelligent Alerts:** accelerate clearing of the alerts backlog by automatically categorizing alert priority (L1, L2, etc.) and providing the reasoning for an alert's auto-dispositioning.
- **Intelligent Typologies:** create customer profiles and identify anomalous behavior in customers to uncover new typologies not covered by existing rule-based systems.

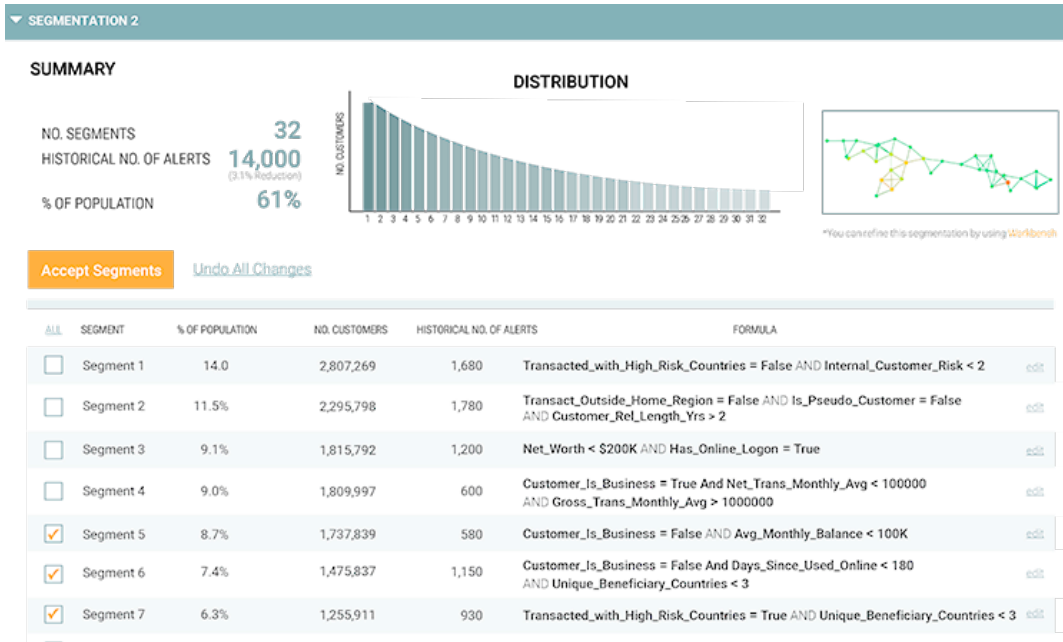


Figure 1: Intelligent Segmentation in AML

By targeting these areas and starting with Intelligent Segmentation, Ayasdi and HSBC reduced false positives by more than 20% while enhancing the bank’s overall risk profile. This was achieved by combining supervised and unsupervised approaches to create more uniform segments, better scenarios and higher thresholds. Deployed in North America with future geographies coming online over the next few quarters, HSBC has set a new bar for sophistication in an AML process.

Fraud

The complexity of modern fraud is stunning. Characterized by a lack of ground truth, most solutions suffer from similar false positive rates as AML while missing innovative or novel schemes until they have extracted significant costs.

In these situations, it is incredibly valuable to accurately determine fraudulent data points. Even in those situations where some of the labels might be available (sparse ground truth), it is often to be expected that there will be new forms of fraud or other behavior that one wants to discover, even if they don’t fit into any previous patterns. This means that they will not be discovered by methods that attempt to match with already occurring phenomena for which there are labels.




Figure 2: Topological Model of a Fraud Signature

Ayasdi's counter fraud applications combine multiple approaches to find existing fraud, as well as the "unknown unknowns" that keep management teams awake at night. The key is to be able to discover potential fraud using unsupervised approaches (anomaly detection) in conjunction with more supervised approaches (hotspot detection), while also learning that the discovered fraud is no longer anomalous. And finally, the reasons for, and the differentiators of, that fraud must be clear and interpretable so they can be actioned and built upon.

This Ayasdi counter fraud approach was implemented by one of the world's largest payers. After deploying Ayasdi's technology this payment organization was able to deliver a 2x improvement in leads, a 7x improvement in turnaround efficiency, and an 8x ROI. It is worth noting that this is a company that was already considered one of the most sophisticated fraud hunters in the world.

Cyber

The world's largest financial institutions are also the world's largest targets, not just for criminals, but also for foreign governments and, increasingly non-state actors. This is the world where the "unknown unknowns" dominate the discussion. This is also the world of truly "big" data – where log files and other information run to billions of rows and thousands of columns in any given week.

Here again, we find the challenge of false positives, combinatorial complexity, and scale. Solutions that depend solely on one analytic approach or even a few such approaches are ultimately vulnerable. Institutions need applications that scale to millions of customers and billions of events, that combine unsupervised learning with supervised techniques, and that can justify their findings.




Figure 3: Topological Model of Intrusion

While working with a major Japanese bank Ayasdi found such “unknown unknowns”. This bank had already implemented state of the art systems from all of the major vendors, yet Ayasdi’s software detected 60 anomalous groups, each with over 300 features, that hid 15 criminals, all in the first week of deployment. The software now runs lights out – allowing the management team to sleep better, knowing that they have mitigated a significant risk to the bank.

Customer Intelligence

Financial institutions are always seeking to grow and optimize their client population. Growth is a function of adding profitable new clients while reducing manageable churn. Optimization in this regard involves allocating resources efficiently to the bank’s clients in such way as to maximize customer satisfaction and profitability.

Using data to generate customer intelligence is inherently a dimensional problem. The amount of data is staggering and, as a result, most institutions don’t fully utilize such data because traditional approaches simply cannot make sense of it. Conventional approaches to analytics rely on business intuition and armies of analysts to keep pace with evolving client behavior, products, market conditions, regulations, and the growing complexity and volume of their data. While these approaches often surface macro-trends in client preferences, they are fundamentally static and they fail to uncover the subtler relationships that exist between your clients and your products. Further, such approaches cannot easily be encoded into those mathematical models and deployed in an automated fashion; encoding and deployment being a fundamental requirement to transform your customer acquisition and product offering capabilities.

Ayasdi’s Customer Intelligence applications represent an innovative approach that allows financial institutions to create behavior-based client profiles to meet these goals. Using client, product, and market-related data as well as external information, Ayasdi’s platform uncovers subtle, precise client sub-segments that can be effectively actioned. The underlying attributes that describe each segment can also be used to inform the development of precise client profiles and predictive models that can be dynamically updated and deployed into operational systems.

AYASDI CUSTOMER INTELLIGENCE 1.0

Private Banking Customers Results

Attrition

Product Recommendation

Advisor Performance

Q

Advisor Performance

ADVISOR NAME	HIGH ATTRITION CUSTOMERS	MEDIUM ATTRITION CUSTOMERS	LOW ATTRITION CUSTOMERS	AUM \$
Anthony Trifero	3	1	28	12.337
Samantha Malik	3	1	24	11.937
Ian Gray	2	1	26	11.617
John Walker	2	1	39	10.527
Otis Selib	1	0	27	9.6438
Jeff Richards	1	0	30	8.1599
Walker Connelly	1	2	21	7.7015
Stephanie Lippens	1	0	30	7.4946
Sophia Faro	1	1	42	5.8436
Zara Khan	1	0	35	5.3044
Bryce O'Callaghan	1	0	29	5.2997
Becky Porter	1	1	28	4.7917
Alex Anderson	1	1	35	2.2446

Figure 4: Advisor performance ranking generated via Topological Data Analysis

Ayasdi's Customer Intelligence capabilities cover three core areas:

- **Client segmentation:** segment clients more granularly by leveraging customer profile and transaction data.
- **Churn prediction:** predict customer churn months in advance in order to service a client before they terminate a relationship.
- **Product propensity identification:** discover products or services that clients are likely to purchase in a compliant and risk-adverse manner while taking into account both a client's and a product's risk profiles.

Ayasdi's capabilities in Customer Intelligence were recently demonstrated at one of Canada's largest banks. The goal was to more precisely identify what clients were likely to churn within the next six months and why, and thereby enable the bank to create an ROI-based mitigation plan for each segment. An additional goal was to determine what customers were likely to add to their investable assets over a similar timeframe.

The project involved several hundred thousand clients and looked at more than 140 attributes ranging from return performance to life events. Within a couple of weeks the bank had identified millions of dollars of opportunity by predicting high risk clients with 70% accuracy through an application that seamlessly integrated with their existing wealth management platforms.

More importantly, the work identified additional opportunities for application of Customer Intelligence technology including wealth advisor profiling, client onboarding, cross selling, and fee optimization.

Regulatory Risk Intelligence

Accurately assessing risk exposure requires a deep understanding of the complex and dynamic interplay of market and macroeconomic variables as well as the ability to continuously update models as conditions change.

Ayasdi has a suite of technologies and applications, such as Ayasdi Model Accelerator, that are purpose built for the challenge of rapidly designing, developing, and deploying accurate and defensible risk models.

The Ayasdi Model Accelerator excels at a variety of modelling challenges but is particularly effective when applied to enterprise stress testing, loss-given default (LGD), Probability of Default (PD), and other regulatory modeling problems. These models typically make use of regression techniques to support the simplicity and explainability requirements associated with regulatory oversight. However Ayasdi's Model Accelerator supports a range of other approaches.

At the heart of Ayasdi's Model Accelerator is a powerful approach to data science that encompasses both unsupervised and supervised learning capabilities. By leveraging these learners, the Model Accelerator can identify groups of variables and explain why they were selected.

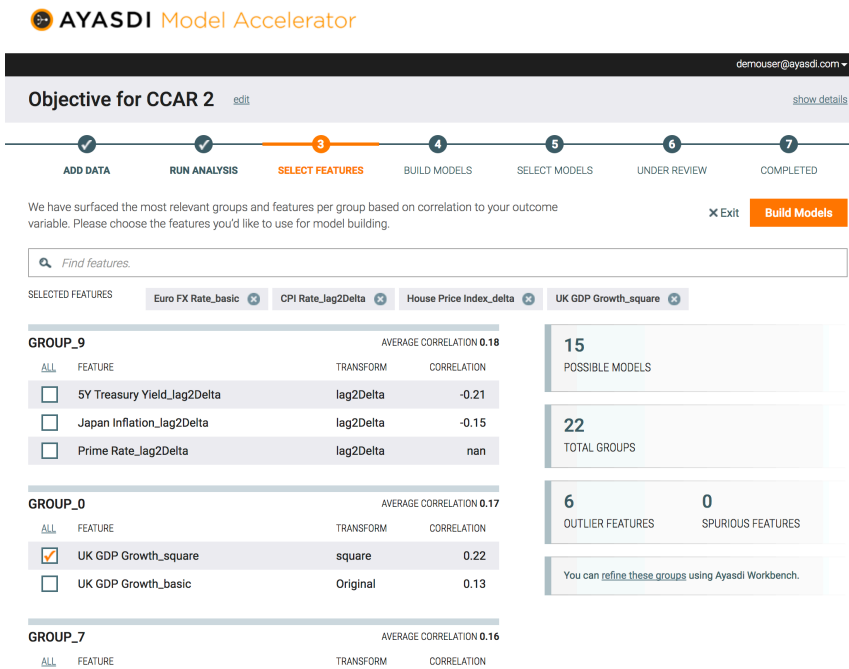


Figure 5: Model Accelerator screenshot of CCAR process

While Ayasdi's Model Accelerator (AMA) has unique capabilities in unsupervised learning, often model developers and line-of-business leaders want explicit control over the model through the selection of key variables. This variable-selection approach is fully supported in AMA, however AMA

goes one step further by documenting *any* bias associated with the selection of certain variables and the elimination of others. This leads to better models and greater transparency.

During the model creation process, everything is documented automatically. The AMA systematically and deterministically selects variables and creates models, thereby implementing a repeatable process that generates consistent, supporting reports on model lineage, variable selection, and cross validation. In addition to recording and documenting initial selections, the AMA also tracks and catalogs the entire modeling and approval process. The result is information usable by model developers and collaborators to facilitate model review and reuse.

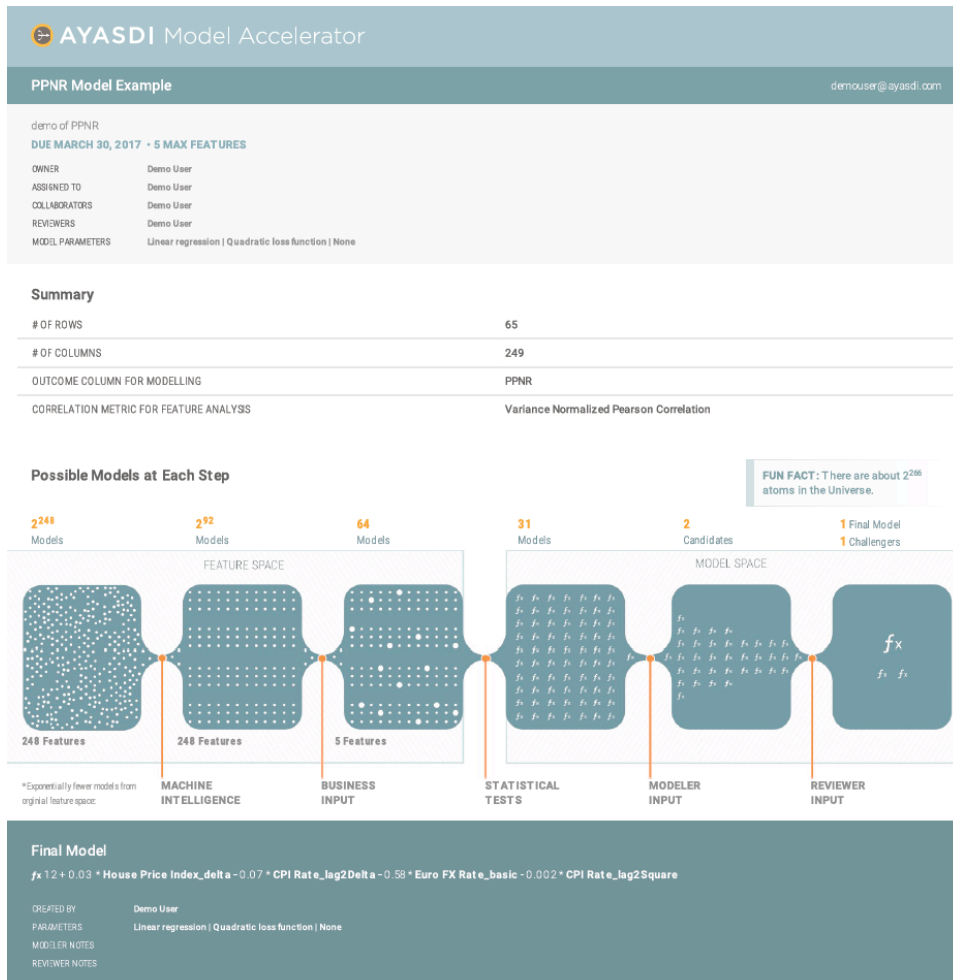


Figure 6: Model Accelerator overview of model creation steps and final linear model for regulator

In addition to the capabilities outlined, Ayasdi’s Model Accelerator application can be used both as a model generation engine and as a model execution engine. It can also be deployed as a model repair and improvement application, wherein its capabilities help detect areas where models are failing. Finally, Model Accelerator always creates a portfolio of models that contains the candidate model (selected using Ayasdi’s unsupervised learning approach) as well as a number of viable challenger models. This is particularly important in a regulatory setting where performance may dictate changing models over time. Further, the ability to demonstrate the rigor of evaluating challenger models is also of value to the enterprise when explaining the models.

Model Accelerator in action: simplifying CCAR/PPNR

Several years ago, one of the world's largest banks was consistently struggling to pass its annual stress test. The bank needed to rapidly create accurate, defensible models that would prove to the Federal Reserve that they could adequately forecast revenues and the capital reserve required to absorb losses under stressed economic conditions. The bank's existing modeling approach left the business unit leads with little time to consider the logic behind variable selection. The result was an inability to confidently defend the models included in filings presented to the Federal Reserve.

Using AMA the bank transformed the Federal Reserve's base variables, ultimately creating over two thousand variables. Ayasdi's Model Accelerator rapidly correlated and analyzed the impact of these variables on each business unit's monthly revenue performance over a six-year period, uncovering statistically significant variables that were highly correlated with each business unit's performance.

Each business leader selected the relevant variables from these generated lists and Ayasdi's software then ran another series of statistical tests to validate these models' ability to predict revenues for the business units.

The business leads then evaluated the candidate model and the challenger models, selecting those that best represented their business units. With a collection of accurate, transparent, and defensible revenue forecast models, the bank could easily clear their most challenging regulatory hurdle.

"Citi's unmatched multinational business footprint creates a complex data analytics landscape. Ayasdi's technology simplifies and accelerates the analysis of thousands of discrete variables and delivers outcomes that enable Citi to tailor services to specific client needs, operate more efficiently, and mitigate risk."

Deborah Hopkins,
Chief Innovation Officer

Model Accelerator in action: automating PD modeling

One of the largest Nordic banks was having issues with mortgage defaults and losing large amounts of money by incorrectly predicting which mortgages were likely to default. The bank used Ayasdi's Model Accelerator to analyze the accuracy of their existing Probability of Default (PD) models. AMA was able to understand the complex population, compute confidence level values for every model prediction, and thereby identify areas in which the PD models did not fit the data well. By automatically flagging low-confidence model predictions for review, the AMA automated the process of locating PD model weaknesses.

Furthermore, operationalizing the new, improved models was easy for the bank. AMA allowed the bank to maintain their existing PD models but create TDA overlays for the problem regions of their

models and implement targeted fixes. The bank was able to easily export both their new models and the TDA-defined features into their existing model management system.


Summary

Ayasdi's suite of applications leverage the power of its machine intelligence platform to deliver applications that distill extraordinary complexity into simple, transparent answers that are interpretable by subject matter experts and regulators alike. Across key risk areas and emerging revenue opportunities, Ayasdi has the power to transform how financial services organizations operate.

About Ayasdi

Ayasdi is the global leader in the development of enterprise-grade, machine intelligent applications for financial services, healthcare and the public sector. Powered by breakthroughs in both mathematics and computer science, the company's software platform has already solved some of the world's most complex challenges.

Based in Menlo Park, CA, Ayasdi is backed by Kleiner Perkins Caufield & Byers, IVP, Khosla Ventures, Centerview Capital Technology, Draper Nexus, Citi Ventures, GE Ventures and Floodgate Capital.



AYASDI

4400 Bohannon Drive, Suite 200
Menlo Park, CA 94025 USA
+1 650.704.3395
sales@ayasdi.com
ayasdi.com | @ayasdi