AI Innovation in AML Transaction Monitoring

Key AML data science challenges and how Featurespace addresses them

FEATURE SPACE

OUTSMART RISK

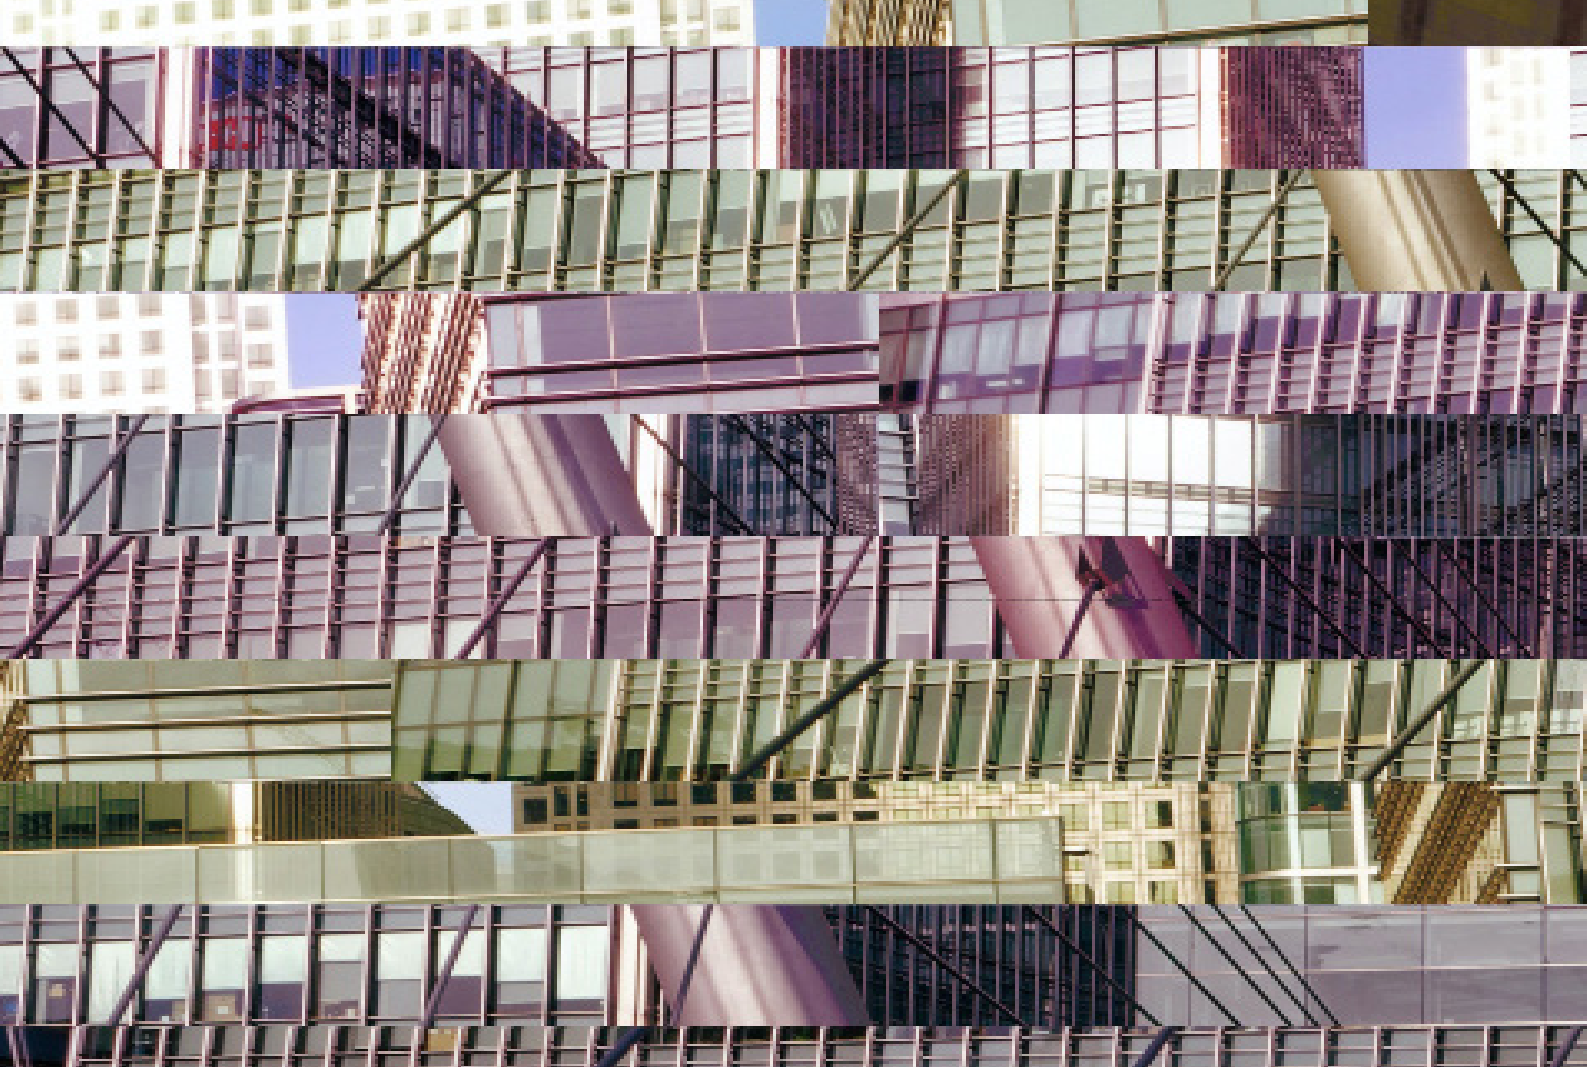

Featurespace has been providing Adaptive Behavioral Analytics to the financial services industry for over a decade, providing individual and peer group behavioral profiles for some of the world's largest banks.

Our expertise in providing true profiling is now replacing traditional rules-based Anti-Money Laundering (AML) transaction monitoring solutions. Featurespace's ARIC™ Risk Hub offers an end-to-end AML transaction monitoring solution that can be deployed as a standalone solution, or as augmented analytics alongside an existing product.

Data science provides the unique approach Featurespace is known for - and a huge amount of research into the most advanced machine learning techniques ensures we give our customers the best machine learning solution possible.

Why use Machine Learning for AML Compliance?

The benefits of machine learning that have long been enjoyed by fraud teams are now being adopted by their AML compliance counterparts. However, there are challenges in this adoption, as machine learning for AML requires a different approach altogether to achieve the same results.

There are some important reasons for this, and Featurespace's expert team has put significant industry and scientific research, domain knowledge, and development resources into investigating and overcoming these challenges.

This report will outline those challenges for compliance teams and explain how Featurespace is enabling its customers to overcome them and see real benefits using innovative machine learning techniques.

Rules-Based Systems or Machine Learning?

Rules-based systems have long been the industry norm for AML transaction monitoring. However, there has recently been a move towards machine learning in this space. Regulators have been pushing financial institutions towards real-time transaction monitoring in recent years, and forward-thinking financial institutions are already exploring and adopting this technology. In particular, payment service providers are finding real-time processing of great risk management benefit due to the instantaneous nature of their services.

The opportunities for uplift are plentiful, because of the range of data and features that can be analyzed simultaneously by a machine learning system.

For example, Featurespace provided a global tier 1 bank with an incredible increase of 133% on the suspicious activity detected. Additionally, ARIC's prioritization model identified all of this activity within the top 5% of alerts.

The power of machine learning, when applied correctly, is immensely beneficial to compliance teams.

Featurespace achieved these results not only because of machine learning but because of the unique approach taken in applying the technology. The behavioral approach ARIC uses to build rich profiles of individuals and their peers provides unparalleled anomaly detection. More effective than most AI solutions, and a far cry from traditional rules-based systems.

A rules-based system considers a small number of feature values, which are individual properties or behaviors, and triggers an alert based on a static threshold. Conversely, a machine learning model predicts the probability based on a multitude of blended feature values, triggering alerts based on probability according to the customer's risk exposure, inherent and residual.

When blending the signals, the model considers how extreme each feature value is, and also how relevant each feature is for predicting money laundering risk in the first place.

A machine learning model's ability to estimate risk accurately when there is part of one signal plus part of another signal is what gives it its analytical edge over rules systems.

The difference in rules-based systems and machine learning models is highlighted below, where the boundaries of the alerting region for the rules system are rigidly horizontal or vertical, whereas the probability contours for the model are flexible and curved.

Rules-based system

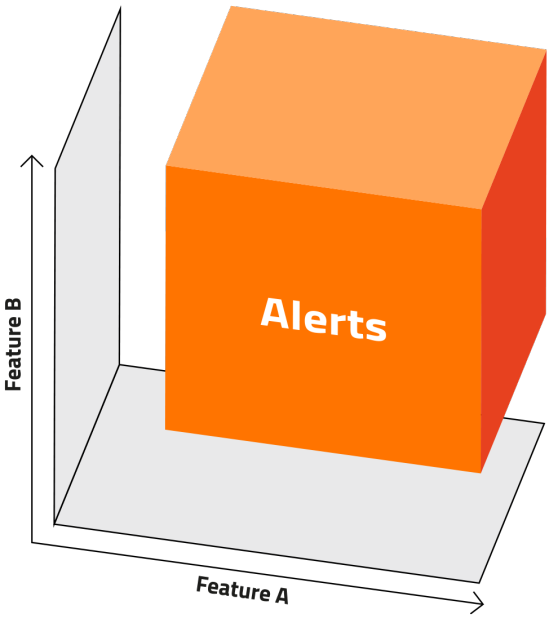

Machine Learning

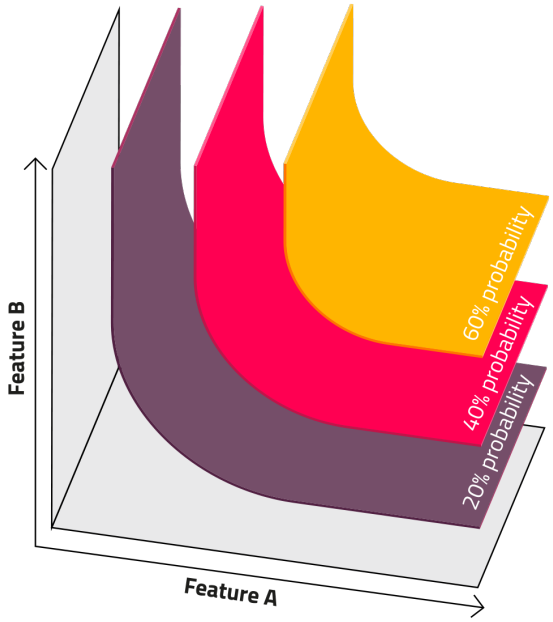

Rules provide an absolute threshold on an individual feature: either the activity satisfies an alert condition, or not.

In contrast, machine learning provides a continuous probability for suspicious behavior.

This produces a learned decision boundary that blends information from many features. High probability regions contain fewer cases, but those alerts are almost purely cases of money laundering. Low probability regions contain more cases, but there are more cases of genuine activity mixed in. The customer can decide the trade-off between certainty and coverage in their alerts.

The machine learning model's logic is too complex for humans to configure by hand, and so is learned by a machine.

Supervised vs. Unsupervised Machine Learning

When talking about machine learning, there are often two types of machine learning discussed: supervised and unsupervised.

Supervised Learning

Supervised learning algorithms are trained using labeled examples. Data scientists commonly use supervised learning in applications where historical data predicts likely future events. The data scientists teach algorithms to map from multiple inputs (typically referred to as features) to an output or label, which they then test on a historic set of data. This trained data is then used to create models for a second prediction phase.

Unsupervised Learning

Unsupervised learning is an approach applied to data that has no historical labels. In contrast to supervised learning, unsupervised learning means that data scientists do not tell the machine learning system the correct answer or answers: the algorithm itself determines what is important and learns the intrinsic structure and relationships between different data elements. Unsupervised learning works by grouping the data into clusters that are similar to one another or identifying outliers or "anomalies" in the data.

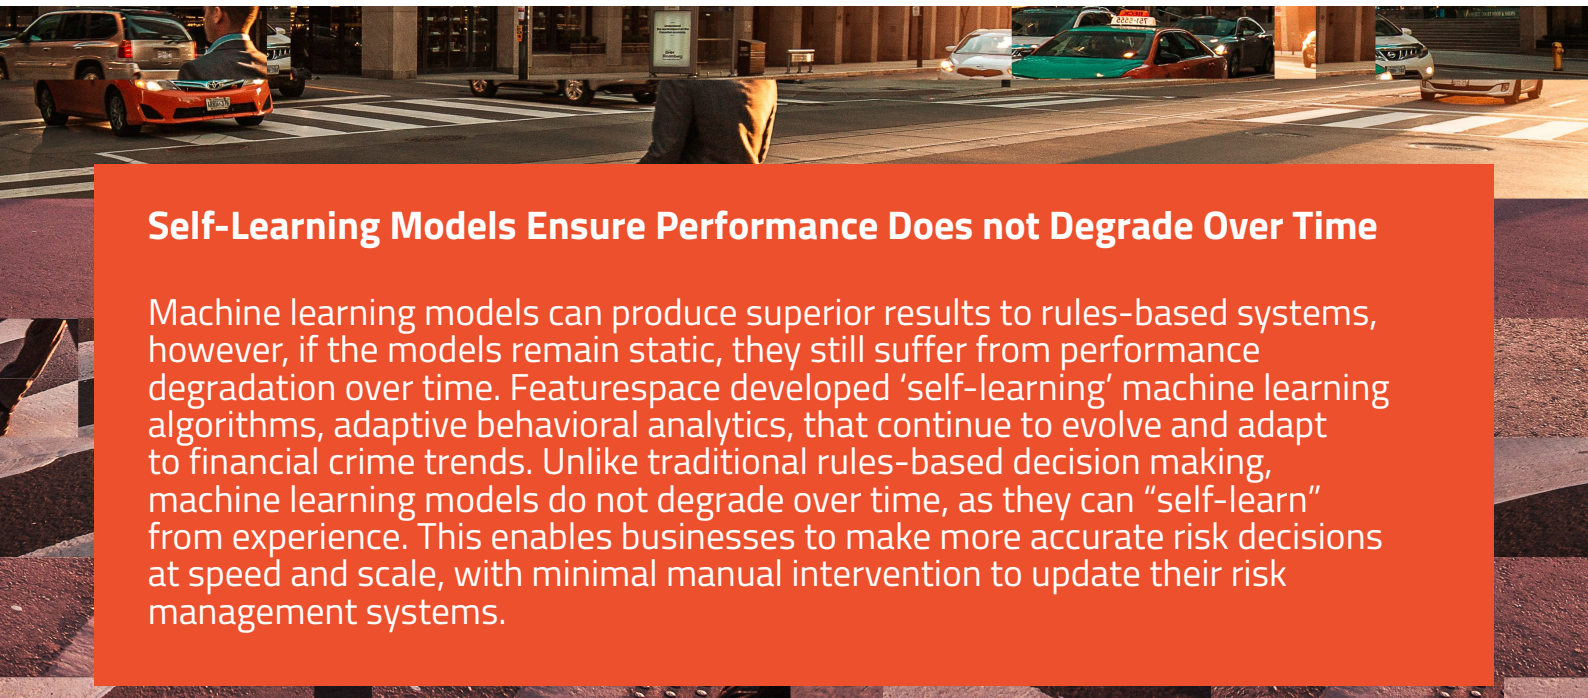

Self-Learning Models Ensure Performance Does not Degrade Over Time

Machine learning models can produce superior results to rules-based systems, however, if the models remain static, they still suffer from performance degradation over time. Featurespace developed 'self-learning' machine learning algorithms, adaptive behavioral analytics, that continue to evolve and adapt to financial crime trends. Unlike traditional rules-based decision making, machine learning models do not degrade over time, as they can "self-learn" from experience. This enables businesses to make more accurate risk decisions at speed and scale, with minimal manual intervention to update their risk management systems.

Training a Supervised Machine Learning Model

Machine learning models are trained by analyzing historic data, where the data points that identify if a behavior is suspicious are known as labels. Without the historic data or the labels attached to the data, it is impossible to train a supervised machine learning model.

The more labeled data is available, the more accurate the trained model will be at estimating money laundering probabilities on future transactions.

Most machine learning models require frequent manual retraining to refresh the training data and improve performance. This takes place sporadically and involves manual work from data scientists. Therefore, the machine learning model is only effective for short periods after retraining.

However, Featurespace's unique Adaptive Behavioral Analytics 'self-learning' models automatically retrain regularly using the most recent training data. This way, the model never degrades, and its predictions reflect the latest trends seen in the data and utilize recent feedback from AML investigators.

Rules-based System	Machine Learning
Raises alerts at static threshold	**Returns probabilities. Alerts raised if probability exceeds threshold**
Each rule considers only a handful of features	**Weighs up strength and relevance of signals from all features**
Simple logic, manually tuned by humans	**Complex logic, trained by computer analyzing trends in historic data**

The Challenges in Using Machine Learning for AML

Unlabeled data is an obstacle for data scientists when creating machine learning models to detect suspicious activity. Labeling data to indicate if behavior is suspicious or not is vital in teaching a machine learning model what to look out for in the future. However, labels for AML data are often sparse.

Label sparsity can be caused by a low number of suspicious activity reports (SARs), which we mitigate by using escalated level 3 alerts as a proxy for SARs.

This provides more data with which to train the models, but still less than would be optimal. Simply put, it is difficult to obtain positive labels identifying suspicious behavior because most transactions are legitimate.

There is also the issue of **label representation bias**, which is where an incumbent system (from which training data is obtained), has missed suspicious behavior when producing alerts.

The data tells us which of the system's alerts were deemed suspicious by the human investigators who reviewed them.

But our data does not tell us which of the entities that were NOT alerted on by the system would have been deemed suspicious by human investigators, had they too been reviewed.

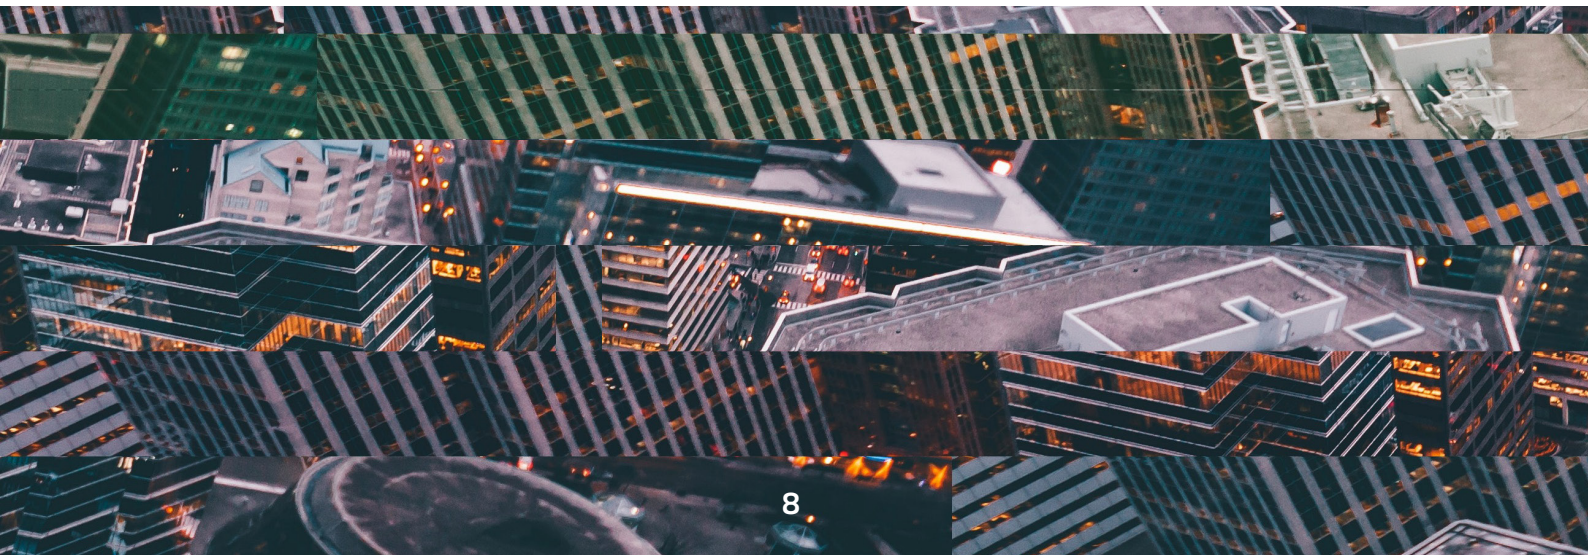

Consider an analogy. Imagine you are a parent who feeds their daughter apples and oranges. You notice that your daughter eats the apples you give her, but rejects the oranges.

So, you learn that your daughter likes apples and does not like oranges.

But does your daughter like pineapples? You cannot know because you have never given your daughter pineapple to try!

Therefore, the data used to train models is incomplete, and the model cannot identify which of the non-alerted training data is actually suspicious behavior that was previously missed.

This can derail a machine learning model that relies on labels to learn what suspicious behavior looks like.

At Featurespace, we have worked with the financial services industry to mitigate these issues with a combination of scientific expertise, human intervention, and exploration.

The Scientist, the Explorer and the Human Expert

They did not walk into a bar. Instead, scientific expertise is used alongside human expertise in the form of rules, and exploration in order to obtain the best possible results.

Human expertise is used in the creation of rules - this is why they have been used in the past, after all, and so it might be unwise to discount these altogether. Particularly with the issue of label sparsity that can hinder the efficacy of machine learning in AML.

This approach is also helpful when there are instances of underrepresentation in data that the human expert can identify, or when regulatory requirements stipulate certain behaviors need to be reported on under all circumstances.

It can also be used to explore new types of suspicious behavior, or when intelligence is shared between teams on specific typologies.

In cases where the machine model does not have sufficient labeled data to make a conclusive decision by itself, we may wish to incorporate some human expertise into the decision logic to arrive at the final decision.

The use of human domain knowledge (in the form of rules) helps address both label sparsity and label representation bias. The human knowledge helps make decisions in cases where the model has insufficient labelled training data to make the decisions by itself.

Exploration addresses the label representation bias issue we covered earlier. It is recommended that a small part of the alert investigation time is reserved for transactions that have not triggered alerts in the past. AML compliance investigators then look at these transactions, identify which behavior should be labeled as suspicious, and their feedback is used to improve the training data in the next automatic retrain.

Using our previous analogy, this is where we offer pineapple to our daughter, to learn whether she likes them or not - applying labels to her behavior.

Exploration is introduced with a temporary rule that looks for a specific type of behavior - for example, human trafficking. If the alerts triggered by this rule are confirmed as suspicious by the human investigators, they will be escalated, resulting in positive labels being added to the next batch of training data.

This is a short-term effect. The intention is that this exploratory feature will reap rewards and improve machine learning models with the enhanced training data. However, making the feature temporary ensures that if the alerts generated are not useful, the impact is only short-term.

A hybrid system can strike the perfect balance between a rules-based system that allows for manual configuration using rules created with human knowledge and a machine learning model solution that creates incredible uplift in analytical performance.

For example, a rules system where the rules condition on the outputs of machine learning models.

This means that where labels have not yet been applied, rules can intervene by adding human expertise into the decision logic. With each automatic retrain, the suspicious behavior the rules have identified are fed into the machine learning model using the training data, gradually decreasing the reliance on rules over models.

Researchers in Featurespace's AI Lab identified that a hybrid solution can be the best approach for many customers by taking public datasets such as San Francisco crime rates.

They then applied modeling and rules principles to an unlabeled portion of the data. Subsequently applying the same logic to AML data, they tailored the models to the specific problem of detecting money laundering.

In the below example, our researchers took the data on San Francisco crime rates to replicate the batches and labeling issues they were likely to incur with AML data.

They were then able to identify that the best way of training the machine learning models was in tandem with rules to create a large uplift on identifying San Francisco crime - and were then able to implement this same logic on AML data.

This graph shows that with the rules (representing the incumbent system), there is little to no fluctuation, no improvement over time.

However, the model on its own has no training data and therefore misses true positives for a period until it is trained.

By blending the two systems, the base level the rules provide is retained, and at the same time, there is an uplift in performance the model provides.

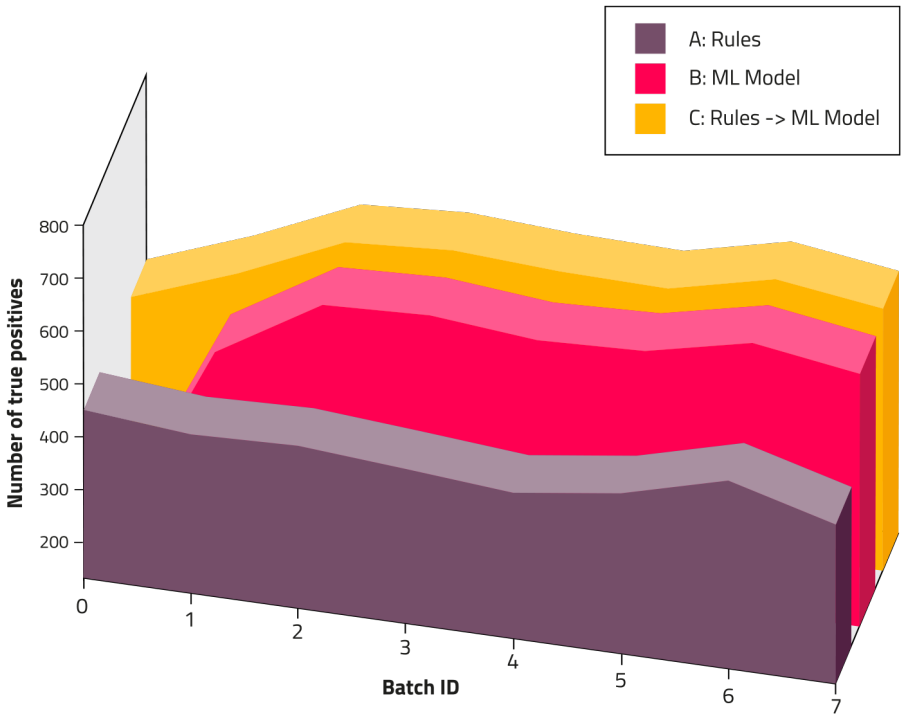

For the most effective performance, the system will increasingly rely on the model over the rules. Each month, the machine learning model becomes more accurate as it ingests training data from the alerts the rules are triggering.

The combined logic and results from the rules are ingested by the machine learning model so that, eventually, the system relies solely on machine learning to catch suspicious activity.

Regulatory Scrutiny

Regulatory scrutiny has also been a delaying factor in the adoption of machine learning for compliance teams. However, as regulators embrace a more encouraging attitude towards innovation in combatting financial crime, forward-thinking financial institutions are starting to use sophisticated technology.

A historical concern with machine learning models is a perception that they are complicated and incomprehensible to human understanding, becoming a 'black box'. For compliance teams, this renders them unusable, as it is a regulatory requirement that investigators can explain why behavior has been flagged as suspicious.

To address these industry concerns, Featurespace developed explainable models within the ARIC Risk Hub. Customers are provided with clear reason codes to illustrate why the machine learning models have generated specific alerts.

ARIC Risk Hub uses a heuristic technique that shows to what extent each feature contributed to the risk score. This produces reason codes for each alert, which can be understood by the investigator and shown to the regulator along with all the relevant customer and transactional data.

Common reason codes provided include: 'High-risk jurisdictions', 'Large withdrawal after high inbound velocity', and 'Network risk with counterparties'

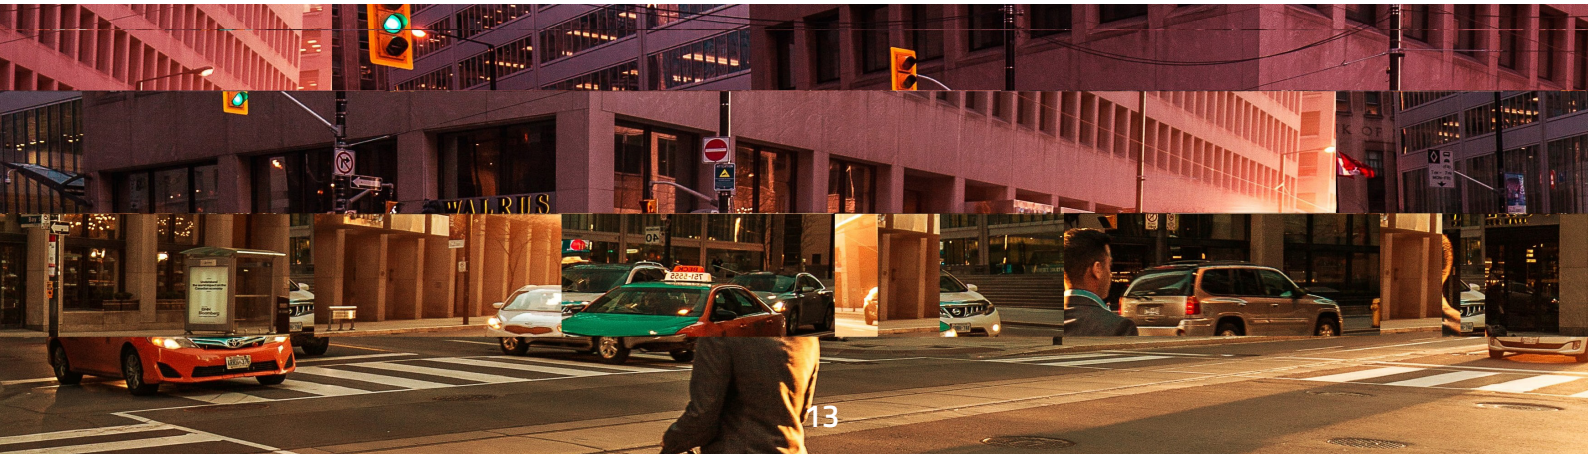

Providing high-quality explanations with the relevant data also improves investigator accuracy, or enables the routing of alerts to investigators with specific areas of expertise.

A robust feedback loop is produced that ensures model results continue improving with ARIC's self-learning capability.

Explainable models ensure that we can regularly check the quality of our proprietary algorithms with expert investigators and industry experts.

Our AI laboratory is constantly finding new and improved ways to make our models explainable to meet business needs.

Featurespace works with global customers across the financial services industry, providing both customized and standardized features based on our industry expertise.

With customized features, custom reason codes are essential for interpreting the varied financial crime typologies that customers encounter.

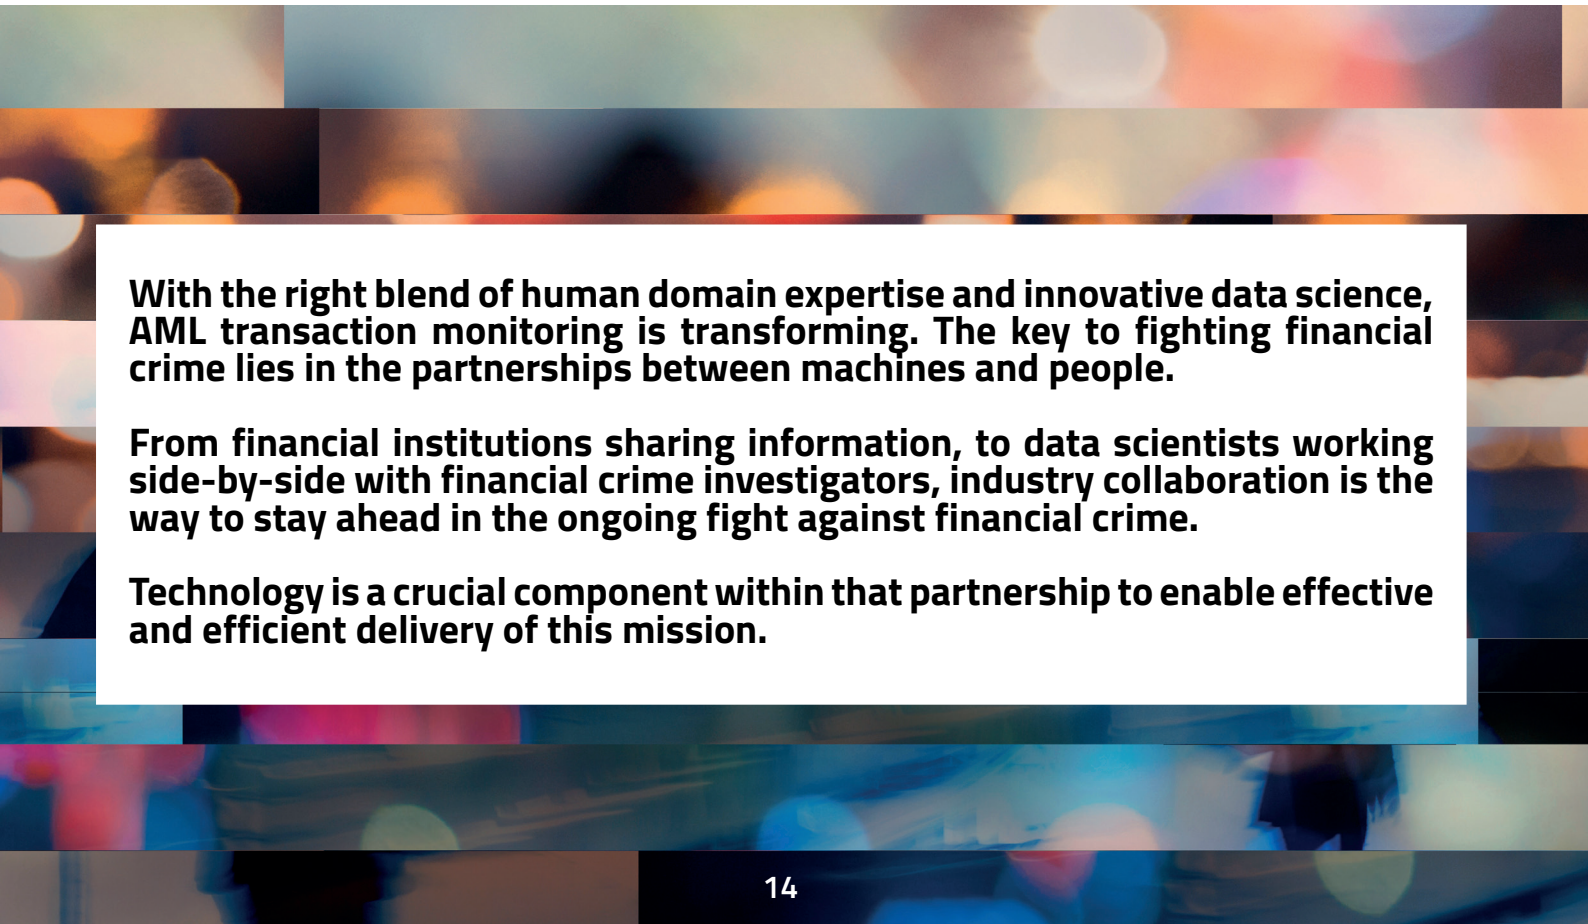

With the right blend of human domain expertise and innovative data science, AML transaction monitoring is transforming. The key to fighting financial crime lies in the partnerships between machines and people.

From financial institutions sharing information, to data scientists working side-by-side with financial crime investigators, industry collaboration is the way to stay ahead in the ongoing fight against financial crime.

Technology is a crucial component within that partnership to enable effective and efficient delivery of this mission.

Want to Outsmart Financial Crime with ARIC™ Risk Hub?

Featurespace is proud to provide an enterprise risk prevention platform that blends expert data science and industry knowledge with the best Adaptive Behavioral Analytics, real-time machine learning models, and profiling rules available.

Through constant and continual innovation, Featurespace provides this award-winning and industry-leading combination to our customers to protect their businesses from financial crime risks.

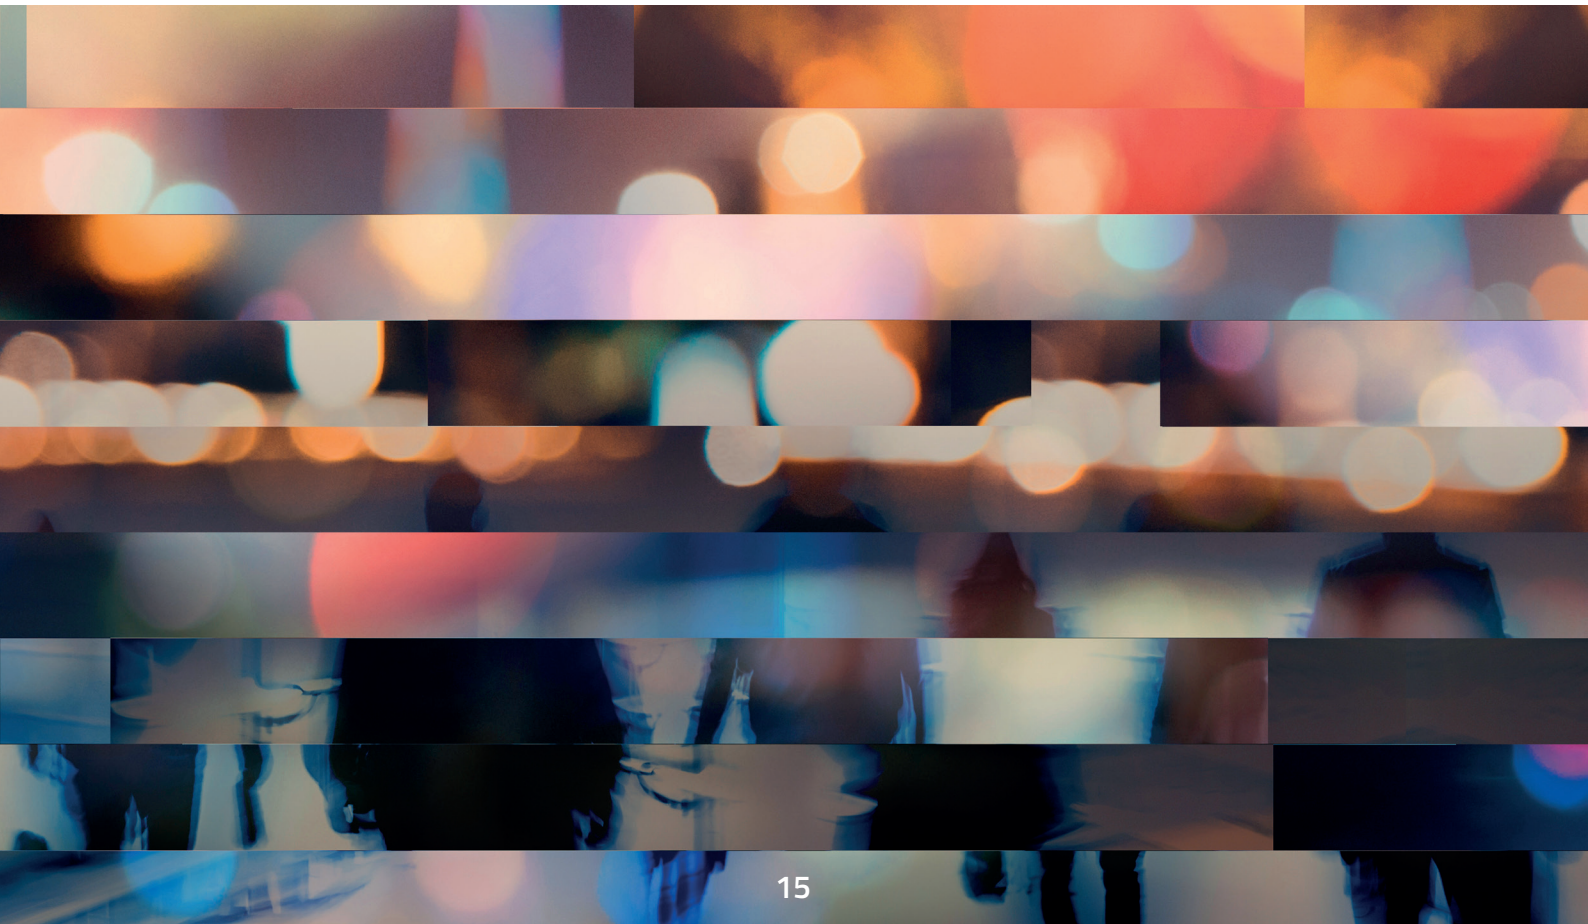

Find out more

Get in touch to discuss a standalone solution or how to enhance your existing system

info@featurespace.com
www.featurespace.com
UK +44 (0)20 3962 8989
US +1 (404) 649 0108

FStech awards 2019 WINNER

THE QUEEN'S AWARDS FOR ENTERPRISE 2018

THE SUNDAY TIMES TECH TRACK 100 2019 HISCOX
'Best Management Team'

Aite
AITE IMPACT MATRIX (AIM)
BEST IN CLASS
"AIM Evaluation: Fraud and AML Machine Learning Platform Vendors.", AITE, 2019

TOPA 2018 BEST SECURITY OR ANTI-FRAUD DEVELOPMENT Category Winner

2018 STEVIE GOLD WINNER AMERICAN BUSINESS AWARDS

AI 100 CBINSIGHTS 2020

The Card & Payments Awards
2019 CATEGORY WINNER
Industry Achievement Award

Cambridge, UK

Featurespace,
140 Science Park,
Milton Road,
Cambridge, CB4 0GF
United Kingdom

Atlanta, USA

Featurespace,
600 Peachtree Street NE,
Suite 420
Atlanta, Georgia 30308
United States of America

London, UK

Featurespace,
110 Bishopsgate,
London, EC2N 4AY
United Kingdom