

**F E A T U R E
S P A C E**

OUTSMART RISK

**Whitepaper:
Featurespace Model Governance
for Anti-Money Laundering**

Index

1) Foreword	3
2) The Increasing Place of Model Governance in AML and Fraud	4 - 5
3) Model Governance for Machine Learning in AML	6 - 7
3.1 Risks Outlined by OCC	7
4) The Featurespace Approach to Model Governance	8 - 15
4.1 Model Performance	9
4.1a Data Validation	10
4.1b Model Validation	11
4.2 Understanding Model Decisions	12
4.3 Model Fairness	13
4.4 Model Stability	14 - 15
5) Case Study: Tier 1 Global Bank	16
6) Conclusion	17

1 Foreword

Featurespace has designed the Model Governance Framework with the input of our customers and regulators. The technical and process information contained in this paper is created to provide confidence in how Featurespace's models are developed to identify criminal activity for financial crime investigation. The underpinnings of this framework capture Featurespace's Machine Learning Equality Policy and our commitment to producing solutions that enable fair access to financial services products.

After Featurespace's or our customers' own models are deployed, governance doesn't stop as it's important to ensure the decisioning process continues to remain effective and fair. Featurespace's adaptive models ensure performance is maintained and our explanations ensure investigators have insight as to why a transaction is suspicious.

Our goal to make the world a safer place can only be achieved by working hand in hand with our customers, investigators and regulators. With fairness and transparency alongside performance, Featurespace's platform coupled with our machine learning models are recognised as the most effective when used to protect organisations and consumers from financial crime.

Dave Excell, Featurespace Founder


2 The Increasing Place of Model Governance in AML

With ever increasing interest in the expanded use of Artificial Intelligence (AI) and machine learning (ML), significant resources are being invested in what has been termed the “Fourth Industrial Revolution”. A phrase first introduced by a team of scientists developing a high-tech strategy for the German government.¹ The Fourth Industrial Revolution is being built on the Third, the digital revolution that has been occurring since the middle of the last century. It is characterized by a fusion of technologies that is integrating the physical, digital, and biological spheres, for increased automation, improved communication, and self-monitoring.²

The momentum created by the Fourth Revolution is flourishing and it is disrupting every industry, not least Anti-Money Laundering (AML) by transforming entire systems of production, management, and governance.

Although rules-based systems and thresholds have traditionally been used to identify financial crime in transactions, criminals are now moving faster and becoming wise to the thresholds that these systems use to identify potentially illicit behavior.

It is now necessary to accelerate and elevate detection methods to match and even surpass criminals’ sophistication.



ARIC Risk Hub’s Adaptive Behavioral Analytics

AML transaction monitoring has proven a difficult task for machine learning to overcome, with often poorly labeled data from a vast number of sources delivered in batches creating a unique data science challenge.

However, Featurespace became the first company to successfully apply machine learning to AML transaction monitoring with its now award-winning³ machine learning model for a Global Tier 1 Bank. Featurespace’s ARIC Risk Hub converges data from multiple sources with holistic monitoring, utilising Adaptive Behavioral Analytics to understand all behavior and spot bad actors.

¹ “Industrie 4.0: Mit dem Internet der Dinge auf dem Weg zur 4. industriellen Revolution - vdi-nachrichten.com”. web.archive.org. 4 March 2013. Retrieved 25 January 2021

² https://en.wikipedia.org/wiki/Fourth_Industrial_Revolution#cite_note-2

³ <https://www.featurespace.com/newsroom/with-featurespace-hsbc-wins-celents-model-risk-manager-award/>

Utilizing behavioral profiling, AML investigators can build more complex pictures of financial criminals than ever before. The differences in behavior between criminals and their peers conducting normal activity can be minute - and can be identified using Featurespace machine learning models.

As pioneers in bringing machine learning to AML compliance, Featurespace has long considered the importance of model governance. Featurespace follows regulatory guidelines addressing services that Featurespace offers, including model governance for machine learning models. In working with customers globally, we have enabled teams to demonstrate compliance while providing an immense uplift in operational efficiency in regions across EMEA, the US, APAC, and LATAM.

Regulations and guidance are expanding and changing to address machine learning rather than simply rule and consortium-based approaches. In future we expect more, rather than less regulation and oversight in these areas in the future. Our approach and discipline focus on the core principles

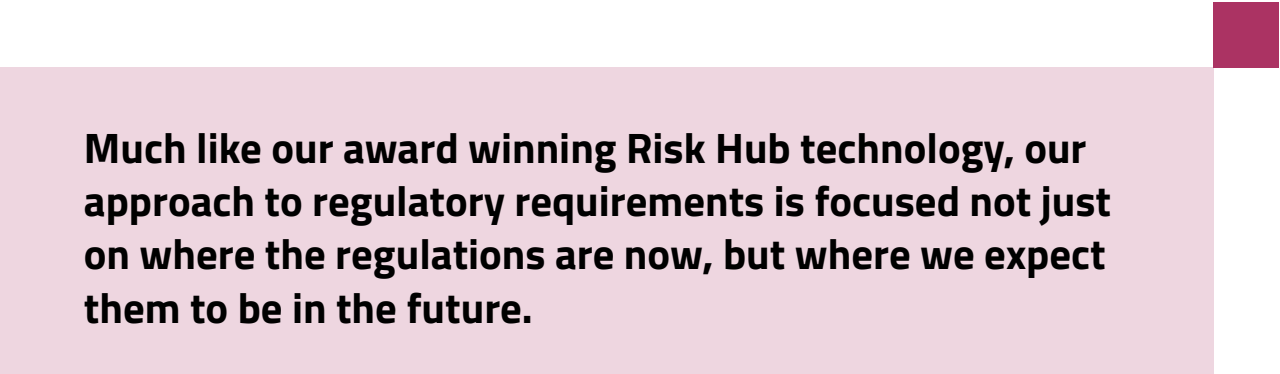
and intent of these legislative regimes, rather than just outcomes or mechanical, prescriptive solutions.

Much like our award winning Risk Hub technology, our approach to regulatory requirements is focused not just on where the regulations are now, but where we expect them to be in the future.

This dynamic approach has enabled Featurespace to satisfy the requirements of the stringent governance teams of the world's largest financial institutions and payment processors.

In addition to enabling innovation, governance facilitates fairness in the application of machine learning. The importance of removing prejudice from behavioral profiling while identifying as much suspicious activity as possible cannot be overstated.

With fairness and innovation in mind, this paper will outline Featurespace's approach to model governance in its application of machine learning models in the field of AML.



Much like our award winning Risk Hub technology, our approach to regulatory requirements is focused not just on where the regulations are now, but where we expect them to be in the future.

3 Model Governance for Machine Learning in AML

Before we explore the governance processes Featurespace employs when creating machine learning models for AML, it is important to understand the risks its customers are facing.

The US Office of the Comptroller of the Currency (OCC) was arguably the first regulator to outline the potential risks of machine learning in AML transaction monitoring. These risks are outlined in its 2011 bulletin¹:

Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency:

“Models can improve business decisions, but they also impose costs, including the potential for adverse consequences from decisions based on models that are either incorrect or misused. The potential for poor business and strategic decisions, financial losses, or damage to a bank’s reputation.”

It goes on to say that ‘model risk can be diminished but not eliminated,’ and that therefore model risk needs to be managed. The OCC highlights several possible risks to the use of machine learning models in banking: when models play a material role is the essence of “model risk”.²

1. “Sound Practices for Model Risk Management: Supervisory Guidance on Model Risk Management”. <https://www.occ.gov/news-issuances/bulletins/2011/bulletin-2011-12.html> 4 April 2011. Retrieved 6 April 2021

2. “Supervisory Guidance on Model Risk Management”. <https://www.occ.gov/news-issuances/bulletins/2011/bulletin-2011-12a.pdf> 4 April 2011. Retrieved 6 April 2021

3.1 Risks outlined by OCC

The risks that financial institutions may face when using machine learning to detect and prevent financial crime have been summarised below:

- 1. The model design may be flawed**
- 2. The choice of sampling may either be poor or limited**
- 3. Algorithms may have mistakes that mean they fail to carry out the purpose of the model**
- 4. Shortcuts or simplifications used to manage complicated problems could compromise the reliability of outputs**
- 5. The quality of data input may be insufficient, leading to sub-optimal, or even useless results**
- 6. The user may not understand the limitations of the model, and expect the model to do too much**

These risks must all be addressed through model governance to apply machine learning in financial services safely and effectively. In the next section, we will explore how Featurespace assists its customers in mitigating the risks of machine learning model building for AML, enabling them to enjoy the full benefits of an innovative risk management solution.

4 Featurespace's Approach to Model Governance

With over 30 years of machine learning expertise stemming from Cambridge University, Featurespace has used machine learning to process billions of transactions for financial institutions, payment processors, and gaming companies.

Inventors of Adaptive Behavioral Analytics and Automated Deep Behavioral Networks, Featurespace has put significant resources into ensuring model governance requirements are met for all machine learning models produced for regulated entities.

Regardless of the type of machine learning model the Featurespace data scientists are building, their strict model governance process focuses on four main questions:

1. Is the model as performant as possible?

2. Can the model's decisions be understood?

3. Is the model as fair as possible?

4. Is the model stable in production?

Model governance centers around two processes - audit of the model before deployment, and monitoring of the model once it is deployed. The model governance approach that Featurespace takes is also dictated by the approach to model building.

Whereas model governance frameworks for off-the-shelf machine learning models for AML transaction monitoring can be standardized from implementation to implementation, Featurespace takes a more tailored approach to model building.

This means that the model governance varies across each deployment, ensuring that models are validated and tested against each customer's specific data. The result is not only a model best suited to our customers' use case but a governance framework that truly takes our customer's requirements into account.

4.1 Model Performance


A machine learning model's performance is determined in the preparation of training data and is tested throughout the model build process.

Featurespace machine learning models are adaptive and self-learning, meaning they require no manual retuning.

With such sophisticated models, optimal performance is always a key requirement for Featurespace data scientists.

The Featurespace approach is based on collaboration, and underpinned by governance at every step.

Model Documentation Process



At the beginning of the planning process, the machine learning model's purpose, performance requirements and limitations are outlined and documented. This impacts some of the technical choices the data scientists will make based on the specific use case and requirements of each customer.

Alongside the standard data schema and offering that Featurespace has prepared for each solution, the data scientists and SMEs then collaborate to select the specific set of features and reason codes. These describe customer behavior based on data analysis and interviews with our customer's in-house financial crime experts. This is to ensure that the very highest-scoring alerts are of interest to those responsible for investigating them. Moreover, those responsible for reviewing the alerts will be able to identify if there are any risk signals present, that the model is not identifying clearly enough. Customer reviews can show valuable context on the background and investigation process that alerted transactions have been subject to, which will identify whether the model is identifying truly suspicious activity.

During the model build, Featurespace carries out data health checks and data validation activities. To create the most effective machine learning models, it is important to spot poor quality or biased data and 'noise' that may affect the training of machine learning models. These checks and results are recorded fully in the model governance documentation.

Testing using quantitative and qualitative methods helps ensure model performance is high before the initial model output is reviewed with our customer. This provides an opportunity to receive feedback on false positives and worthwhile alerts, with rationale that helps our data scientists tune and refine features and reason codes.

Taking this collaborative approach ensures that our customers understand the model build process. Rather than providing an unexplainable, obscure box, customers can become familiar with the model outputs, which benefits later use post-deployment.

4.1a Data Validation

Data validation can be thought of in two parts: offline (on historic data) and online (in production).

Offline data validation is part of the data health check and exploratory data analysis that take place on the historic dataset that is used to understand the data to build and test models. In this phase, data are analyzed with great attention, and potential mistakes in data format or content are communicated to the customer that can therefore fix their data extraction process.

Online data validation is mainly applied via data schema validation. ARIC is configured with a specific project's data schema that determines the expected format and content for each event type that is part of the data stream. This includes things like attribute names and types, their minimum and maximum lengths, matching with regular expressions (for fields like dates or timestamps that must be provided in a specific format). Each incoming event is validated with respect to the schema definition. Events that pass the validation are ingested, whereas those which fail the checks are rejected and placed in a "failed events" queue. This ensures that the engine only receives events in the expected format, and at the same time it informs the customer of potential flaws in their data stream.

The Featurespace engineering team has been developing new tools to further enhance the data validation process.

The aim of the research is to be able to train a set of constraints, or 'expected bounds' on the historic dataset that can be then used, once live, to ensure that live data are aligned with the historic ones. For example, the tool can learn the proportion of events for each of the defined event types, the

number of events that have null values in any specific attribute, the observed cardinality of categorical fields, the most common values of the same, the total volume per period, and so on.

This information is computed on batches of data and compared with what was learned to be 'normal'. If the new quantities are outside the configured boundaries, the tool generates data validation alerts.

4.1b Model Validation

Once the model has been built, the data science team presents it to an independent model validation team. This team is comprised of Featurespace data science experts who have no knowledge of the project. This means they can interrogate the model build process without context, removing any bias. This process takes into consideration and replicates each customer's model validation standards.

The team carries out functional and analytics testing of all models executed both during the model build process and on an ongoing basis after deployment.

This includes:

- 1. Simulation of extreme cases such as unusual events, missing data, or duplicated data. This tests the soundness and functional robustness of the models**
- 2. Stability metrics compare how the model will cope with missing or mislabeled data**
- 3. Score distribution analysis, which tests how the model segments the events in the data**
- 4. Performance assessment online and offline that show the sensitivity and specificity of the model**

Alongside data validation, model validation aids Featurespace data scientists in ensuring the most performant machine learning models for each customer.

4.2 Understanding Model Decisions

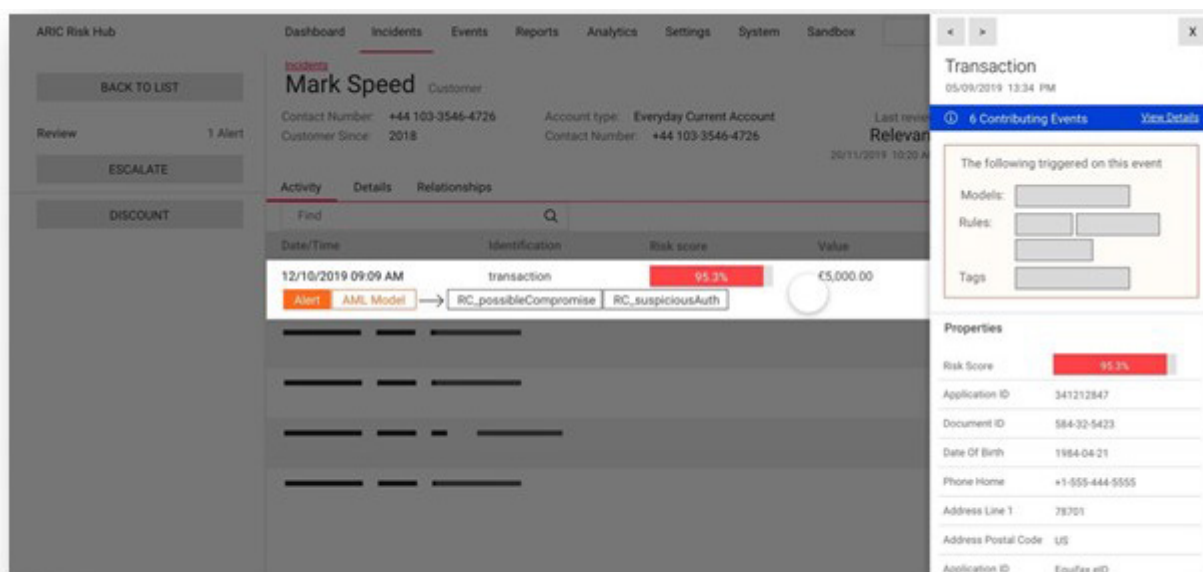
At Featurespace, model explainability is not simply a case of adhering to regulation. Making the output of machine learning models understandable also provides investigators with the context they need to analyze alerts further. If the investigators understand why each alert is raised, they are in a better position to work the alerts.

Therefore, providing explainable models has a two-pronged benefit. Regulatory adherence and investigation support. To this end, Featurespace has put significant research time into ensuring that customers are provided with clear reason codes to illustrate why the machine learning models have generated specific alerts. ARIC Risk Hub user interface

There is a perceived risk with explainability that the more explainable a model is, the less performant it tends to be. At Featurespace we prioritize top analytical performance, so we had to solve this issue.

Our researchers found the best technique to provide explainability while still providing the highest performing machine learning models. The heuristic technique used shows to what extent each feature contributed to a risk score. This produces human-readable reason codes for each alert, which can be shown to the regulator along with relevant customer and transactional data.

ARIC Risk Hub user interface



Alert includes explainability-related features demonstrated with reason codes or sent to an external case management solution.

4.3 Model Fairness

As with everything Featurespace does, ensuring discrimination is left at the door is a key part of our machine learning model build process. One may think that using machines to process data removes the bias that humans inherently carry, but as it is humans that produce the data and program the machines, further steps must be taken to ensure machine learning models are fair and ethical.

Our models are trained without data attributes that identify a protected population. These protected attributes will be identified on a case-by-case basis with

our customers whether they are using a standard or bespoke data schema to ensure they are not incorporated into the model. Prior to deployment, we test our models to ensure there is no unintended disparate impact on those protected populations.

A thorough analysis of historical data can also reveal any biases hidden within the data, for example from a previous decisioning system.

When testing our models for ethicality, we ensure the following questions are answered:

- 1. Is the model serving a clear purpose?**
- 2. Are working practices robust?**
- 3. Is the work clear and transparent?**
- 4. Are we using data responsibly?**
- 5. Is the data proportionate to the need of the machine learning model?**
- 6. Does the model align to Featurespace's policies, and those of our customer?**

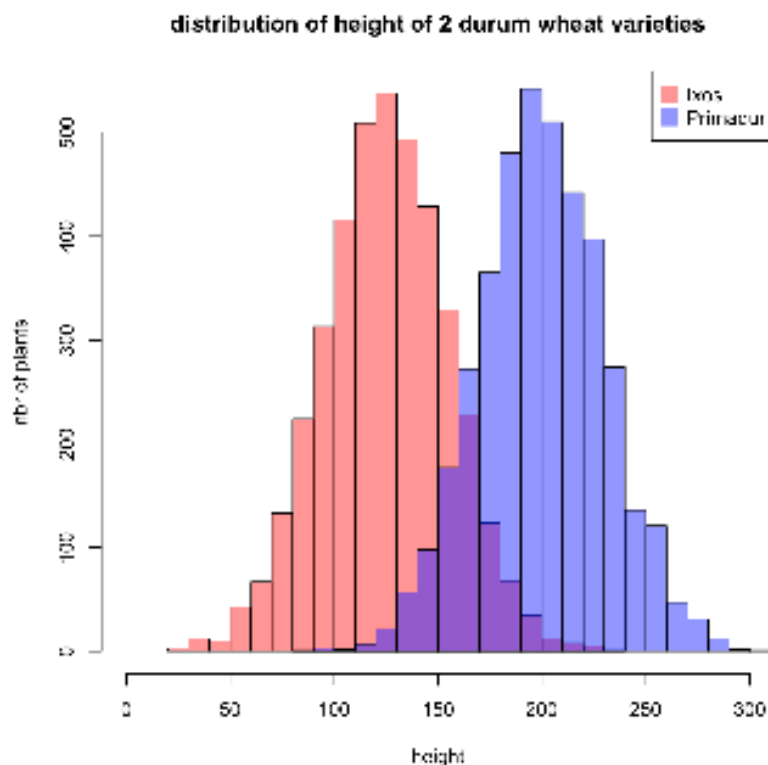
4.4 Model Stability

Every Featurespace model is subjected to a battery of stability and sensitivity tests to ensure the model is performing under production conditions and changing circumstances. We are rigorous in these analyses before declaring a model production-ready.

This testing includes the simulation of major operational issues and disaster scenarios, such as data outages, as well as the simulation of extreme cases, which may happen on such events as Black Friday sales or pandemic lockdowns. All Featurespace models are designed to gradually adapt to shifts in behavior, once it becomes clear that


the behavioral change represents a new normal. This has, for example, allowed our systems to quickly adapt to the new realities of COVID-19, as our case study with TSYS demonstrated.

Another major part of stability testing is feature distribution monitoring. Although Featurespace models are capable to adjust and adapt over time, we still want to ensure we prevent any performance degradation if there are significant shifts in feature distributions. Feature distributions are displayed in a histogram, which shows what values we expect a feature to have.



This simple histogram example showing heights of durum wheat.

"The R Graph Gallery" <https://www.r-graph-gallery.com/histogram.html> Retrieved 7 April 2021



The stability of Featurespace's adaptive machine learning models allowed our systems to quickly react to the new realities of COVID-19, as our case study with TSYs demonstrated.

There are certain values we expect our machine learning features to show. Monitoring shifts in feature distribution helps us detect potential instability in the machine learning model.

The main precaution to prevent big shifts in feature distribution is good feature engineering. This is something that Featurespace data scientists pay a lot of attention to during model build. Nevertheless, there may be real world circumstances or other data anomalies that could still lead to shifts beyond our control.

After the model has been built and deployed, continuing to monitor the features, and catching any shifts as soon as possible is key to maintaining stability. This can be run offline and online once just before going live to generate the first report, and then continuing at periodic intervals.

Each time the task is run, some statistics around the feature distributions are computed: proportion of most common tokens for categorical features, proportions of "bin size" (the value parameters - durum wheat height in our previous example), in the histogram of numerical features, proportion of true and false in boolean features (which only require a true or false answer), and proportion of null values for each feature.

The statistics of the latest run are compared to those of the previous report, and the shift between these is measured. If this is above certain configurable thresholds, then alerts are generated.

Associated reports are always generated even if there are no alerts and can be used by either the customer or by Featurespace, if we are hosting and monitoring the system, to assess the ongoing behavior of the feature distributions.

5 Case Study: Tier 1 Global Bank

In search of the best AML solution, this Tier 1 Global Bank selected Featurespace following a head-to-head challenge to demonstrate the benefits of machine learning coupled with automation.

A machine learning model to reduce alerts, minimize false positives, and automatically prioritize alerts was built and deployed in the Asia-Pacific region.

As it was a large global bank, there were many layers of complexity in the model build. Whereas many model governance frameworks in such complex builds for large organizations can become lengthy and complex themselves, Featurespace was proactive in planning to ensure the smoothest process possible.

Model governance requirements were identified early on in the project, with

additions made closer to deployment as new teams were brought in to ensure a deployment that would meet regulatory requirements in multiple jurisdictions.

The flexibility of the Featurespace team and processes ensured that these additions could be made, and the model produced was award winning.

Reasons for selecting Featurespace encompassed our unsurpassed capability to deliver excellent results and support other financial crime risk controls requirements.

Equally important in selection was Featurespace's ability to describe the machine learning models and decisions as applied to insurance customer data (i.e., a non-proprietary approach to delivering machine learning models).

6 Conclusion

Model governance is not only a regulatory matter. Compliance is a major part of building machine learning models for financial crime teams, but keeping strict processes and documentation also allows us to serve our customers better in their financial crime investigations.

Even more important, it is the ethical thing to do, and Featurespace strives to ensure an ethical approach to all aspects of the organization by capturing the fairest way to conduct business in our documents and processes. By constantly reassessing and updating these frameworks, we can not only help our customers beat the criminals, but we can also help them do this the right way.

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



THE QUEEN'S AWARDS
FOR ENTERPRISE
2018



AITE IMPACT MATRIX (AIM)

BEST IN CLASS

"AIM Evaluation: Fraud and AML Machine Learning Platform Vendors.", AITE, 2019



Technology Fast 50
2020 UK REGIONAL WINNER
Deloitte.



BEST SECURITY OR
ANTI-FRAUD DEVELOPMENT

Category Winner

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