Mphasis
The Next Applied

# Anomaly Detection Using Azure Quantum QIO (Quantum Inspired Optimization) Platform

Whitepaper by Ashutosh Vyas, Senior Data Science Manager - Mphasis NEXT Labs |
Lauren Roberts, Data Scientist - Mphasis Datalytyx | Saurabh Gupta, Senior Data Scientist - Mphasis NEXT Labs |
Rohit Patel, AVP – Data Science and Quantum Computing - Mphasis NEXT Labs

# Contents

# 1.
# Introduction

An anomaly is a rare event or observation in the behavior of a system. Such anomalies, if not detected, may lead to financial losses for the business and result in customer dissatisfaction and mistrust in business operations and transactions. The anomaly detection methodologies help in identifying anomalies in the behavior of a system/process, thereby actuating subsequent decision support systems.

A few real-life examples of anomaly could be:

- In the banking domain, anomaly can be considered as unauthorized transactions through the system. This may include money laundering, abnormal volume of transactions, etc.
- In industrial production domain, anomalies could be irregular machine behavior, under performance of the employees, poor production output, etc.
- In agriculture domain, anomalies could be irregular development of plants, sudden deterioration in plant health, sudden reduced production volume, etc.
- In healthcare domain, anomalies could be unwanted symptoms in patients under medication, sudden deterioration of organ parameters, slow recovery in health, etc.

## 1.1 Type of Anomaly Detection Systems

In general, anomaly detection systems can be classified into two major categories, i.e., Event-based anomaly detection and Monitoring-based anomaly detection. These categories are based on the type of data captured as part of the organization's monitoring and compliance strategy and business requirements.

### 1.1.1 Event-based Anomaly Detection

Event-based Anomaly Detection algorithms are designed for use cases where each observation or event is independent of its previous observation or event. A few use cases could be:

- Credit card fraud detection
- IoT sensor anomaly detection
- Network intrusion detection
- Retail customer behavior drift detection
- Security surveillance

### 1.1.2 Monitoring-based Anomaly Detection

Certain problems in the anomaly detection domain requires constant monitoring of the actions or processes in the business operations. Constant monitoring helps in sequence-based analysis of the pattern, as individual events could look non-anomalous, but the pattern reflects the anomalous activity. A few of the business use cases which come under monitoring-based anomaly detection are as follows:

- Healthcare monitoring
- Weather monitoring
- IT infrastructure monitoring
- Industrial production line monitoring

In the coming sections, we will discuss the criticality of anomaly detection for businesses and how quantum inspired optimization can bring in more value to the performance of the anomaly detection algorithms.

# 2.

# Business Challenges & Drivers to Build Robust Anomaly Detection Systems

Anomaly detection is a critical system for business operations. The fundamental aim to identify anomalies is to protect businesses from incurring financial losses due to unnoticed and unauthorized actions and processes. Building a robust anomaly detection system is hindered by the following factors:

- **Poor understanding of the data:** Due to the high volume, velocity and variety of transactions happening on the system, it becomes difficult to build systems to understand and analyze complex patterns in the data. This leads to unnoticed and unauthorized usage of the system.
- **Fewer audits:** Majority of the anomalous actions are recognized once they are reported by the users or customers. Audits, if any, due to the complexity of analysis of the data, systems and workflows, are fewer in number. This leads to non-availability of the right data to build a robust system/model.

Anomaly detection can add a lot of value to businesses in the areas of customer satisfaction, operational performance, risk, case resolution, etc. Following are the benefits a business could leverage by integrating anomaly detection modules in their operations:

## 2.1 Customer Satisfaction

Anomaly detection systems can help a business achieve high customer satisfaction through:

- **Safeguarding customers' assets:** Anomaly detection systems protect the customers' assets from being compromised - mitigating the financial losses

- **Managing access to assets:** Anomalies generally block customers' access to assets until anomalies are resolved. This can lead to financial and personal hardships for the customers. By detecting and notifying the anomalies, anomaly detection systems help avoid such scenarios.
- **Increasing customer faith:** Anomaly detection systems pro-actively detect and manage anomalies leading to a reduction in the occurrence of fraudulent activities, hence building customer faith and company image-impacting future business positively.

## 2.2 Business Operational Performance

Anomaly detection systems can help improve the operational performance of a business through:

- **Reduction in operational cost:** Unchecked anomalies, other than causing immediate financial loss, leads to operational overhead in investigating and handling anomalies such as fraud. Anomaly detection systems can prevent the occurrence of anomalies and thereby reduce operational costs.
- **Reduced human intervention:** Anomaly detection systems proactively stop and help resolve anomalies. This reduces the human intervention required to prevent anomalies.
- **Faster response time:** Anomaly detection systems can help resolve the anomalies in real time or reduce the time to root cause analysis.
- **Robust operational workflows:** Backtracking and root cause analysis of hidden behavioral patterns using anomaly detection systems can lead to the identification of potential loopholes in the operations which helps improve the design and re-structuring of data flow architectures to reduce potential threats.

# 3. Machine Learning-based Anomaly Detection Systems

Present legacy rule-based anomaly detection systems which work in predefined rules and statistic thresholds for anomaly detection, have a major drawback of generating a huge number of false positives. Thus, tremendous human intervention is required to resolve different false positive cases. This makes the non-anomalous transaction process slow and laborious. This problem can be resolved by introducing a machine learning-based anomaly detection system. Before proceeding to how to design such a system, the following section explains some of the technical challenges to achieve the same.

## 3.1 Technical Challenges

Anomaly detection is a complex classification problem. The major technical challenges in the development of a good performing anomaly detection solution are as follows:

- **Appropriate feature extraction:** Selection of features is an important and difficult exercise in anomaly detection. Domain understanding is required to capture hidden patterns and relationships that can affect model development. Good features correspond to a better model performance.
- **Defining normal behaviors:** Event-based data or monitoring sequences need to be analyzed to develop an understanding of the behavior of a system/application. This further helps in defining or identification of normal behavior patterns or events of the system.
- **Handling imbalanced distributions:** There exists a wide imbalance in terms of the amount of information available for non-anomalous rows in comparison to anomalous rows. This imbalance ratio is mostly due to system/application incapability to capture anomalous behavior and treat it as a normal pattern. Specific strategies are required to handle this imbalance.
- **Threshold for an anomaly:** Deciding what a suitable cut-off point is for a result to be anomalous. The data is subject to random variation, so the algorithm must distinguish between slight deviations from a normal pattern and an actual anomaly.
- **Proportion of data for training:** The number of samples used for training the algorithm must be chosen to avoid overtraining, but to keep enough information in the data to learn what normal patterns look like.
- **Bias of the algorithm:** The algorithm used must avoid having a bias against a particular group, which could lead to disproportionately labeling that group as anomalous.

## 3.2 Quantum Annealer-based Anomaly Detection

There are two approaches for probabilistic modeling - Discriminative and Generative models. Discriminative models perform conditional probability evaluation to make predictions and generative models use joint probability evaluation to make predictions.

The behavior learning paradigm prefers the usage of generative models because it captures the inter-relationships between model features to learn the data distribution and a wide range of class information and parameters. These algorithms use stochastic sampling of underlying distribution to get the point estimation values for the parameterized distributions.

Also, semi-supervised anomaly detection approaches have demonstrated the capability to learn and understand the normal behavioral patterns of the data and once trained, they raise alarms to indicate abnormal patterns. Thus, learning the underlying statistical distribution of the data becomes an important objective of the solution approach.

State-of-the-art classical machine learning-based anomaly detection systems adopt generative semi-supervised machine learning approach to identify anomalies. These approaches try to learn and understand the normal behavior patterns in the data using an unsupervised algorithm and then build a supervised threshold classifier using the tagging information of the patterns and

output generated from the unsupervised algorithm for all the patterns. The classical deep learning architecture uses Restricted-Boltzmann-Machine, Autoencoders, GANs, etc., in deep learning domain to learn and understand the statistical patterns of normal behaviors in the data. The deep learning models train a vast number of weights to learn statistical data distribution. The training process requires a high volume of input data to learn the patterns.

### 3.2.1 Why Quantum Annealer?

In this whitepaper, we illustrate an alternative semi-supervised generative approach to build an anomaly detection system using quantum annealer for machine learning objective. ***Classical hardware implementations of quantum annealer-based machine learning can approximate complex density function from comparatively smaller data sets using quantum-inspired sampling techniques and properties of probabilistic search space exploration.***

We have used Azure Quantum optimization solvers' algorithms to perform stochastic learning by generating samples from the underlying data distribution akin to classical sampling techniques such as Markov Chain Monte Carlo (MCMC). Annealing-based sampling brings improved quality of the point estimation by reducing the correlation between the samples. This improves the precision of the trained model, thus reducing the number of false positives at the time of inference generation.

The quantum-inspired approach has the advantage of better training KPIs (Key Performance Indicators) and faster training time of the model. This stems from the fact that MCMC sampling strategies are sequential in nature. Thus, point estimation from the underlying data distribution becomes a lengthy process and these strategies require a burn-in time to stabilize before generating samples and due to this, the initial samples are rejected. All these drawbacks are not present in Quantum-Inspired Optimization (QIO) solvers as they perform parallel sampling and do not require any specific sampling strategies or burn-in time.

## 3.3 Mphasis Solution Approach

In this section, we present the Mphasis solution architecture for the event-based anomaly detection problem and illustrate the use of QIO algorithms for model training.

We opted for a semi-supervised ensembled modeling technique to detect event-based anomalies. We used a collection of Quantum-Restricted Boltzmann Machine (QRBM) models to learn the underlying data distribution, where a smart sample selection technique is used to provide data to each QRBM model. This strategy helps in learning the different aspects and behaviors in the data and avoiding overfitting of the model.

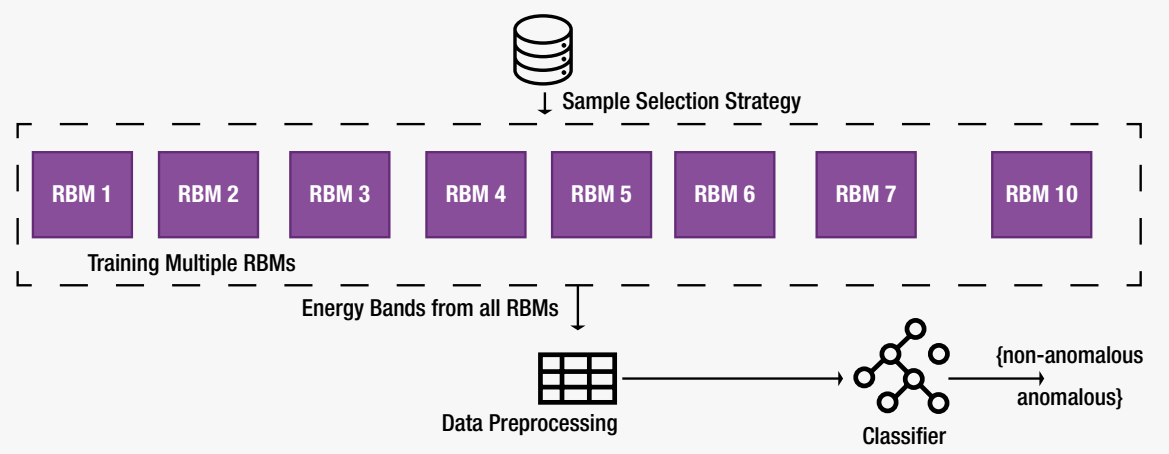


*Figure 1: Solution Architecture*

The above architecture describes the different components we have used for design and development of event-based anomaly detection solution. The training procedure is as follows:

**Data type:** The solution approach is semi-supervised in nature and thus requires tagged data for training purposes. The solution can handle high imbalances in the data.

**Data segregation:** As the architecture is semi-supervised, non-anomalous data is used for RBM training.

**Strategic sampling for each RBM:** Data distribution-based sampling strategy is used to generate samples from the event-based data set. Sampling data sets with % of shared data points are used to train each RBM.

**Learning data distribution using RBM:** Each sample dataset is provided to each RBM for learning the associated data distribution. This type of ensemble-based approach is utilized to control the overfitting possibilities and to develop a generalized learning approach.

**Energy-band data collection:** RBM is a generative neural network training approach which tries to regenerate the input data. The model is considered as trained if the model can generate the output like the input given. The architecture of RBM can be considered as an undirected graph, and using the Hamiltonian energy expression, an associated energy of the graph can be calculated. The energy of the graph changes with the changing input to the architecture. This energy is also called free energy and the expression is as follows:

$$F(v) = -a^{T}v - \sum_{j} \log\left(1 + \exp(b + W^{T}v)\right)_{j}$$

In the above expression, “v” represents the input vector, “a” is the bias matrix for the visible layer in RBM, “b” is the bias matrix for the hidden layer in RBM and “W” is the cross-sectional weight matrix between visible and hidden layer in the RBM. “j” represents the hidden node in the hidden layer.

We can thus compute a scalar energy value associated with each input we provide to RBM while training. These scalar energy values are collected for all the datasets that we provide to different RBMs.

**Energy data normalization:** The collected free energy values for anomalous and non-anomalous data from all the RBMs are collected and normalized.

**Energy threshold classifier:** A supervised learning algorithm is used to identify a threshold between the non-anomalous data associated energies and anomalous data associated energies. This classifier is used to finally generate inference to identify anomalous/non-anomalous events.

### 3.3.1 RBM Training Using Azure Quantum Optimization Platform

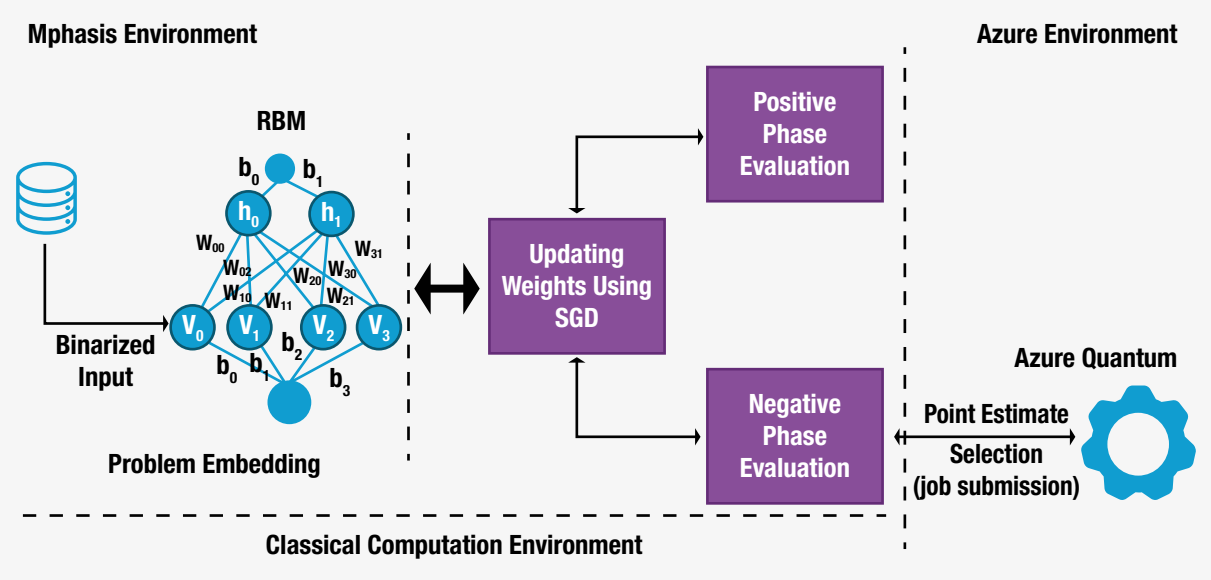


*Figure 2: RBM Training Process Using Azure Quantum Optimization Platform*

The above workflow demonstrates the process of training the RBM using Azure Quantum optimization solvers for probabilistic sampling and stochastic gradient descent to train the weights of the RBM model.

Restricted-Boltzmann Machine model is a generative model which uses probabilistic sampling to train. The RBM model tries to learn the underlying data distribution by using Boltzmann distribution to approximate the actual data distribution. The expression of Boltzmann distribution is:

$$p_i \propto e^{-\mathcal{E}_i/kT}$$

Where "pi" is the probability of a state "i" of a system whose energy at state "i" is "$\mathcal{E}$i" where "k" is Boltzmann's constant and "T" is the temperature of the system.

Now to embed the RBM network into Boltzmann's distribution, RBM network energy needs to be computed. To perform this task, we utilize the resemblance of Ising model and RBM architecture. The RBM neural network can be treated as an undirected connected graph. The structure is as follows:

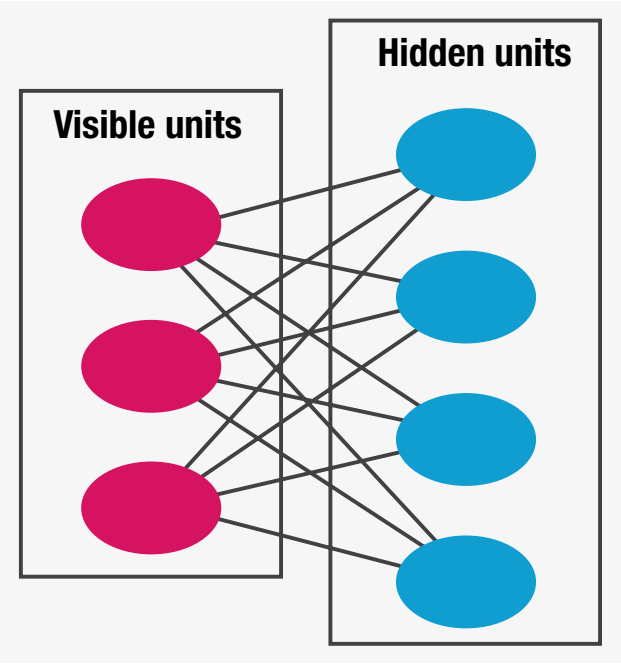


*Figure 3: RBM Architecture*

Where visible layer is exposed to incoming input data and hidden layer is treated as a latent layer and the node in the RBM architecture can take only two values. The above architecture resembles the Ising model architecture, which is:

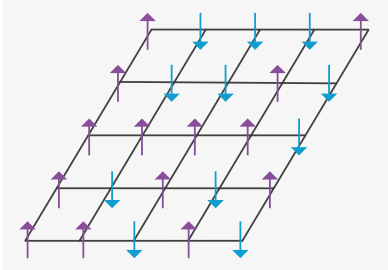

*Figure 4: Ising Model*

The Ising model is a mathematical model of ferromagnetism in statistical mechanics. The nodes in the model can take two values – either -1 or 1, which represents the spins of the magnetic particle. The energy associated with this model is expressed using the Hamiltonian expression:

$$H(\sigma) = -\sum_{(i\,j)} J_{ij}\ \sigma_i \sigma_j - \mu \sum_j h_j \sigma_j,$$

This expression has “J” as the weight matrix associated with adjacent nodes “sigma” and individual bias of each node as “h”. Keeping the above expression in consideration, RBM network energy can be computed as:

$$\sum_{(i\,j)} w_{ij}\, v_i\, h_j + \sum_i a_i\, v_i + \sum_i b_j\, h_j$$

Where “$w_{ij}$” is the cross-sectional weight matrix of the RBM architecture between hidden and visible layer nodes, “$a_i$” is the bias of visible nodes “v” and “$b_j$” is the bias for hidden nodes “h”.

The analogy between the energy expressions helps to use Boltzmann’s distribution for input data distribution approximation.

The RBM performs training of models by using stochastic gradient descent for updating the weights of the network and probabilistic sampling on associated Boltzmann distribution to estimate the values of the hidden layer nodes, i.e., 0 or 1. The process of probabilistic sampling and weights update is an iterative process, and it terminates when the RBM can generate the same incoming input with high accuracy.

The quantum-inspired solver accepts the problem expressed in terms of Ising model where variables can take only two values and are represented as nodes in the Ising model. The solver explores the configuration space to identify the optimal values of the nodes, i.e., either 0 or 1 by expressing the search space in terms of associated configurations Hamiltonian energies.

The RBM architecture resembles the Ising model architecture as nodes in both, where the networks take only two possible values, and both the models have adjacency weight matrix and bias weights. Thus, the RBM network can be mapped to Ising model expression, thus solvers in the Azure Quantum optimization platform can perform probabilistic sampling on associated Boltzmann distribution by expressing the RBM architecture configuration in Hamiltonian energy space exploration to return point estimation of the hidden layer nodes.

The point estimation value returned by quantum-inspired solvers is further utilized to perform stochastic gradient descent operation to update the network weight by back propagation. This process continues until the output of the RBM model resembles the input with high accuracy.

### 3.3.2 Azure Quantum QIO Workflow

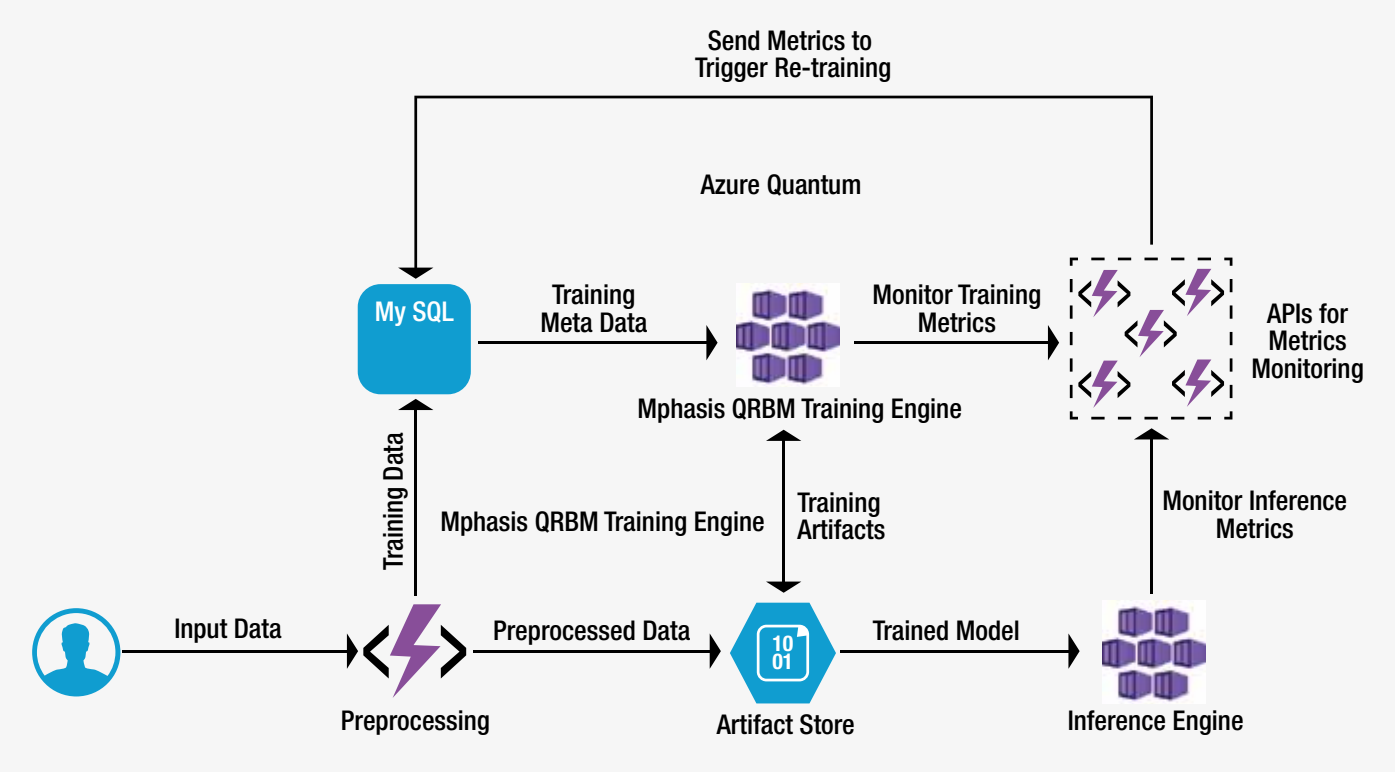


*Figure 5: Integrated Azure Cloud and Azure QIO Workflow for Hybrid Quantum-inspired ML*

Azure provides different architectural components to develop and deploy a quantum-inspired anomaly detection solution. The basic reference diagram is shown above. The cloud-based solution requires different components to perform input data preprocessing, model training and generate inferences on real-time streaming data. The following points describe the Azure components and their utility in the solution architecture:

- **Serverless Azure function app** is used to perform the required preprocessing of the input data and evaluate algorithm performance using different matrices and parameters.
- **Azure Kubernetes Service** is used to perform training on the historical data and inference generation on the real-time streaming data as these modules require heavy computation power for data processing.
- **Azure blob storage service** is used to store the trained model and statistical parameter files of the finalized model that would be used to generate inference. This input data and the meta data for the entire application can be stored in **My SQL Db** service provided by Azure.
- **QIO solvers** are part of the Azure Quantum Optimization suite of tools used to perform sampling based on physics-inspired processes while training the model. The **Azure QIO solvers** are called via **Azure Kubernetes Service** in the **Hybrid Quantum-Inspired ML workflow** described above.

## 3.4 Experiment & Results – Credit Card Fraud Detection Using Azure Quantum Optimization Solvers

**Data Description:** Credit card transaction dataset containing legitimate and fraud transactions from the duration 1st Jan 2019 - 31st Dec 2020. It covers credit cards of 1000 customers doing transactions with a pool of 800 merchants. The original dataset contains over 1 million transactions and we have used a small section of it to train our models, as quantum annealers can approximate underlying density function with a relatively smaller dataset.

**Experiment 1: Comparing Different QIO Solvers**

- Each sample size: 22000 approx.
- Number of samples: 9
- Number of frauds in each sample: 100-150 approx.
- Used for training: 8
- Number of binary independent columns: 95
- Test sample: 1
- Class ratio: 1:147 approx.

**Results:**

| Performance/Azure Solvers | | Azure Parallel Tempering | Microsoft QIO Simulated Annealing | Azure Quantum Monte Carlo | Azure Population Annealing | Azure Substochastic Monte Carlo |
|---|---|---|---|---|---|---|
| Training Results | Accuracy | 100% | 100% | 100% | 100% | 100% |
| | Precision | 100% | 100% | 100% | 100% | 100% |
| | Recall | 100% | 100% | 100% | 100% | 100% |
| | F1 Score | 100% | 100% | 100% | 100% | 100% |
| | Training Time (mins) | 21 | 22 | 24 | 32 | 25 |
| Testing Results | Accuracy | 99.54% | 99.94% | 99.95% | 99.93% | 99.96% |
| | Precision | 100.00% | 100.00% | 100.00% | 100.00% | 100.00% |
| | Recall | 99.54% | 99.94% | 99.95% | 99.93% | 99.96% |
| | F1 Score | 99.77% | 99.97% | 99.98% | 99.97% | 99.98% |

**Experiment 2: Azure Quantum - Quantum Monte Carlo Solver Results**

We ran numerous tests with different parameters on Azure Quantum - Quantum Monte Carlo Solver. Following are the best results we got on the tuned parameters.

- Each sample size: 22000 approx.
- Number of samples: 9
- Number of frauds in each sample: 100-150 approx.
- Used for training: 8
- Number of binary independent columns: 95
- Out of sample testing: 1
- Class ratio: 1:147 approx.

- **Parameters:**
  - sweeps = 2, trotter_number = 10, restarts = 72, beta_start = 0.001, transverse_field_start = 10, transverse_field_stop = 0.1, seed = 22

**Results:**

Total time to train the model = 24 mins.

| Sl. No. | KPI | Training | Testing |
|---|---|---|---|
| 1 | Accuracy | 100% | 99.95% |
| 2 | Precision | 100% | 100.00% |
| 3 | Recall | 100% | 99.95% |
| 4 | F1 Score | 100% | 99.98% |

# 4. Solution Benefits

## 4.1 Classical vs Quantum-inspired RBM Classifier

**Experiment 1:** The comparison of Classical RBM and Azure Quantum - Quantum Monte Carlo Solver annealer-based RBM classifier:

- **Dataset Description:** Credit card transaction dataset containing legitimate and fraud transactions from the duration 1st Jan 2019 - 31st Dec 2020. It covers credit cards of 1000 customers doing transactions with a pool of 800 merchants.
- **Training data size**
  - Class 0 (non-anomalous) - 25873 transactions
  - Class 1 (anomalous) – 1501 transactions
  - % of anomalous transactions – ~5.5%
- **Architecture setup:**
  - Quantum annealer-based:
    - Number of QRBM: 1
    - Hidden Nodes: Equal to visible layer nodes
  - MCMC-based:
    - Number of QRBM: 1
    - Hidden Nodes: Equal to visible layer nodes

- **Training results:**

| Sl. No. | KPI | Azure Quantum Results | Classical Results |
|---|---|---|---|
| 1 | Precision | 100.00% | 100.00% |
| 2 | Recall | 99.86% | 97.00% |
| 3 | Accuracy | 99.96% | 99.80% |
| 4 | F1 | 99.98% | 98.50% |
| 5 | Training Time | 4 secs | 35 mins approx. |

The above table clearly shows the improvement in the training performance of the Azure Quantum-powered RBM because of the quantum-inspired sampling while training the model.

- **Test results:**
  - Total transactions: 9794 transactions
  - Class 0 (non-anomalous): 9344 transactions
  - Class 1 (anomalous): 450 transactions

The graph below represents an improvement in the performance of quantum-inspired approach in comparison with classical approaches of training the RBM.

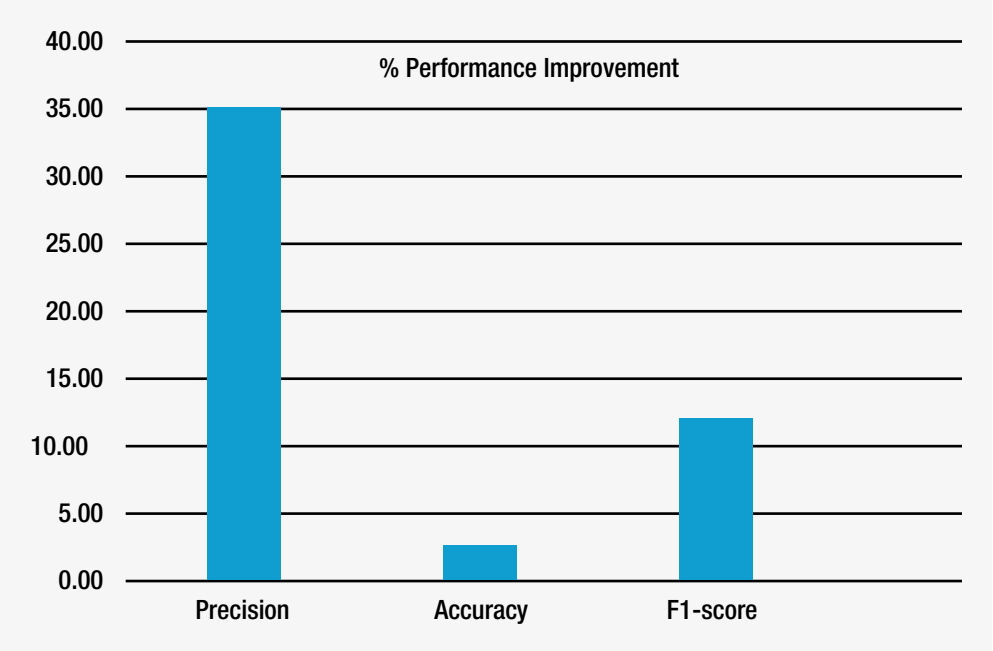


*Figure 6: Performance Results*

## 4.2 Other Benefits

Our anomaly detection solution based on Azure Quantum's optimization solvers delivers the following benefits:

- **Improved KPIs** such as improved training accuracy, shorter training times and faster inferences
- **Real-time anomaly detection** through API (Application Program Interface) with QIO in Azure Quantum backend
- **Capture more complex patterns** using Quantum-inspired Algorithms
- **Production-ready cloud-based** integrated Azure cloud and Azure Quantum pipelines for solution delivery
- **Requires smaller datasets** in comparison to classical state-of-the-art algorithms with comparable performance KPIs

# 5. Conclusion

The above experiments clearly state that quantum-inspired based generative model training of the RBM is more efficient than its classical counterpart of MCMC sampling in terms of training accuracy and out-of-sample performance. We can conclude that the utilization of quantum-inspired sampling techniques can bring in better behavior learning and understanding of the underlying patterns to reduce false positives and build a more robust anomaly detection algorithm. Quantum annealer-based sampling is efficient in terms of better sample generation and reduced sampling time. This capability could further be utilized in exploring more complex density functions estimation and building better approximation models for behavior modeling in different application domains including IoT sensor behavior modeling, system maintenance scheduling, retail customer behavior modeling, etc.

# Authors

## Ashutosh Vyas

*Senior Data Science Manager, Mphasis NEXT Labs*

Ashutosh has 6+ years of experience in the data science domain. He has worked on multiple projects of pattern recognition, time series forecasting, regression modelling, NLP, classification and optimization in Life science, Finance, FMCG and Media domains. He completed his MTech in 2015 from IIIT-B.

He has expertise in Bayesian methods of machine learning and had been working in quantum ML and quantum optimization for the past 2 years and has developed multiple algorithms in image classification, anomaly detection domain using quantum systems that leverage quantum gates and quantum annealer to process information and learn the patterns. At Mphasis, he works as a senior data science manager with an ethos of developing customer-centric and robust solutions.

## Lauren Roberts

*Data Scientist, Mphasis Datalytyx*

Lauren holds a PhD in statistics, specializing in Bayesian time series modeling for healthcare applications. She has worked as a data scientist for 2 years on projects involving time series, anomaly detection, and data visualization, as well as gaining experience in data engineering and CI/CD.

## Saurabh Gupta

*Senior Data Scientist, Mphasis NEXT Labs*

Saurabh Gupta holds a masters' degree from IIT Kanpur with majors in Analytics and Machine Learning. Currently, he is working as a Senior Data Scientist. He has worked on several AI/ML projects with expertise in statistical analysis, image processing, deep learning, and quantum machine learning.

## Rohit Patel

*AVP – Data Science and Quantum Computing, Mphasis NEXT Labs*

Rohit has 11+ years of experience in IT industry with 6+ years of experience in Data Science. He holds a BTech in ECE and PGDM degree in Finance. He has been co-leading the Quantum Computing practice at Mphasis NEXT Labs for 2+ years. He has developed solutions in AI and Quantum Computing in the areas of Logistics, Life Sciences, and Process Optimization, among others.

## About Mphasis

Mphasis' purpose is to be the *"Driver in Driverless Car"* for Global Enterprises by applying next-generation design, architecture and engineering services, to deliver scalable and sustainable software and technology solutions. Customer centricity is foundational to Mphasis, and is reflected in the Mphasis' Front2Back™ Transformation approach. Front2Back™ uses the exponential power of cloud and cognitive to provide hyper-personalized ($C = X2C^{2}_{TM} = 1$) digital experience to clients and their end customers. Mphasis' Service Transformation approach helps 'shrink the core' through the application of digital technologies across legacy environments within an enterprise, enabling businesses to stay ahead in a changing world. Mphasis' core reference architectures and tools, speed and innovation with domain expertise and specialization, combined with an integrated sustainability and purpose-led approach across its operations and solutions are key to building strong relationships with marquee clients. Click here to know more. (BSE: 526299; NSE: MPHASIS)

**For more information, contact: marketinginfo.m@mphasis.com**

**USA**
460 Park Avenue South
Suite #1101
New York, NY 10016, USA
Tel.: +1 212 686 6655

**UK**
1 Ropemaker Street, London
EC2Y 9HT, United Kingdom
T : +44 020 7153 1327

**INDIA**
Bagmane World Technology Center
Marathahalli Ring Road
Doddanakundhi Village
Mahadevapura
Bangalore 560 048, India
Tel.: +91 80 3352 5000

www.mphasis.com