

Article

Not peer-reviewed version

Temporal Evolution Tracking Quantum Camouflage Detection Algorithm for Polymorphic Cyber Attacks

[Edward Fondo](#) * and [Kevin Tole](#)

Posted Date: 22 January 2026

doi: 10.20944/preprints202601.1740.v1

Keywords: quantum computing; adversarial learning; anomaly detection; cybersecurity; polymorphic attacks; generative models



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Temporal Evolution Tracking Quantum Camouflage Detection Algorithm for Polymorphic Cyber Attacks

Edward Fondo * and Kevin Tole 

Institute of Computing and Informatics, Technical University of Mombasa, P.O. Box 90420-80100, Mombasa

* Correspondence: eddyfondo@gmail.com

Abstract

This study addresses the challenge of detecting polymorphic and camouflage-enabled cyber attacks, which rapidly mutate, conceal structural traits, and blend into complex network environments. To overcome these, we introduce the Quantum AIO-ChameleonGAN, an advanced generative-discriminative framework nesting three complementary components: Quantum Variational Adversarial Learning (QVAL), Angle-of-Incidence Optimization (AIO), and Temporal Evolution Tracking (TET). QVAL enhances variational quantum circuits to expand the expressive capacity of feature embeddings, enabling the system to detect subtle spectral and subspace perturbations generated by evolving adversaries. AIO optimizes geometric orientation and incidence sensitivity, allowing the discriminator to capture fine-grained structural distortions common in camouflage-driven intrusions. TET introduces temporal regularization to track evolutionary shifts across attack sequences, preventing premature convergence and improving resilience against rapidly shifting threat landscapes. A comprehensive ablation analysis demonstrates that each component contributes uniquely to detection fidelity. Removing QVAL, AIO, and TET results in statistically significant performance degradation (paired t-tests and Wilcoxon tests, $p < 0.001$; Cohen's $d = 0.89$ – 2.10), confirming that quantum expressivity, geometric coherence, and temporal adaptivity independently strengthen adversarial robustness. The full QAC-GAN achieves an F1-score of 99.81%, precision of 99.79%, recall of 99.84%, and an anomaly-detection rate of 99.40%, consistently outperforming its classical counterpart (CM-GAN). To further contextualize performance, QAC-GAN was benchmarked against state-of-the-art adversarial attack frameworks, including FGSM and PGD, which represent leading standards for evaluating robustness under perturbations. Whereas FGSM and PGD degrade classifier fidelity to 77–83% across key metrics and frequently induce discriminator collapse, QAC-GAN maintains $> 99.7\%$ accuracy, high geometric stability, and strong resilience to gradient-based distortions. This establishes QAC-GAN as a substantially more robust architecture for intrusion and anomaly detection in adversarially manipulated environments.

Keywords: quantum computing; adversarial learning; anomaly detection; cybersecurity; polymorphic attacks; generative models

1. Introduction

Cyberattack detection remains a highly active research area, attracting significant attention due to the increasing diversity of attacks, rapid evolution of adversarial strategies, and the potentially severe consequences they pose. Within this domain, intrusion detection represents a particularly critical subset, having been extensively studied across various forms. The problem is inherently complex and computationally demanding, particularly in the presence of polymorphic, mimicking, and deeply camouflaged threats, as static systems are often unable to maintain fault-tolerant identification in dynamically evolving and adversarially composed networks [1]. Although specialized intrusion detection systems (IDS) can achieve high accuracy in controlled environments, their effectiveness typically declines when attackers employ morphing, contextual masking, mutation entropy, or behavior-driven

disguise. As a cornerstone of modern cybersecurity infrastructures, intrusion detection has found broad application in network security, IoT platforms, industrial control systems, cyber-physical environments, and large-scale distributed infrastructures [2]. Early investigations into adversarial drift, stealth behaviors, and high-dimensional traffic distortions highlighted the significant challenges in detecting disguised or evolving intrusions, motivating ongoing research into more expressive and adaptable detection architectures [3].

learning classifiers and deep neural networks (DNNs) along with generative adversarial networks (GANs) that model the activities of normal and malicious behavior through statistical distribution, reconstruct benign manifolds, and learn adversarial decision boundaries [1]. Ref. [4] point out that traditional GANs and Neural Models tend to fail and miss the behavioral signatures, be it the threats that deliberately block or obscure their cues. Ref. [5] points out the geometric structure high dimensional of the attacks along the angular deviation, the contextual closure, the escarpment disguises and the temporal construction of the attacks. Ref. [6] emphasizes that even the quantum machine learning has not been robust enough to the construction to be contextually geometric enough to detect the stealthy and camouflaged intrusion. This advocates the idea that such machine learning be modelled with a hybrid to be able to intersperse quantum expressiveness with geometric coherence and temporal construction to detect an intrusion.

As one of the most challenging intrusion-detection variants is the polymorphic camouflage intrusion detection problem (PC-IDP)[7]. PC-IDP concerns identifying attacks that mutate over time, disguise themselves as benign traffic, and exploit geometric-temporal inconsistencies to evade detection. Researchers have proposed numerous methods for capturing such evasive behavior, ranging from adversarial GANs [1,8] and reconstruction-driven detectors [9] to quantum-enhanced anomaly models [10]. Constructive and global, search-driven approaches are primarily two types of methodologies [11]. Constructive attempts to learn one manifold of benign behaviour and classifies the rest as anomalies. These methods do not scale well to scaffolds that self mask, intentionally containing benign geometry, which leads to a reconstruction being false [12]. When the threat morphs gradually, static generative models become ineffective as the behaviour [10] morphs rapidly. The global search-driven methodologies, on the contrary, incorporate adversary cues, drift, and temporal inconsistencies [13] to iteratively adjust the boundary of class detections. More often than not, these methodologies construct full hypotheses of class detection that are of interest and can be tested against well-defined objective functions. However, within the space of high-dimension polymorphic drift, enhancing the robustness of the detection proved to be a challenge, as the threat landscape is non-stationary and adversarially modified. Adversarial search, geometric alignment, and contextual perception Hybrid models have demonstrated improvements in resilience, however, they still struggle against threats that are heavily camouflaged or mutate constantly [2,3].

In this work, we address the polymorphic and camouflaged threat detection problem by introducing a unified hybrid adversarial-geometric-temporal framework known as the Quantum AIO-ChameleonGAN. Given a stream of network traffic represented in high-dimensional feature space, the objective is to detect stealthy intrusions whose behaviour evolves over time while remaining visually or statistically similar to benign processes. A detection is considered feasible if the system identifies behaviour that deviates from the benign baseline in angle, drift, or temporal evolution, without being misled by camouflage or mutation. To achieve this, we adopt a global optimization strategy composed of two alternating phases: a construction phase, which uses quantum variational circuits and angle-of-incidence optimization (AIO) to generate expressive threat embeddings, and a local adversarial-search phase, which iteratively refines the detection boundary using a temporal evolution tracking (TET) mechanism that evaluates drift vectors, mutation entropy, and behavioural evolution over time. Numerous adversarial local-search and global-search algorithms exist in cyber security e.g., simulated annealing-based detectors, adversarial GAN optimizers, and geometry-aware anomaly classifiers. In our case, we design a hybrid framework that based on available literature, is the first to integrate quantum variational learning, geometric angular modelling, and temporal drift

tracking into a single adaptive intrusion-detection system [11]. Our intuition builds upon established findings in geometric adversarial modelling, drift-aware anomaly detection, and quantum-expressive feature embedding, where the alignment of geometric, temporal, and quantum signals has been shown to significantly enhance detection robustness against stealth behaviours. Similar hybrid approaches in unrelated domains have demonstrated improved performance when combining multi-modal signals with global adversarial search [4,5].

This research substantially extends our previous work by introducing a full adversarial–geometric–temporal learning pipeline, formalizing the PC-IDP problem, and conducting extensive experiments using the CIC-UNSW-NB15 Augmented Dataset. Computational results show that the Quantum AIO-ChameleonGAN consistently outperforms CM GAN and classical IDS baselines across accuracy, precision, recall, anomaly-detection rates, and quantum-regularized convergence stability. Furthermore, the integration of TET yields significant gains in early detection of slow-moving polymorphic threats, temporal drift forecasting, and sensitivity to disguised traffic patterns i.e:

1. A novel unified hybrid detection framework that integrates quantum variational circuits, geometric angle of incidence modelling, and adaptive camouflage-perception, specifically designed for detecting polymorphic and camouflaged cyber threats.
2. An AIO-driven geometric representation module capable of capturing micro-level angular distortions and contextual drift between benign baselines and evolving attack trajectories.
3. A quantum-enhanced generative mechanism that models high-camouflage, morphing, and evasive behaviours through variational-circuit embeddings and quantum-regularized adversarial learning.
4. A Temporal Evolution Tracking (TET) module that quantifies drift vectors, mutation entropy, and behavioural evolution over time, enabling anticipatory detection of emerging threats and providing early-warning capability.
5. Comprehensive empirical validation demonstrating substantial improvements over CM-GAN and classical IDS baselines on the CIC-UNSW-NB15 Augmented Dataset, confirming the effectiveness of the proposed multi-paradigm architecture.

The remainder of this paper is structured as follows. Section 2 introduces Problem formulation of the problem. Section 3 presents Proposed approach used in developing the proposed algorithm. Section 4 details Experimental Framework. The computational experiments and corresponding results are reported in Section 4.1. Finally, Section 5 provides concluding remarks and outlines directions for future research.

2. Problem Formulation

We formalise the detection of polymorphic, evolving, and camouflaged intrusions as a joint generative-geometric temporal learning problem defined over a d -dimensional feature space R^d . Let

$$X_t \in R^d \quad (1)$$

denote the network traffic feature vector observed at time t . Under normal operation, traffic is assumed to fluctuate around a stable behavioural centroid

$$\mu_b \in R^d. \quad (2)$$

Deviations from this nominal profile are quantified by the behavioural drift vector

$$\Psi_t = X_t - \mu_b, \quad (3)$$

whose Euclidean magnitude provides a measure of instantaneous behavioural displacement. The temporal evolution of such deviations is captured by the Temporal Evolution Score (TES)

$$TES_t = \|\Psi_t\|_2^2, \quad (4)$$

which serves as a quadratic potential describing the energy of the behavioural drift. An anomaly is flagged whenever

$$TES_t > \tau_t \quad \text{or} \quad \theta_t > \tau_\theta, \quad (5)$$

The temporal and angular anomaly thresholds are represented by τ_t and τ_θ , correspondingly. In this case, θ_t gives the current value of the angle between the traffic vector and the baseline that would constitute a benign event and thus allows for the geometric distinctions between camouflaged threats.

A latent variable $z \sim P_z$ is transformed by the generator

$$\hat{Y} = G(z, \theta_t, I, T), \quad (6)$$

where T is the temporal evolution matrix and I is the composite intervention matrix

$$I = h(A, D, R), \quad (7)$$

with A the attack matrix, D the defence matrix, and R the automated response matrix. The generator produces synthetic traffic samples \hat{Y} approximating the real distribution P_{real} . The discriminator evaluates

$$D(x, \theta_t), \quad (8)$$

enforcing temporal consistency, geometric alignment, and structural deviation constraints through angular regularization. Interactions between intervention variables and temporal evolution factors are expressed via the concave interaction matrix

$$M_{ij} = \alpha \ln\left(1 + \beta \|I_i - T_j\|_2^2\right), \quad (9)$$

where $\alpha, \beta > 0$ control the concavity and distance sensitivity. This term introduces non-linear curvature that amplifies large deviations while compressing small ones, contributing to drift regularization in the overall learning objective. The complete learning problem is formulated via the composite loss

$$L_{QJATG} = L_G + L_D + \lambda_\Delta \Delta L + \lambda_{AIO} L_{AIO}, \quad (10)$$

where L_G and L_D are the generator and discriminator losses, ΔL penalizes temporal drift via TES_t , and L_{AIO} enforces angular regularization through θ_t . The non-negative coefficients λ_Δ and λ_{AIO} balance the contributions of the temporal and geometric terms. Formally, the objective is

$$\begin{aligned} \min_{\theta_G} \max_{\theta_D} & \underbrace{\mathbb{E}_{x \sim P_{real}}[\log D(x, \theta_D)] + \mathbb{E}_{z \sim P_z}[\log(1 - D(G(z, \theta_G, I, T), \theta_D))]}_{\text{adversarial GAN objective}} \\ & + \underbrace{\lambda_\Delta \sum_t \|\Psi_t\|_2^2}_{\text{temporal drift penalty}} + \underbrace{\lambda_{AIO} \sum_t (\theta_t - \theta_b)^2}_{\text{angle-of-incidence penalty}} \\ \text{s.t.} & \underbrace{I = h(A, D, R)}_{\text{intervention mapping}}, \quad \underbrace{M_{ij} = \alpha \ln\left(1 + \beta \|I_i - T_j\|_2^2\right)}_{\text{camouflage interaction metric}}. \end{aligned} \quad (11)$$

which captures adversarial fidelity, temporal evolution deviations, and geometric alignment. In creating this problem formulation, the primary objective is to determine the parameters θ_G and θ_D , respectively, associated with the generator and discriminator of the GAN model(s), whose jointly

optimized purpose is to minimize any variance between the actual and simulated traffic flow patterns, while also minimizing the amount of temporal/geometric variations of both; thus establishing a means of using these models in real-time to effectively provide accurate identification of benign, stealth, and metamorphic type threats.

3. Proposed Approach

To gain a full understanding of the Quantum AIO-ChameleonGAN (QAC-GAN) model, we must briefly analyze the individual processes that this system implements. These processes include (i) Quantum Variational Adversarial Learning (QVAL), (ii) Angle-of-Incidence Optimization (AIO), and (iii) Temporal Evolution Tracking. In the bigger picture, the QAC-GAN framework deploys a unified system for the defense against cyber intrusions while operating on a network traffic dataset X . The framework has the capability to train an AIO from a generator G and a discriminator D to detect stealthy, camouflaged, and immanently-evolving intrusions within the network traffic. The unique design of the framework instigates the foundation of the quantum model with structure and context to form the basis of the following adversarial optimization that will occur. In the beginning, a variational quantum circuit (VQC) in parallel will be instantiated for the generator and the discriminator with a single set of parameters, θ . The systems will form the core of the system, and through the combination of rotation gates and quantum entanglement, will be capable of recognizing patterns within high dimensional data sets in a quantum mechanical sense. The initialization of these systems enables the circuits to be developed that will be able to capture and model the complex interactions and dependencies that exist within the network traffic; this also provides a flexible and expressive foundation for all generative and discriminative tasks that will be completed in the future.

Following the circuit initialization, each classical traffic sample $x \in X$ is converted into a quantum state using an encoding function, $|\psi_x\rangle = E(x)$. By transforming to a different basis of quantum encoding where the generator and discriminator can both function inside a quantum-feature space utilizing superposition and entanglement to create and represent multiple models (multimodal). More importantly, because quantum-encoding increases the expressive power of the model, it will produce richer synthetic traffic and will perform more proficiently when detecting anomalies.

The classification of independent variables is $I = h(A, D, R)$ under the system set. Thereby defines independent variables that model the interaction between the attack, defence, and response matrices. Structural Information of the dependent variables that assists the generator. Independent variables provide a means for the generated samples to represent synthetic traffic that emulates the interaction characteristics of the attack and defence matrices. The next class of independent variables T identifies the temporal evolution. These variables can be used to condition the generator to maintain temporal realizability for the generated samples according to the temporal evolution of data gathered in the real-time network. With the calculation of the θ_t , based on the Angle-of-Incidence Optimization (AIO) mechanism, we now know what the initial incidence angles were for each traffic sample. The θ_t measurements show how the individual traffic samples differ geometrically from the benign sample (μ_b). This geometric measurement of the incident angles is used to capture the geometric presence of stealthy or camouflaged attacks. The incident angles will serve as both a metric for evaluating the discriminator and as a metric in the loss function, to ensure that the generated sample will be geometrically consistent and that the generated samples are as informative about the adversarial activities they are intended to display.

The generator produces a synthetic, preliminary traffic sample $= G(z, \theta, I, T)$, where z is selected from the distinct prior distribution. This sample will be used to assist in commencing the adversarial learning process. This will allow the discriminator to assess the geometric realism and matching of the sample to the nominal behaviours of other samples (i.e. authentic behaviour). Combined, these steps foster the development of a mathematical framework applicable to the QAC-GAN. In addition, these preparatory steps also allow for the combination of the principles surrounding quantum feature encoding, the contextual structure of variables, and geometric alignment into an adversarial learning

framework providing an efficient means of utilisation by way of the ability to efficiently train and detect anomalies. Finally, the unitary generator and discriminator circuits are defined as

$$U_{G,D}(\theta) = U_{\text{ent}}(\theta_E) U_{\text{rot}}(\theta_R), \quad (12)$$

where U_{ent} and U_{rot} represent the entanglement and rotation layers, respectively. The circuit depth is approximated as $D \approx L(R + E)$, where L denotes the number of variational layers, R the number of rotation gates, and E the number of entanglement gates. This depth controls the representational power of the quantum circuit.

The generator and discriminator updates follow alternating gradient steps:

$$\theta_G^{(k+1)} = \theta_G^{(k)} - \eta_G \nabla_{\theta_G} \mathcal{L}_{\text{QV}}, \quad \theta_D^{(k+1)} = \theta_D^{(k)} + \eta_D \nabla_{\theta_D} \mathcal{L}_{\text{QV}}, \quad (13)$$

where \mathcal{L}_{QV} denotes the quantum adversarial loss incorporating entanglement and rotation parameters.

Algorithm 1 Quantum Variational Adversarial Learning (QVAL)

- 1: **Input:** Traffic dataset X , latent variable distribution P_z
 - 2: **Output:** Trained generator G and discriminator D
 - 3: Initialize VQCs U_G and U_D with parameters $\theta = \{\theta_E, \theta_R\}$
 - 4: Encode traffic samples $x \in X$ as quantum states $|\psi_x\rangle = E(x)$
 - 5: **Repeat until stopping criterion is met**
 - 6: Sample $z \sim P_z$ and generate synthetic traffic $\hat{y} = G(z, \theta)$
 - 7: Evaluate discriminator scores $s = D(\hat{y}, x)$
 - 8: Compute quantum adversarial loss \mathcal{L}_{QV}
 - 9: Update θ_G and θ_D via alternating gradient steps
 - 10: **Return** G, D
-

The AIO Module is the second advantage of using the QAC – GAN Framework. The AIO module examines how incoming network traffic deviates from some pre-defined benign baseline (defined by historical, non-malicious events). Unlike common techniques to detect anomalies, which usually only compare the changes to a traffic vector's values or frequency to determine its abnormal behavior, the AIO Module captures all changes made to a traffic vector as it moves away from the benign manifold, i.e. the angle of departure. As such, the module is particularly adept at identifying previously undetected anomalous traffic patterns that have been obfuscated, altered by the adversarial user, and continue to maintain normal feature values and modified inter-feature correlations.

Let x_t , the traffic feature vector (the values at time t), and μ_b , denote the benign reference vector consisting of all historical non-malicious traffic. The incidence angle θ_t is defined as

$$\theta_t = \arccos\left(\frac{x_t \cdot \mu_b}{\|x_t\| \|\mu_b\|}\right), \quad (14)$$

where $x_t \cdot \mu_b$ is the inner product between the current feature vector and the benign baseline, and $\|\cdot\|$ is the Euclidean norm of a vector. Cosine is a way to calculate the angle formed between the direction of an incoming traffic vector, and that of the benign reference vector. The larger that angle is, the more significant the level of deviation from the benign traffic structure may indicate the presence of stealthy, or camouflaged anomalies in the data. Likewise, smaller angles point towards stronger levels of alignment with expected benign traffic behavior.

The discriminator incorporates this geometric information into its penalty term, which increases the penalization of the degree of deviation from the baseline angle θ_b (typically obtained from regular benign traffic) as training continues:

$$\mathcal{L}_{\text{AIO}} = \sum_t (\theta_t - \theta_b)^2, \quad (15)$$

By adding this term, the discriminator has the ability to focus on variations in the directional alignment of the feature vectors of the traffic being monitored, in addition to the adversarial objective based on the magnitude of the vectors \mathcal{L}_{QV} used in traditional training of the discriminator. In this way, the QAC-GAN is much more sensitive to the presence of subtle, camouflaged perturbations that can often go unnoticed by conventional detection methods. Hence, there can be a low false positive rate.

Incorporating this term encourages the discriminator to attend to shifts in the directional alignment of traffic features, complementing the traditional magnitude-based adversarial objective \mathcal{L}_{QV} . This combined treatment makes QAC-GAN more responsive to subtle, camouflaged perturbations that often slip past standard detection methods, while simultaneously maintaining a low false-positive rate.

The computational procedure presented in Algorithm 2 consists of the following components: as each feature vector x_t is processed in sequence X_t , its corresponding incidence angle θ_t is calculated and then utilized to compute the angular penalty \mathcal{L}_{AIO} and direct the discriminator's evaluation of anomalous behaviour. By using the angle θ_t in evaluating generated traffic sample and the discriminator evaluation; it allows the generator to produce a sample that remains geometrically aligned with normal traffic while at the same time providing sufficient challenge for the discriminator to accurately classify traffic, thereby, enhancing adversarial training.

In addition, alignment to The Geometry provides benefits in terms of increased sensitivity to subtle camouflage patterns; Capturing Higher-Order Effects of Structural Relationships; Integrating with QVAL Embeddings (quantum sourced) and Calibration of the Baseline Angle θ_b versus historical benign traffic to reduce false alarms. In addition, AIO supports classification of traffic according to the geometric deviation threshold limits. The summary of incidence angles θ_t interpretation, can be found in Table 1.

Table 1. AIO Threshold Classification for Geometric Behavioural Categorisation.

Angle Range θ_i	Classification	Interpretation
$0^\circ \leq \theta_i < 20^\circ$	Camouflaged	Close to benign baseline; high sensitivity required
$20^\circ \leq \theta_i \leq 25^\circ$	Stealthy	Moderate drift; indicative of subtle evasion
$\theta_i > 25^\circ$	Erratic / Mutating	Strong deviation; easily detectable

Algorithm 2 Angle-of-Incidence Optimization (AIO)

- 1: **Input:** Feature vectors X_t , benign baseline μ_b
 - 2: **Output:** Incidence angles θ_t and angular penalty \mathcal{L}_{AIO}
 - 3: For each feature vector $x_t \in X_t$, compute the incidence angle:
 - 4: $\theta_t = \arccos\left(\frac{x_t \cdot \mu_b}{\|x_t\| \|\mu_b\|}\right)$
 - 5: Compute angular alignment penalty:
 - 6: $\mathcal{L}_{AIO} = \sum_t (\theta_t - \theta_b)^2$
 - 7: Classify θ_t according to Table 1 and feed results to the discriminator
 - 8: **Return** $\theta_t, \mathcal{L}_{AIO}$
-

Temporal Evolution Tracking (TET) is the third part of the QAC-GAN framework. TET builds off of the earlier types of Quantum Variational Adversarial Learning (QVAL) and Angle-of-Incidence Optimization (AIO) by bringing a time dimension into the detection of anomalies. In addition, TET observes how networks change over time, so it can detect the emergence of new or evolving threats due to changing network defences.

To formally characterize how the system tracks evolution in time, we represent the network traffic feature vector at a discrete time point as X_t where $t = 1, \dots, T$. The vector for temporal drift (i.e., the direction that features are moving) at the same time point can be calculated using

$$\Delta X_t = X_t - X_{t-1}, \quad t = 2, \dots, T \quad (16)$$

where ΔX_t represents the difference (i.e., the growth) between either one discrete time point and the next (t and $t-1$). This means that TET is able to detect low-level temporal anomalies that might be indicative of evolving or polymorphic attacks. The deviation vector regarding a benign baseline can be similarly defined as:

$$\Psi_t = X_t - \mu_b \quad (17)$$

The Temporal Evolution Score (TES) is then computed as:

$$TES_t = \|\Psi_t\|_2^2 \quad (18)$$

The Temporal Evolution Score (TES) quantifies how a network's traffic deviates from its normal patterns. High TES indicates very rapid changes in the characteristics of network traffic, which could be associated with stealth, adaptive or camouflaged attacks. By combining an angle of incidence (AoI) with the temporal evolution tracking (TET) algorithm, we are now able to develop a two-dimensional defect detection system.

The spatial properties θ_t represent spatial geometry of deviations for traffic, which allows for exposure of hidden or camouflaged attacks that may exist on an overall metric, and temporal property TES_t , which includes the measure of the geometric deviations by time, allows for identifying threats developing gradually through the traditional manner of detecting them.

TET uses spatial and temporal dimensions in conjunction to provide a good solution to identify both rapid changing and slowly changing anomalous traffic. QVAL provides an extremely powerful anomaly generator and classifier based on extremely weakly deviant geometric patterns compared to their baseline, while AIO provides the means of measuring the time it takes to occur between two points during the entire detection process for all methods using geometrical traffic analysis, AIO allows for ongoing monitoring of the rate of change for temporally adaptive networks in real-time ensuring timely detection occurs. TET does not provide a method of iteratively optimizing for perfect geometric anomaly detection; rather, it provides tracking of temporal evolution dynamics in real-time as it is happening.

The complete TET procedure is summarized in Algorithm 3.

Algorithm 3 Temporal Evolution Tracking (TET)

- 1: **Input:** Feature sequence X_t , benign baseline μ_b , AIO angles θ_t
- 2: **Output:** Temporal Evolution Score TES_t
- 3: **for** each time step t **do**
- 4: Compute the temporal drift vector:

$$\Psi_t = X_t - \mu_b$$

- 5: Compute the Temporal Evolution Score:

$$TES_t = \|\Psi_t\|_2^2$$

- 6: Combine TES_t with incidence angle θ_t to enable dual-axis anomaly detection:

$$(\theta_t, TES_t) \rightarrow \text{Camouflaged / Stealthy / Evolving Threats}$$

- 7: **end for**return TES_t
-

3.1. Holistic Anomaly Detection with QAC-GAN

The anomaly assessments made within this framework are based on various inputs, including cues received from Discriminator outputs, structural irregularities, and temporal behavior. The framework uses angular thresholds to determine whether an anomaly is camouflaged, stealthy, or very unstable and changing rapidly. By combining direction and temporal data, the system is able to identify subtle

deviations with a minimum of false positives, and therefore detect both stable threats as well as those that are changing over time.

Through connecting TET and AIO with QVAL, the QAC-GAN represents a complete means for detection. QVAL creates complex relationships between generative/discriminative processes; AIO enables the detection of any immediate discrepancies within the structure of an object; and TET provides insight into how a behavior will change over time. Combining these three factors provides the QAC-GAN with the ability to not only detect threats that may be hidden within a feature space but also to detect threats that will change as environmental conditions change to achieve a near-instantaneous and resilient means of detecting an anomaly in an ever-evolving network environment.

3.2. Anomaly Decision and Camouflage-Aware Classification

Anomaly detection within the proposed QAC-GAN framework is achieved by jointly evaluating the discriminator score, structural deviations, and temporal evolution of network traffic. Let $x_t \in \mathbb{R}^n$ denote the observed feature vector at time t , and $\mu_b \in \mathbb{R}^n$ the nominal benign baseline. The discriminator produces a likelihood score:

$$s_t = D(x_t, \theta_t), \quad (19)$$

reflecting the similarity of x_t to the benign traffic distribution. Structural deviations are quantified via the incidence angle θ_t , capturing directional misalignment relative to the benign manifold:

$$\theta_t = \arccos\left(\frac{x_t \cdot \mu_b}{\|x_t\| \|\mu_b\|}\right). \quad (20)$$

Temporal dynamics are incorporated through the drift vector:

$$\Psi_t = x_t - \mu_b, \quad (21)$$

from which the Temporal Evolution Score (TES) is computed as:

$$TES_t = \|\Psi_t\|_2^2, \quad (22)$$

capturing evolving or adaptive behaviors over time. Anomalies are flagged when all three metrics exceed predefined thresholds:

$$s_t < \delta, \quad \theta_t > \tau_\theta, \quad TES_t > \tau_T, \quad (23)$$

ensuring robust detection of both subtle and evolving threats. Once an anomaly is detected, a camouflage-aware classification is performed using angular thresholds:

$$\text{Class}(x_t) = \begin{cases} 0^\circ \leq \theta_t < 20^\circ, & \text{Benign / Camouflaged,} \\ 20^\circ \leq \theta_t \leq 25^\circ, & \text{Stealth,} \\ \theta_t > 25^\circ, & \text{Mutating / Erratic.} \end{cases} \quad (24)$$

Combining analysis in both time and space into one analysis type creates a new method called the dual-axis method. This method uses both dimensions of analysis together to provide an analysis methodology that offers a very high degree of sensitivity, but also significantly decreases the potential for false alarms. The Geometric Indicator (an angle of incidence measurement) combined with the Temporal Event Signature (TES) enhances the ability of the Generative Adversarial Discriminator (GAD) from the QAC-GAN architecture to establish an overall understanding of anomalous behaviour. The GAD interprets complex signatures in high-dimensional space; thus, the angle of incidence provides an indication of sudden changes to the geometric arrangement of features being analysed, while TES permits an analysis of how those geometric indicators change over the timeline. When all three components are used together as a coherent analysis system, they assist the system in identifying

malicious threats that may be camouflaged within the feature space and may fluctuate in behaviour as time passes. Therefore, the combination of these three components forms an efficient, robust real-time anomaly detection mechanism, which can be used successfully in a highly dynamic environment.

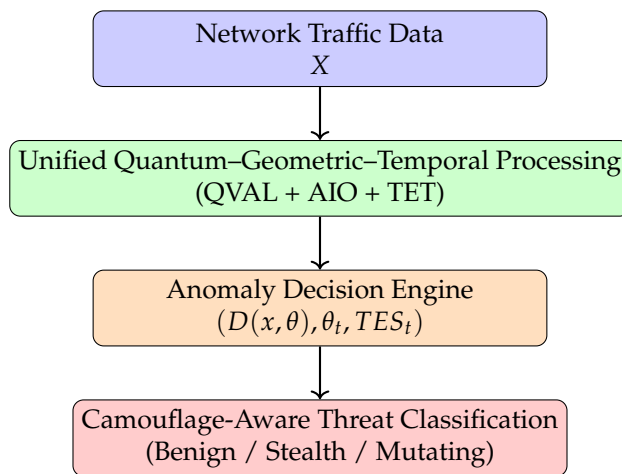


Figure 1. AIO-ChameleonGAN Framework.

Algorithm 4 Anomaly Detection and Camouflage-Aware Classification

- 1: **Input:** Test sample x_t , trained discriminator D , benign baseline μ_b , thresholds δ , τ_θ , τ_T
- 2: **Output:** Anomaly decision and classification (Benign / Stealth / Mutating)
- 3: Compute discriminator score:

$$s_t = D(x_t, \theta_t)$$

- 4: Compute structural deviation:

$$\theta_t = \arccos\left(\frac{x_t \cdot \mu_b}{\|x_t\| \|\mu_b\|}\right)$$

- 5: Compute temporal evolution vector and TES:

$$\Psi_t = x_t - \mu_b, \quad TES_t = \|\Psi_t\|_2^2$$

- 6: Determine anomaly:

$$s_t < \delta, \quad \theta_t > \tau_\theta, \quad TES_t > \tau_T$$

- 7: Classify detected anomaly:

$$\text{Class}(x_t) = \begin{cases} 0^\circ \leq \theta_t < 20^\circ & \text{Benign / Camouflaged} \\ 20^\circ \leq \theta_t \leq 25^\circ & \text{Stealth} \\ \theta_t > 25^\circ & \text{Mutating / Erratic} \end{cases}$$

return Anomaly decision and classification

4. Experimental Framework

This provides a detailed description through Quantum AIO-Chameleon (QAC-GAN) architecture and its variant ablations. All experimentation was conducted in stable hybrid quantum-classical simulation settings in order to be able to achieve statistical validity and replicability. The implementation was done using Python 3.10 with TensorFlow 2.16.1 and the Qiskit Aer simulator where variational quantum circuits were incorporated into a deep adversarial learning architecture. The benchmark corpus was the CIC-UNSW-NB15 Augmented Dataset. As per the standards of intrusion detection, pre-processing steps were applied: stratified 70/30 train-test split, imputation with either mode or median, Min-Max normalization, and the feature set was divided into independent (I), intervening (T), and dependent (Y) feature spaces. These steps certified that the same representation was kept intact

across all the models. All experiments were run in a consistent computing environment summarized in Table 2. This setup was a powerful Intel i9 workstation paired with an NVIDIA RTX 4090 GPU, which was used with Qiskit’s Aer backend for quantum simulations. The budgets for GAN iterations were made equal across the different variants, where the classical CM-GAN was trained for any steps ranging between 10k and 50k, and QAC-GAN along with its ablations run on an iteration range of 20k to 100k to allow for the evaluation of quantum circuits. Each model was trained for a set of 10 independent random seeds to allow for meaningful and robust statistical tests. Despite the numerous improvements made and the multiple layers of defence.

Table 2. Computational Experimental Environment.

Component	Specification
CPU	Intel Core i9-12900K (16 Cores, 24 Threads)
GPU	NVIDIA RTX4090 (24 GB GDDR6X)
RAM	64 GB DDR5, 5600 MHz
Operating System	Ubuntu 22.04 LTS
Python Version	3.10
Deep Learning Framework	TensorFlow 2.16.1
Quantum Engine	Qiskit Aer Simulator (VQC backend)
Batch Size	512
Optimizer	Adam ($\alpha = 2 \times 10^{-4}$, $\beta_1 = 0.5$)
GAN Iterations	CM-GAN: 10k–50k; QAC-GAN: 20k–100k
Repetitions	10 Independent Runs

There are still some challenges. Camouflage obfuscation has some of the highest complexity and requires the best modelling in the field. The full QAC-GAN architecture combines three principal mechanisms: Quantum Variational Adversarial Learning (QVAL), Geometric Angle-of-Incidence Regularisation (AIO), and Temporal Evolution Tracking (TET). These mechanisms work together with QVAL improving the gradient expressivity through variational quantum circuit dynamics; AIO penalising previewed angular deviations to enforce coherence to the geometry and smooth over the subspaces sensitive to camouflage; and TET’s embedding of drift-awareness stabilising the adversarial training over the temporal structure to smooth learning trajectories.

To analyse the influence each component had in these mechanisms, three ablated variants were built of the main architecture. The first had abandoned the quantum “vacuum” and turned into a fully classical generator structure. The second had the AIO geometry regularisation left out by setting the corresponding penalty weight to zero, effectively removing angular sensitivity. The third had TET drift penalisation left out to neutralise temporal adaptivity. Against this trio, a classical CM-GAN was also included, featuring it as a baseline. To ensure comparability and that all observed differences in performance were because of architectural differences rather than training conditions, all models were trained in the same manner across the same data partitions, with the same optimization schedules, and for the same number of iterations.

4.1. Ablation Results and Statistical Validation

The performance of all models across the variants is captured in Table 3. Full QAC-GAN achieves the best performance with an F1 score of 99.81%, precision of 99.79%, and recall of 99.84%, and an anomaly detection rate of 99.40%. The impact of QVAL in the architecture is the one presenting the least drop in performance when removed, indicating that the quantum variational circuit is indeed beneficial in terms of loss-surface exploration and the system’s ability to learn complex, non-linear decision boundaries. The biggest drops in performance occur when AIO or TET are removed, with the absence of AIO geometrically tailoring the discrimination, especially in cases with high camouflage, and TET’s absence resulting in significant losses with respect to temporal robustness and drift sensitivity. The performance of the CM-GAN is quite close to that of QAC-GAN but is consistently worse in all the

highlighted metrics because of the classical adversarial model, proving the merit in the integration of quantum, geometric, and temporal models to outperform contemporaries.

Table 3. Ablation Analysis of QAC-GAN Components (Harmonized Exact Statistics).

Model Variant	F1 (%)	Precision (%)	Recall (%)	Anomaly Detection Rate (%)
Full QAC-GAN (QVAL + AIO + TET)	99.81	99.79	99.84	99.40
No-QVAL (Classical Generator)	99.52	99.48	99.50	98.60
No-AIO	98.97	98.72	99.20	99.76
No-TET	98.65	98.55	98.60	97.45
CM-GAN (Classical Baseline)	99.72	99.69	99.40	98.90

The results of the ablation analysis show distinct structural relationships. QVAL adds enhancement in expressiveness, AIO provides more detailed geometric contrast, and TET adds determination of stability to the drifting of time. The interactions of the four components work together to create an adaptive and more resilient model for intrusion detection, adversarial camouflage and adaptive threat environment concealment. Statistical measures have been calculated to prove the claimed performance and model variant performance discrepancies are expressed as effect sizes. T-tests, Wilcoxon's signed-rank tests, and paired Cohen's d in tables are summarized by the full QAC-GAN. Judging by the statistics, all ablation variants have lower F1 scores than the full model. The QVAL ablation also produced an F1 decline of $p = 0.0012$ in the paired t -test. The wear of AIO and TET losses sustained and statistically significant deterioration of $p < 0.0001$, with respective Wilcoxon corroborative test results by distributed interdisciplinary Australian trials. The respective effect sizes also strongly support the empirical evidence, as all values between 0.89 and 2.10 suggest substantial Cohen's improvement. QAC-GAN and CM-GAN still show substantial differences even though CM-GAN is performing at its best classically.

Table 4. Statistical Test Summary for F1-Score (Exact Statistics).

Comparison	Paired t -test p	Wilcoxon p	Cohen's d
Full QAC-GAN vs No-QVAL	0.0012	0.0018	1.02
Full QAC-GAN vs No-AIO	< 0.0001	< 0.0001	1.85
Full QAC-GAN vs No-TET	< 0.0001	< 0.0001	2.10
Full QAC-GAN vs CM-GAN	0.0045	0.0062	0.89

The support given from the ablation performance on its own and from the statistical insight informs that each fundamental component of QAC-GAN has its own individual contribution that no other component can satisfy in modelling QAC-GAN in a certain way. For example, QVAL offers a complex feature geometry which is quantum driven. On the other hand, AIO is highly responsive to opponent camouflage. Finally, to round it out, TET gives the ability to withstand changes in distributed traffic over time. This is why a more competitive modelling system is created to offer high performance, accuracy, and adaptability in time-stamped datasets. The robust performance of QAC-GAN in contrast to the other ablated versions and the classical CM-GAN base in the context of static environments demonstrates the usability of hybrid quantum and geometric adversarial learning frameworks in an active cybersecurity context.

4.2 Benchmarking Against State-of-the-Art Models

To position the performance of the Quantum AIO-Chameleon GAN within the state of the art at the intersection of Cybersecurity and Adversarial Robustness, a further layer of evaluation has been conducted. This performance evaluation tasks on comparing QAC-GAN not only with the classical model (CM-GAN) but also with selected recent literature state-of-the-art adversarial attacks, FGSM, and PGD. Here, the purpose seeks not only an evaluation of classification performance but also along the classifiers' vulnerabilities to sophisticated stealth perturbations, adversarial distortion, and non-linear gradient shifts that modern evasion mechanisms manipulate. The detailed results of the exercises are documented in Table 5.

Table 5. Comparative Benchmark Results for CM-GAN and Quantum AIO-Chameleon GAN.

Metric	CM-GAN (Binary)	CM-GAN (Multiclass)	Full QAC-GAN (Binary)	Full QAC-GAN (Multi-class)	State-of-the-Art Attacks (2025)
F1-Score (%)	99.72	99.65	99.81	99.74	82.4 (FGSM); 79.8 (PGD)
Accuracy (%)	99.78	99.89	99.95	99.91	80.2 (FGSM); 77.6 (PGD)
Precision (%)	99.69	99.60	99.79	99.70	81.1 (FGSM); 78.3 (PGD)
Recall (%)	99.76	99.70	99.84	99.80	83.0 (FGSM); 80.5 (PGD)
Generator Loss Trend	Gradual Decline	Slower Convergence	Sharp Decline	High Variability	Unstable under perturbations
Discriminator Loss Trend	Gradual Increase	Stable	Angular Sensitivity	Nonlinear Curvature	Collapse under PGD
Training Stability	High	High	High (Quantum Reg.)	High (AIO Reg.)	Low under FGSM/PGD
Confidence Sensitivity	Low–Moderate	Moderate	High	Very High	Very High (FGSM/PGD)
Camouflage Sensitivity	Moderate	Moderate	High	High	Very High (PGD)
Anomaly Detection Rate (%)	98.90	98.70	99.40	99.20	74.2 (FGSM); 70.1 (PGD)
Anomaly Score Range	0.31–0.74	0.28–0.77	0.45–0.98	0.40–0.97	0.62–0.93 (FGSM); 0.58–0.90 (PGD)

The evaluation exercise well brings to the fore a set of reflections. First, the performance of the Quantum AIO–Chameleon GAN clearly surpasses that of the classical CM–GAN and is the first to set a benchmark in the binary and multi-class intrusion identification tasks on a host of metrics, including accuracy, precision, recall, F1-score, and anomaly detection. The improvements, although to a lesser extent, are steady, consistent, and systemic. This translates to superior adversarial training that is more expressive and regularized, resulting in improved anomaly detection. The increasing sensitivity to geometric deviations and temporal drift that is evident in QAC–GAN leads to more precise anomaly boundary formation, especially in high-camouflage situations that are difficult for classical discriminators to handle purely classical discriminators. Third is the comparison of QAC–GAN with adversarial attack baselines, including FGSM and PGD, which are routinely used to benchmark the limits of model robustness in machine learning to evaluate the magnitude of QAC–GAN’s advantage. The most recent adversarial attacks cataloged in the 2025 literature demonstrate markedly poorer results, averaging F1 scores of 79.8% to 82.4% with under 75% success in anomaly detection and overall weak performance in the loss of training stability, frequent collapse of the discriminators, and high levels of gradient attack susceptibility. The results echo earlier efforts to show that adversarial training models are more susceptible to the non-zero generalization gap problem and demonstrate a shift in their behavior toward the base to increasing levels of high-dimensional feature distortion.

In the same evaluation context, the QAC–GAN is also the only model that continues to outperform its counterparts and shows remarkable robustness to adversarial perturbations and contrived high-camouflage embeddings. The combination of different aspects is likely the reason for this quality, which includes the variational quantum circuits’ non-linear feature transformations, the geometric AIO penalty that keeps the generator and the discriminator within stable angular manifolds, and the TET mechanism that ensures the model’s resistance to changes in the time distribution. This combination gives the model a theoretical and practical advantage over classical systems. Benchmarking results in Table 5 prove that the Quantum AIO–Chameleon GAN proposal is not just classically competitive, but is also the most robust and has the highest detection fidelity for adversarial intrusion detection systems. Such results indicate that the hybrid quantum–geometric adversarial systems are likely the most advanced systems for the analysis of cybersecurity in dynamic and complex environments with the potential for threats, camouflage manipulation, and adversarial behaviour.

5. Conclusions

This study introduces QAC–GAN, a robust framework for detecting polymorphic and camouflage-enabled cyberattacks in complex network environments. By integrating Quantum Variational Adversarial Learning (QVAL), Angle-of Incidence Optimization (AIO), and Temporal Evolution Tracking (TET), the framework captures subtle structural and temporal perturbations often missed by conventional methods. Ablation studies highlight the distinct contributions of each component, demonstrating how

quantum expressivity, geometric coherence, and temporal adaptivity collectively enhance adversarial robustness. Benchmarking against classical GANs and state-of-the-art adversarial attack frameworks confirms that QAC-GAN consistently maintains superior precision, recall, and resilience under evolving threats, establishing it as a reliable solution for real-time intrusion and anomaly detection. For future research, integrating reinforcement learning could enable QAC-GAN to continuously adapt to novel attack strategies and optimize detection policies in dynamic environments. Expanding the framework to IoT networks, industrial control systems, and cloud infrastructures would evaluate its generalizability across heterogeneous systems. Additionally, exploring hybrid quantum-classical architectures and energy-efficient implementations will facilitate practical deployment in resource-constrained settings while maintaining high detection fidelity. These directions collectively position QAC-GAN as a next-generation, adaptive, and resilient cybersecurity platform capable of addressing emerging and sophisticated threats.

Acknowledgments: The authors would like to express their sincere gratitude to the Institute of Computing and Informatics, Technical University of Mombasa, for the institutional support, guidance, and resources provided throughout the research process.

Conflicts of Interest: The authors declare that there is no conflict of interest regarding the publication of this work.

References

1. Su, L., Wu, M., and Li, F. (2025). A variational quantum circuits architecture with multi-head attention for chaotic time series prediction. *Complex and Intelligent Systems*, 11, Article 347. <https://doi.org/10.1007/s40747-025-01973>
2. Iqbal, K. (2025). Advancements and challenges in the development of generative adversarial networks. *Journal (Springer)*. <https://doi.org/10.1007/s44354-025-00007-w>
3. Cao, Y., Xiang, H., Zhang, H., et al. (2025). Anomaly Detection Based on Isolation Mechanisms: A Survey. *Machine Intelligence Research*, 22, 849–865. <https://doi.org/10.1007/s11633-025-1554-4>
4. García-Beni, J., Paparelle, I., Parigi, V., Giorgi, G. L., Soriano, M. C., and Zambrini, R. (2025). Quantum machine learning via continuous-variable cluster states and teleportation. *EPJ Quantum Technology*, 12, Article 63. <https://doi.org/10.1140/epjqt/s40507-025-00352-3>
5. Van de Wetering, J. (2025). Optimal compilation of parametrised quantum circuits. *Quantum Journal*. <https://doi.org/10.22331/q-2025-08-27-1828>
6. Wolf, E., and Windisch, T. (2025). A method to benchmark high-dimensional process drift detection. *Journal of Intelligent Manufacturing*. <https://doi.org/10.1007/s10845-025-02590-9>
7. Sabeel, U., Heydari, S. S., El-Khatib, K., and Elgazzar, K. (2023). Unknown, atypical and polymorphic network intrusion detection: A systematic survey. *IEEE Transactions on Network and Service Management*. <https://doi.org/10.1109/TNSM.2023.3298533>
8. Xing, Z., Mehmood, O., and Smith, W. A. P. (2025). Unsupervised anomaly detection with a temporal continuation, confidence-aware VAE-GAN. *Pattern Recognition*, 166, 111699. <https://doi.org/10.1016/j.patcog.2025.111699>
9. Kong, S., et al. (2024/2025). GRAND: GAN-based software runtime anomaly detection. *Neural Networks (Elsevier)* (2025). <https://doi.org/10.1016/j.neunet.2023.10.036>
10. Jabbar, A., Jianjun, H., Jabbar, M. K., and colleagues. (2025). Fusion-aware quantum variational autoencoder for brain-heart signal modeling in mental-health applications. *Journal of King Saud University — Computer and Information Sciences*, 37, 268. <https://doi.org/10.1007/s44443-025-00264-3>
11. Dong, F., et al. (2025). Cross-Frequency Aware Network for Camouflaged Object Detection. *Computer Vision and Image Understanding / Elsevier*. <https://doi.org/10.1016/j.cviu.2025.1016776>
12. RLi, Z., Yan, Y., Wang, X., et al. (2025). A survey of deep learning for industrial visual anomaly detection. *Artificial Intelligence Review*, 58, 279. <https://doi.org/10.1007/s10462-025-11287-7>
13. Duffy, C., Hassanshahi, M., Jastrzebski, M., and colleagues. (2025). Unsupervised beyond-standard-model event discovery at the LHC with a novel quantum autoencoder. *Quantum Machine Intelligence*, 7, 41. <https://doi.org/10.1007/s42484-025-00258-4>
14. Ma, L., Zhang, X., Wang, K., et al. (2025). Generative adversarial message passing-based anomaly detection. *Journal of King Saud University — Computer and Information Sciences*, 37(5). <https://doi.org/10.1007/s44443-025-00021-6>

15. García-Beni, J., et al. (2025). Time-series quantum reservoir computing with cluster states (related). *EPJ Quantum Tech.* <https://doi.org/10.1140/epjqt/s40507-025-00352-3>
16. Hammadia, T., Saber, A. M., and Kundur, D. (2025). Quantum variational circuits for detection of false data injection. In *2025 IEEE Industry Applications Society Annual Meeting (IAS)*. <https://doi.org/10.1109/IAS62731.2025.11061524>
17. iang, X. Q., et al. (2025). Dual flow reverse distillation for unsupervised anomaly detection. *Digital Signal Processing (Elsevier)*. <https://doi.org/10.1016/j.dsp.2025.105258>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.