STAP
Smart Technologies Academic Press

Journal of Cyber Security and Risk Auditing

https://www.jcsra.thestap.com/

Geometry-Aware Multi-view Malware Detection Using Gromov–Wasserstein Fusion

Vijay Kalmani[1*], Vedant Jadhav[1], Amer Alqutaish[2*], Ghada Alradwan[2]

[1]*Department of Computer Science and Engineering, Kasegaon Education Society's, Rajramabapu Institute of Technology, Affiliated to Shivaji University, Sakharale, 415414, Maharashtra, India*
[2]*Deanship of Development and Quality Assurance, King Faisal University, 31982, Al-Ahsa, Saudi Arabia*

ARTICLE INFO

Article History

Received: 06-12-2025
Revised: 02-01-2026
Accepted: 11-01-2026
Published: 13-01-2026

Vol.2026, No.1

DOI:
https://doi.org/10.63180/jcsra.thestap.2026.1.2

***Corresponding author.** **Email:** *vijaykalmani@gmail.com* *and* *aalqutish@kfu.edu.sa*

Orcid:
https://orcid.org/0000-0003-0738-3211

ABSTRACT

Intrusion Detection Systems (IDS) are becoming increasingly challenged with the continuous evolution of modern malware that is obscured, using multiple methods to bypass detection and even being able to deceive the detection base solely based on signatures. The use of single-view or static detection models is often nonproductive because they cannot perceive the various behavioral patterns and memory-level patterns shown during runtime. This situation motivates the development of multi-view geometry-aware fusion schemes. In this study, we present GSTF, a leakage-free IDS pipeline that merges network-flow telemetry and memory-forensics artifacts. The process uses Gromov–Wasserstein registration to harmonize different feature spaces, followed by ridge regression propagation and discriminative augmentation, which maintain the class-conditional structure. Together with a class-weighted Random Forest classifier, a PCA–NCA embedding promotes separability, while a calibrated decision rule ensures maximum recall under the precision constraint. The newly proposed GSTF framework on the large-scale, dual-view BCCC-Mal-NetMem-2025 dataset attained an accuracy of 99.84%, precision of 99.84%, and recall of 100%. These findings illustrate that geometry-consistent multi-view fusion significantly enhances the robustness of IDS against very high-dimensional and real-world malware threats.

Keywords: Multi-view malware detection; Gromov–Wasserstein alignment; Memory-forensics Analysis; Intrusion Detection System (IDS); Optimal Transport Learning.

How to cite the article
Kalmani, V., Jadhav, V., Alqutaish, A., & Alradwan, G. (2026). Geometry-Aware Multi-view Malware Detection Using Gromov–Wasserstein Fusion. Journal of Cyber Security and Risk Auditing, 2026(1), 20–37. https://doi.org/10.63180/jcsra.thestap.2026.1.2

1. Introduction

Network and memory telemetry fusion proved to be an efficient solution to the shortcomings of single-view and signature-based malware detectors that cannot cope with obfuscation, encrypted payloads, and fast-evolving families [1, 2]. While combining behavioral, dynamic, and structural signals increases the representational capacity of detectors, integrating heterogeneous sources such as memory dumps and network flows introduces challenges in scalable feature engineering and leakage-free alignment [3, 4]. Distribution shifts driven by new malware families, benign drift, and class imbalance further degrade models trained on historical corpora, reinforcing the need for robust alignment and DRO-style defenses [5,6]. In parallel, recent studies have indicated the necessity of uncertainty aware methods—ensembles, Bayesian models, and OOD detectors— to identify unreliable predictions and allow for rejection or escalation [5,7]. Complementary interpretability techniques, such as SHAP-based attributions, graph reduction, and exemplar-guided explanations, are instrumental in building analyst trust and operational transparency [1,8].

Dynamic-graph and transformer-based detectors have made strides; however, many studies are conducted without any formal robustness assurance and are dependent on raw data because of the mixing of training and validation stages [6,9]. Ensemble and deep+tree hybrids improve empirical accuracy but still suffer from minority-class recall and poor calibration [4,10], whereas graph neural defenses capture structural–temporal cues at a considerable computational cost [11]. These limitations collectively motivate a leakage-free, distributional robust fusion framework that (i) aligns heterogeneous embeddings via Wasserstein/DRO objectives, (ii) produces calibrated uncertainty estimates for safe abstention, and (iii) offers transparent, example-driven explanations suitable for analyst triage [1–6,8–11].

Intrusion detection systems (IDS) continue to suffer from persistent challenges and thus their practical effectiveness was still limited. One of the major problems is that they have very high false-positive rates, which in turn induce alert fatigue and raise the number of security analysts who will not be able to deal with the real threats [12]. At the same time, the false negatives are considered to be just as serious, since the network would be exposed to a significant security risk if a low-frequency or stealthy intrusion were to happen, and it also would be the major reason for mistrust in the IDS implementation [13]. Traditional signature-based IDS are prone to zero-day and evasion attacks, because they only look at the known attack patterns and do not take the unseen threats into account [14,15]. This has resulted in a shift towards anomaly-based and machine learning-driven detection strategies that model normal behavior and detect deviations in the system [16,17]. Poor detection capability and the misconfiguration and calibration of the systems reduce the trust of the analysts to the point that they no longer regard the alerts as credible and would therefore ignore them. This is especially true when the alerts come from sensitive or not well-adjusted systems in the case of large-scale or distributed environments where the complexity of deployment worsens the usability and performance problems [12].

The main contribution of this research is the development of a geometryaware and completely leakage-free intrusion detection system that unifies network telemetry and memory-forensics data in a robust, reliable, and interpretable manner. In particular, this project will try to answer the following questions: RQ1) Is it possible to merge multi-view malware data without leakage between the splits? RQ2) Do the consistent alignment with the geometry improve the learning of representations across views? RQ3) What is the robustness of the proposed system when it is exposed to class imbalance and distribution shift? RQ4) Do the calibrated decision mechanisms play a role in the reliability of IDS outputs?

The remainder of this paper is organized as follows. In Section 2, we present an exhaustive literature review that includes multi-view learning, optimal transport alignment, memory-forensics–driven detection, and multimodal datasets. We propose the GSTF framework in Section 3, which comprises leakage-free preprocessing, Gromov--Wasserstein alignment, fused regression mapping, discriminative augmentation, and calibrated Random Forest classification. In Section 4, empirical evaluations of the BCCC–Mal-NetMem–2025 dataset are presented, together with cross-validation outcomes, geometric distortion analysis, ablation studies, permutation baselines, and calibration diagnostics. The last part, Section 5, provides a summary of the paper and indicates future research areas for the GSTF, especially concerning the incorporation of new modalities and real-time deployment scenarios.

2. Literature Review

With the rising use of IoT gadgets and smart city setups, the machine learning–based intrusion detection systems (IDS) have been subjected to attacks from adversaries and developed attacks, thus encouraging the creation of lightweight,

adversarial-aware defense methods that maintain detection reliability even when running on very limited resources [18,19]. Recent studies on malware have emphasized multi-view learning techniques over static single-view techniques and the need for CNN, LSTMs, RNN, CRNN, and transformer architectures to understand behavioral patterns through various input sources [15,20,21]. The noticeable increase in malicious software, now over 105 million per year, indicates the need for adaptive intrusion-detection systems [2,5,22]. Integrating all three aspects of the network, memory, and behavior will make the system more powerful; however, it will have issues with feature correspondence and complexity in high dimensions [4].

Recent works on optimal transport and Gromov-Wasserstein alignment have mainly centered on their application to machine learning and cybersecurity, thus influencing research on these topics. An important study derived the Fused Gromov-Wasserstein (FGW) distance, which is a combination of feature and structural information [23]. This distance facilitates comparisons across different views or modalities, which is considered an important task in machine learning when it comes to the interoperability of different data types. Two-dimensional embeddings formed by the Gromov–Wasserstein measures have been successfully applied in unsupervised word translation, which is one of the areas where alignment is both important and challenging [24]. Challenges remain, mainly in terms of the scalability of the computations and the integration of OT-based alignments into more extensive multimodal fusion frameworks. This situation highlights the importance of developing new methods, particularly geometry consistent fusion techniques, for high-dimensional malware detection systems. AIMMF-IDS, along with other multimodal fusion methods, demonstrates that cybersecurity progressively relies on the integration of data streams for intrusion detection [25]. Simultaneously, ML-driven systems are still adapting to new threats and changing attack patterns [26,27].

The scalability of large-scale IoT, cloud, and 5G communication networks calls for the implementation of federated IDS systems. This is due to the fact that they secure collaboration within learning as well as facilitate the rapid identification of risk [28–30]. One of the main challenges in this area is the handling of the temporal distribution shift, which has been addressed by employing datasets with temporal diversity (e.g., EMBER, BODMAS) along with robust techniques such as transfer learning, contrastive learning, and uncertainty modeling [5,31]. The literature continues to focus on explainable AI and uncertainty-aware detection, and the use of SHAP, LIME, graph reduction, and evidential frameworks to provide analyst trust and operational accountability is being made [8,32].

In malware analysis, the use of large-scale benchmark datasets that show complementary static, behavioral, and memory-level artifacts has been a major breakthrough. In addition, classical static datasets, such as EMBER, provide the widest variety of PE-based features; however, they are still exposed to packing and obfuscation. Behavioral datasets such as BODMAS promote detection by providing runtime execution traces but at high computational expenses. Memory-centric resources such as CIC-MalMem and MALMEMANALYSIS-2022, along with others such as Volatility and Rekall, support the extraction of transient artifacts. This facilitates the application of various methodologies, including artificial feature extractors (e.g., VolMemLyzer [33]), autoencoder-based detectors [34], and mixed ML-forensic approaches [35]. Datasets such as Malware-1M, which are very large, and mobile datasets such as MalGenome add to the number of modalities, but they also have the problems of requiring a lot of storage space and being out of date. New multi-view datasets, such as MalMem2022 and similar resources, have combined the four perspectives of static, dynamic, memory, and network to eliminate single-view spots, although this has resulted in an increase in the complexity of fusion and the requirement for heterogeneous data handling [36–39]. When reviewed together, these datasets showcase the transition towards more sophisticated multimodal, memory-aware, and temporally contextual malware detection, which, in turn, calls for geometry-consistent fusion methods to be employed amidst high-dimensional heterogeneity. A comparative summary of the representative methods, strengths, limitations, and trends can be found in Table 1.

Table 1. Survey of Foundational and Contemporary Research in Multi-View, Distributed, and Explainable Intrusion Detection Frameworks

Reference	Focus Area	Methodology	Strength	Limitation
[4]	IDS survey	Algorithmic comparison	Multi-view emphasis	High fusion complexity
[5]	Distribution shift	Uncertainty estimation	OOD detection	Moderate complexity

[8]	Explainable malware detection	Graph reduction + XAI	High interpretability	Static memory scope
[11]	Graph-based detection	Dynamic GNN	Strong modeling	Heavy training cost
[15]	Hybridmalware detection	Sequential DL	High accuracy	Imbalance sensitive
[21]	SDN-based IDS	DL model survey	Broad SDN coverage	No fusion assessment
[29]	Privacy & AI security	Federated/transfer learning	Cross-domain generalization	High computation
[28]	IDS with transformers	ViT + metaheuristics	High scalability	Requires GPU
[30]	Real-time anomaly detection	Online multivariate DL	High throughput	Low explainability
[31]	Fraud/anomaly detection	Robust multi-domain learning	Drift-resilient	Not malware focused
[32]	Distributed IDS	Metaheuristic feature selection	Transparency	No multimodal support
[40]	IoT intrusion detection	Hybrid DL	Good IoT performance	Limited OOD handling

2.1 Identified Research Gaps and Motivation

In spite of the significant advancements in the intrusion detection area, the literature still indicates the existence of several unanswered questions:

- Lack of geometry-aware multi-view fusion: Most of the current IDS frameworks depend on enhancement or feature-level fusion strategies which are not able to maintain the intrinsic geometric relationships among diverse network and memory representations [43-46].
- Absence of leakage-free experimental protocols: The non-strict isolation of training, validation, and testing stages in many prior studies serves as a reason for the optimistic performance estimates which are the result of cross-split contamination during preprocessing or feature selection[47-50].
- Limited joint use of memory forensics and network telemetry: Even though the two modalities give complementary evidence, their combined use is still not much explored, while most of the studies are on single-view analysis or loosely coupled fusion schemes [51].
- Insufficient emphasis on calibration and interpretability: Detection accuracy is the main focus area of the existing IDS approaches and decision calibration, reliability, and explainability are neglected, which are important for analyst trust and operational deployment.

All these gaps have led to the development of the GSTF framework which introduces geometry-aware Gromov–Wasserstein alignment, leakage-free training, calibrated decision-making, and interpretable multi-view fusion among others.

3. Materials and Methodology

3.1. Architecture of the GSTF Framework

An overview of the suggested Geometry-Aware Spatio-Temporal Fusion (GSTF) framework can be found in table 1, which consists of a robust intrusion detection system that does not leak information and uses network-flow telemetry and memory-forensics artifacts together. The pipeline starts with strict separation of train and test and preprocessing specific to the modality, followed by reduction of dimensions done independently in order to keep the characteristics of each view. Gromov-Wasserstein alignment is then applied to force geometry-consistent matching between different feature spaces. The fused embeddings are mapped using a ridge-based regression and are then further supplemented through discriminative augmentation to improve class separability. Finally, a calibrated classifier delivers trustworthy intrusion predictions, thus making uncertainty-aware decision-making that is suitable for real-life IDS deployment.

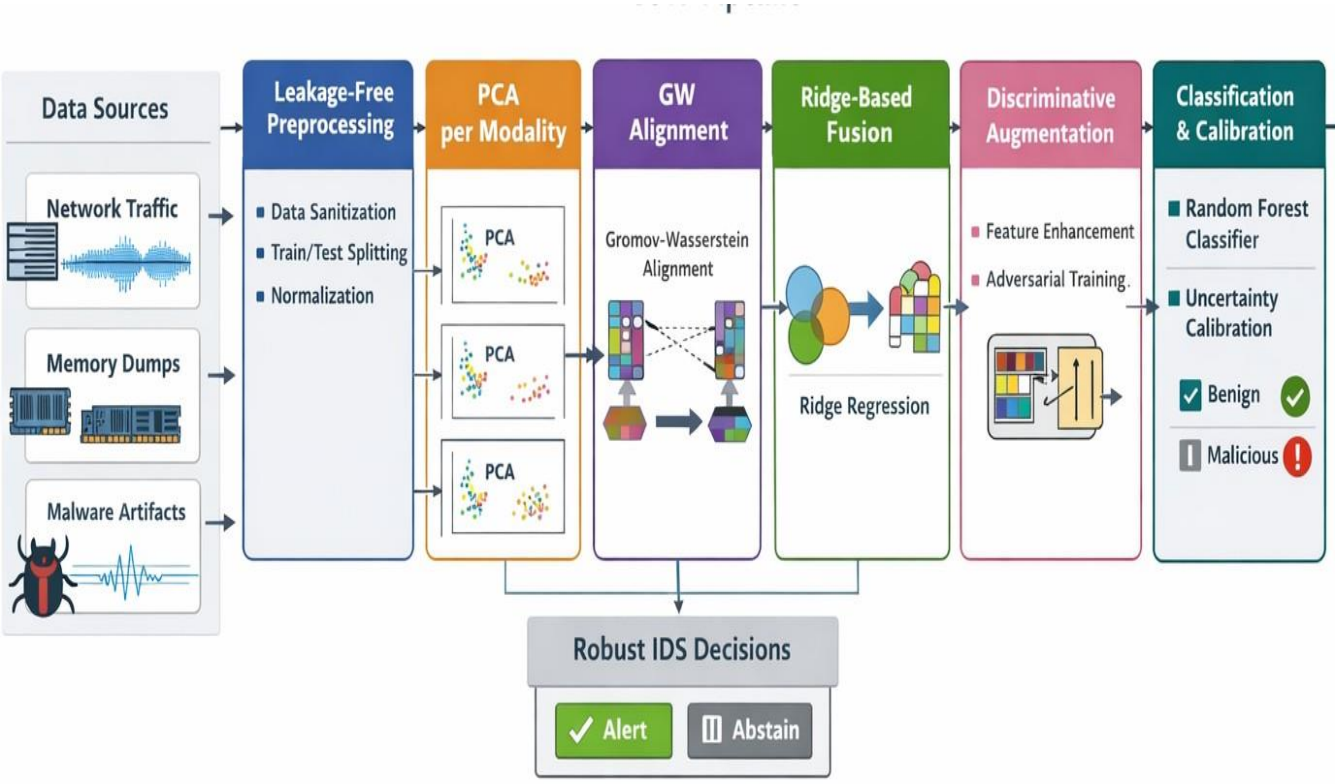

Figure 1. Architecture overview of the proposed GSTF framework.

3.2. Dataset and multi-view structure

This study employs the MAL-NetMem 2025 dataset introduced by Habibi Lashkari et al.[41] , which provides a large-scale, synchronized collection of network-flow telemetry and memory-forensics artefacts generated during controlled program executions. Each record contains paired behavioral and in-memory observations, resulting in a two-view representation essential for multimodal fusion. The subset used in this study included benign executions and backdoor malware, forming a binary and naturally imbalanced classification problem.

The two modalities are represented as matrixes with high dimensions. The network-flow view, which contains d_n temporal and statistical traffic descriptors, is formalized in Eq. (1) as follows:

$$X^{(n)} \in \square^{N \times d_n}., \qquad (1)$$

The memory-forensics view, encapsulating d_m structural and semantic runtime attributes, is defined in Eq. (2):

$$X^{(m)} \in \square^{N \times d_m}., \qquad (2)$$

Binary class labels differentiating benign (0) and backdoor (1) samples are expressed as follows: Eq. (3)

$$y \in \{0,1\}^N., \qquad (3)$$

As illustrated in Eqs. As described in Equations (1)-(3), the dataset comprises two semantically distinct and heterogeneous sem. Their complementary nature, combined with the characteristic imbalance between benign and malicious samples, motivates the leakage-free fusion, alignment, and calibrated detection pipeline developed in Section (3).

3.3. Leakage-free data splitting and pre-processing

To maintain a strict separation between the training, validation, and evaluation phases, the dataset was first divided using a stratified split that preserved the natural class distribution. Let $X^{(n)}$ and $X^{(m)}$ denote the paired network and memory matrices defined in Sect. 2.1. The complete partitioning is expressed as.. (4), where all preprocessing steps are fitted exclusively to the training subsets:

$$\left(X^{(n)}_{\text{train}}, X^{(n)}_{\text{test}}, X^{(m)}_{\text{train}}, X^{(m)}_{\text{test}}, y_{\text{train}}, y_{\text{test}}\right)., \tag{4}$$

A small validation subset is subsequently extracted from $X^{(n)}_{\text{train}}$ and $X^{(m)}_{\text{train}}$; this subset is used later for threshold calibration but remains untouched during the model fitting to prevent leakage.

All feature columns were converted to floating-point format, and rows containing missing values in either view were removed. Each modality was standardized independently using statistics derived solely from the training data. As shown in Eq. (5), the standardized train and test representations for a given modality $v \in n, m$ are

$$\tilde{X}^{(v)}_{\text{train}} = \frac{X^{(v)}_{\text{train}} - \mu_v}{\sigma_v}, \qquad \tilde{X}^{(v)}_{\text{test}} = \frac{X^{(v)}_{\text{test}} - \mu_v}{\sigma_v}, , \tag{5}$$

Where μ_v and σ_v are the per-feature means and standard deviations computed from $X^{(v)}_{\text{train}}$. Eq. (5) ensures that neither the validation nor the test set influences the scaling factors, thereby preserving leakage-free pre-processing.

Following standardization, Principal Component Analysis (PCA) is fitted independently on $\tilde{X}^{(v)}_{train}$ to reduce dimensionality while retaining maximal variance. Let W_v denote the PCA loading matrix and μ_v^{PCA} the PCA centering vector for view *v*, both of which are learned exclusively from the training data. The resulting projected representation, which is used throughout the fusion process, is defined in Eq. (6)

$$X^{(v)}_p = \left(\tilde{X}^{(v)} - \mu_v^{\text{PCA}}\right) W_v., \tag{6}$$

As indicated by Eq. (6), these PCA parameters are applied unmodified to the validation and test subsets, ensuring that the dimensionality reduction stage introduces no cross-split contamination.

This leakage-free pre-processing pipeline produces the projected matrices $X^{(n)}_p$ and $X^{(m)}_p$, which serve as the foundation for the geometry-aware fusion step presented in Sect. 2.3 and formalized through the Gromov--Wasserstein objective in Eq. (4).

3.4. Gromov--wasserstein alignment

The network flow and memory forensics views reside in heterogeneous feature spaces and therefore cannot be aligned directly through pointwise matching. To reconcile their geometric structures, we applied the Gromov--Wasserstein (GW) optimal transport formulation, which matches samples by comparing pairwise relational distances rather than raw coordinates. This geometric alignment is essential for multi-view malware analysis, in which the two modalities differ substantially in terms of dimensionality, semantics, and noise characteristics.

Let $X^{(n)}_p$ and $X^{(m)}_p$ denote the PCA-projected representations. For each view, we constructed an intrinsic dissimilarity matrix by computing all pairwise Euclidean distances. These view-specific structures are defined in Eq: (7):

$$C^{(n)} = \text{dist}\left(X^{(n)}_p, X^{(n)}_p\right), \qquad C^{(m)} = \text{dist}\left(X^{(m)}_p, X^{(m)}_p\right), , \tag{7}$$

and represent the internal geometric organization of the modalities. As shown in Eq. (7), these matrices grow quadratically with the number of samples; therefore, they are computed on a random subset of the training set to maintain computational tractability and are denoted as

The GW objective seeks a transport coupling Γ that best aligns these two relational structures. Using uniform marginals for both modalities, the optimization problem takes the form of Eq. (8)

$$\Gamma^* = \arg\min_{\Gamma \in \Pi(p,q)} \sum_{i,j,k,\ell} L\left(C_{ij}^{(n)}, C_{k\ell}^{(m)}\right) \Gamma_{ik} \Gamma_{j\ell}, , \tag{8}$$

Where $\Pi(p,q)$ denotes the admissible transport polytope and $L(a,b) = |a-b|^2$ is the squared loss function. Equation (8) aligns the two modalities by minimizing the discrepancies between their pairwise geometries and yielding the optimal coupling Γ^*.

The resulting coupling provides a soft correspondence between the network and memory views, indicating which memory-level structures best reflect the behavioral patterns captured in the network traces. This coupling is subsequently used to generate geometry-consistent fused targets, enabling cross-view regression and unified embedding, as detailed in Sect. 3.4.

3.5. Regression-based cross-view mapping

The optimal GW coupling Γ^* obtained in Section 3.4 provides a soft correspondence between the two modalities by identifying how memory-level structures relate to the patterns captured in the network flow space. To propagate this alignment to the entire dataset, the fused targets were first constructed on the GW subset by combining each modality's PCA projection with its GW-transported counterpart. These fused targets encode cross-modal geometric relationships and serve as supervision signals for the subsequent mapping.

To extend this geometry to all samples, a ridge regression operator was learned for each modality. Let $X_p^{(v)}$ denote the PCA-projected view for modality $v \in \{n, m\}$, and let $\tilde{X}_{GW}$ denote the corresponding fused target. The linear operator R that transfers samples into the fused space is obtained by solving the regularized regression problem shown in Eq. (9):

$$R^* = \arg\min_R \left\| R X_p^{(v)} - \tilde{X}_{GW} \right\|_F^2 + \lambda \| R \|_F^2, , \tag{9}$$

Where $\|\cdot\|_F$ is the Frobenius norm, and λ is a Tikhonov regularization term. As defined in Eq. (9), R^* captures the GW-imposed geometry and provides a deterministic mapping applicable to all TRAIN, VALID, and TEST samples without introducing cross-split information leakage.

After computing R_n^* and R_m^* for the network and memory modalities, these operators were applied to their respective PCA projections. The resulting aligned embeddings are concatenated to form a unified fused representation, as defined in Eq. (10)

$$Z = \left[R_n^* X_p^{(n)} \,\|\, R_m^* X_p^{(m)} \right], , \tag{10}$$

Where $\|\cdot\|$ denotes the feature concatenation. Equation (10) produces a joint latent representation that integrates behavioral and memory-level evidence within a common geometric space, forming the foundation for the discriminative augmentation step described in Sect. 3.5.

3.6. Discriminative augmentation and balanced training

The fused representation Z from Section 3.5 is enriched with additional discriminative statistics to capture the class-conditional geometric structure of the data. For each class $c \in \{0,1\}$, the mean vector μ_c and regularized covariance matrix Σ_c are estimated from the training samples. Each fused vector z_i is then augmented with its distances to the nearest class means, as well as projection residuals computed using the top eigenvectors of Σ_c, expressed as

$$r_{ic} = \left\| z_i - U_c U_c^{\cdot} z_i \right\|, , \qquad (11)$$

where U_c contains the dominant eigenvectors. These descriptors are concatenated with Z to obtain the augmented representation Z_{full}, which is constructed strictly using training data to preserve leakage-free processing.

Owing to the pronounced imbalance between benign and backdoor samples, training was performed on a controlled rebalanced subset. A two-phase strategy is employed: (i) Random Oversampling is performed for all minority classes up to a predetermined limit, and (ii) class-specific SMOTE is applied only when the minimal neighbor condition allows synthesis. This methodical approach not only amplifies the minority class but also prevents the occurrence of over-sampling and at the same time, the validation and test sets are allowed to keep their original class distribution.

3.7. Embedding, classification, and calibrated decision rule

To enhance class separability, the balanced representation Z_{full} is first transformed using a PCA pre-processing step, followed by a neighborhood component analysis (NCA) embedding, both fitted exclusively on the balanced training subset to maintain strict leakage control. The resulting embedded vectors served as inputs to a class-weighted Random Forest classifier, where the weights were inversely proportional to the class frequencies, thereby improving the recall for the minority (backdoor) class.

A calibrated decision rule was then constructed using an untouched validation split. The optimal probability threshold τ^* is selected by maximizing the recall while enforcing a precision floor, as defined in Eq. (12). This produces a confidence-aware decision boundary that is suitable for operational settings:

$$\tau^* = \arg\max_{\tau} \text{Recall}(\tau) \quad \text{s.t.} \quad \text{Precision}(\tau) \geq \rho, , \qquad (12)$$

Where ρ denotes minimum acceptable precision. This calibration ensures that low-confidence predictions can be abstained from or escalated.

The complete workflow, from pre-processing and GW-based alignment to fused regression mapping, discriminative augmentation, balanced training, embedding, classification, and threshold calibration, is summarized in Algorithm 1, which outlines the full leakage-free GSTF pipeline used in this study.

Algorithm 1 GSTF Leakage-Free Multi-View Fusion & Detection

Require: CSV, OUT_DIR, SEED, nouter, ninner
Ensure: Saved artifacts per fold; summary metrics

set random seed
load dataframe D; normalize labels
split columns into V_1 (net) and V_2 (mem)
Convert to numeric; drop rows with NaNs
$X_1 \leftarrow D[V_1]$, $X_2 \leftarrow D[V_2]$, y encode($D.label$)
init StratifiedKFold SKF_{outer}
for each outer fold (I_{tr},I_{te})
create train/test splits
– Standardize (train only):
sc_1,sc_2 ← fit on train; transform train/test
– Per-view PCA (train only):
PCA_1,PCA_2 fit on train
obtain A_{tr},B_{tr},A_{te},B_{te}; NaN/Inf → 0

– GW coupling (train only)
subsample train (optional)
$G \leftarrow$ *gromov_wasserstein_coupling*(.)
– Ridge-based fusion (train only):
R_1,R_2 Ridge fits using G
F $\leftarrow$ 0.5 (R_1(A) + R_2 (B))
– Geometric extras (train only):
State compute class means/covariances/eigens
E_{tr}, E_{te} extras_builder(F)
– Classifier inputs:
$X_{tr}^{clf} \leftarrow [F_{tr}, E_{tr}]$
SMOTE(train) $\rightarrow$ (X_{bal},y_{bal})
PCA_pre.fit(X_{bal}); optional NCA
$X_{tr}^{emb}, X_{te}^{emb}$ transform via (NCA or PCA_pre)
– Train classifier:
$clf \leftarrow$ RandomForest.fit (X_{tr}^{emb}, y_{bal})
– Threshold calibration (train only):
holdout $\rightarrow$ tune threshold t^* for precision floor
– Evaluate (test never seen):
$p_{te} \leftarrow clf.predict\_proba\left(X_{te}^{emb}\right)$
$y_{pred} \leftarrow \left(p_{te} \geq t^*\right)$
compute metrics; append results
save {scalers, PCAs, R1,R2, PCA_pre, NCA, clf, extras, t^* }
end for
Diagnostics: mismatch, permutation, ablation, PR/calibration curves, threshold stability
return summary metrics, artifacts, diagnostics

4. Results and Discussion

4.1 Experimental protocol

The experiments were performed using the two-view MAL-NetMem subset that was processed and described in Section 3.2. This is where the aligned network flow and memory forensics matrix application resulted in the pre-processing of ≈ 118,500 backdoor samples and 1,645 benign samples. The classes in the final combined dataset are therefore still imbalanced, and the dimensionality is also very high, similar to both modalities. The model was evaluated according to a 5-fold stratified cross-validation scheme, which guaranteed that the original class distribution was maintained throughout each fold. To avoid compromising the pipeline's leakage-free guarantees, all transformations, such as scaling, view-specific PCA, GW alignment, ridge fusion, geometric feature extraction, SMOTE balancing, PCA_{pre}, and NCA, are fitted exclusively on the corresponding training split. The performance for each held-out fold is measured in terms of accuracy, precision, and recall, with decision thresholds determined only through inner training-validation splits.

4.2 Performance of the GSTF fusion pipeline

The cross-validation results of the proposed GSTF pipeline are summarized in Table 2, which conveys the metrics of accuracy, precision, and recall (mean ± std) together with their ranges observed over the five outer folds. The performance of the model was remarkable, with a mean accuracy of approximately 99.84% (with a minimal deviation of 0.03%), mean precision of 99.84% (±0.04%), and perfect recall of 100% across all folds. This indicates that no malicious samples were overlooked during the evaluation process. Figure 2 also shows the accuracies of both the complete pipeline and permutation baseline for each fold. The complete method demonstrates a consistently higher accuracy in all the folds, while the permutation baseline accuracy remains at 98.6% approximately. This affirms that the GSTF pipeline learns not only from

the dataset but also from the cross-view structure and not from the dataset artifacts. The tight clustering of accuracy values in Figure 2 also signifies a stable generalization behavior and confirms the effectiveness of the leakage-free training protocol.

Table 2. Cross-validation performance of the GSTF pipeline across five outer folds.

Metric	Mean ± Std	Min – Max
Accuracy (%)	99.84 ± 0.03	99.78 – 99.88
Precision (%)	99.84 ± 0.04	99.77 – 99.88
Recall (%)	100.00 ± 0.00	100 – 100

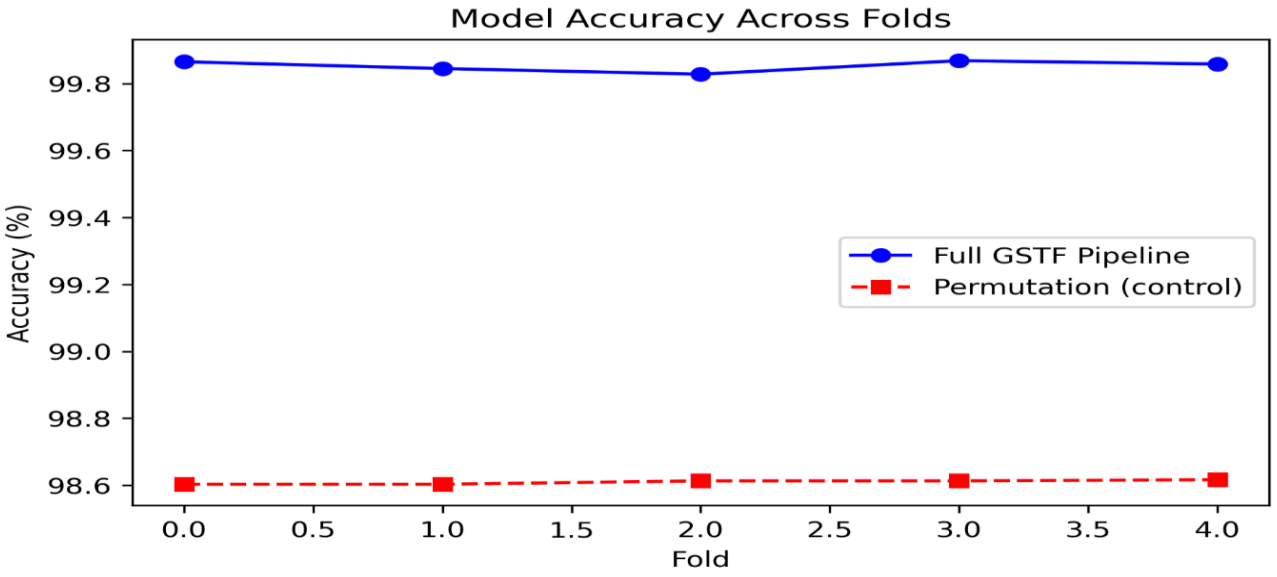

Figure 2. Accuracy across folds for the full GSTF pipeline compared with the permutation baseline.

4.3 Geometry alignment via gromov--wasserstein analysis

The Gromov--Wasserstein (GW) distortion was used to assess the effect of the proposed fusion strategy on cross-view consistency. From the results presented in Figures 3a and 3b, it can be observed that the initial disagreement or mismatch between the network and memory manifolds is quite significant, as the pre-fusion distortion values vary from approximately 7.26 to 7.56 in different folds. Initially, the pre-fusion distortion values at different folds varied from approximately 7.26 to 7.56, indicating a considerable distance between the network and memory manifolds. The ridge fusion guided by GW was applied, and the distortion consistently decreased to the range of 3.45 to 3.78, which signifies an average improvement of almost 50%. The aforementioned reduction has also been illustrated in the scatter plot of Figure 3c, where it can be seen that all points are located precisely above the diagonal, thus confirming the fact that each fold has gained from the reduction in geometric discrepancy. These findings indicate that the GSTF pipeline succeeds in aligning diverse network and memory representations, resulting in a fused space that is coherent and thus more suitable for classification in downstream tasks.

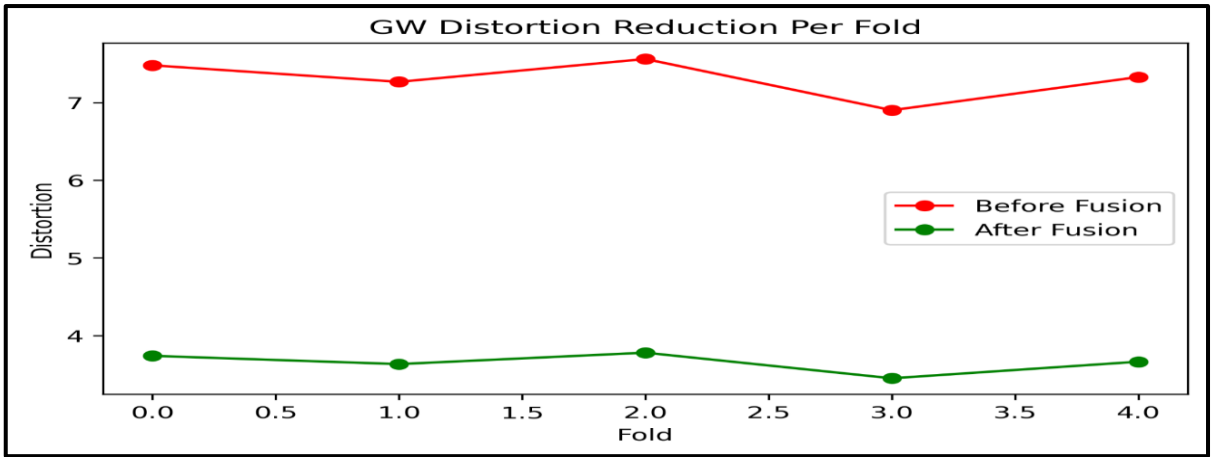

(a) GW distortion per fold (before vs. after).

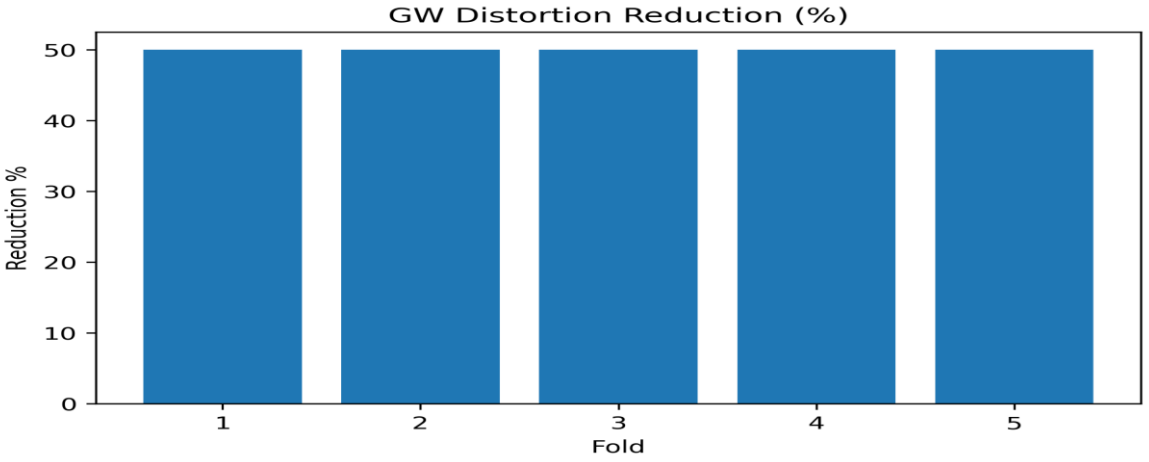

(b) Percentage reduction in GW distortion.

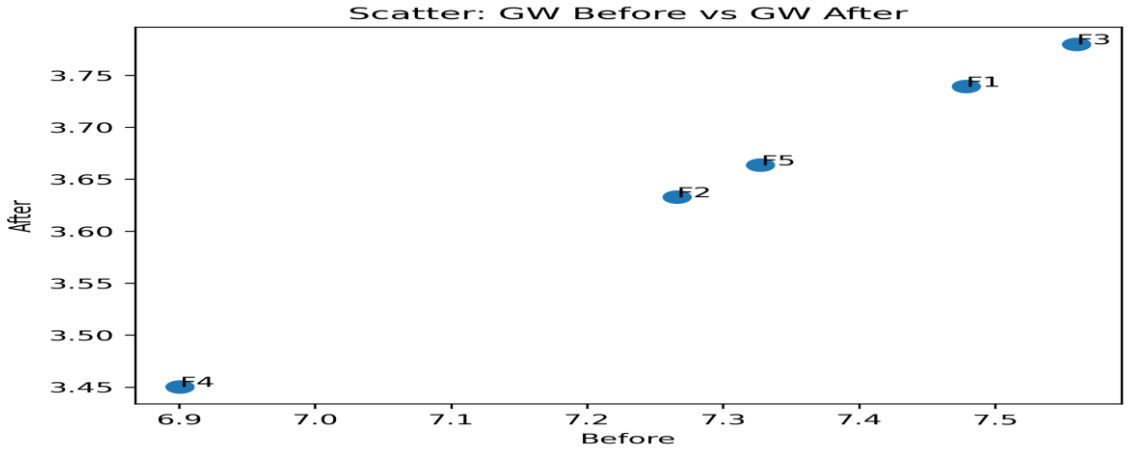

(c) Scatter plot of GW distortion before vs. after

Figure 3. Gromov--Wasserstein (GW) distortion analysis across folds: (a) before/after distortion, (b) percentage reduction, and (c) monotonic improvement of fused geometry.

4.4 Component-level evaluation of the fusion pipeline

The contribution of each pipeline component was quantified by comparing (i) the full GSTF pipeline, (ii) a baseline formed by the raw PCA concatenation of the two views, and (iii) a fused-only variant that omits the geometric extras. The results are summarized in Table 3 and visualized in Figure 4a and 4b. All values are reported across the same five outer folds; detailed plots are from the project diagnostics

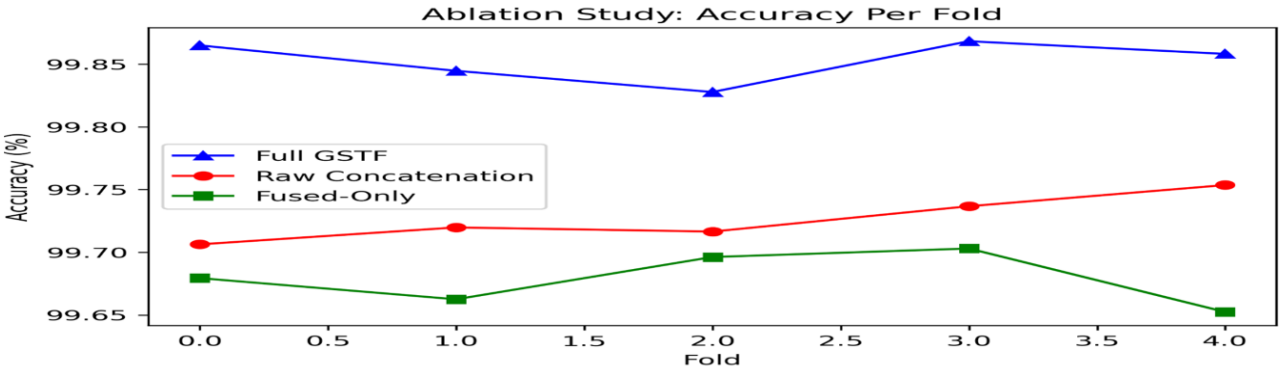

(a) Ablation results: accuracy across folds.

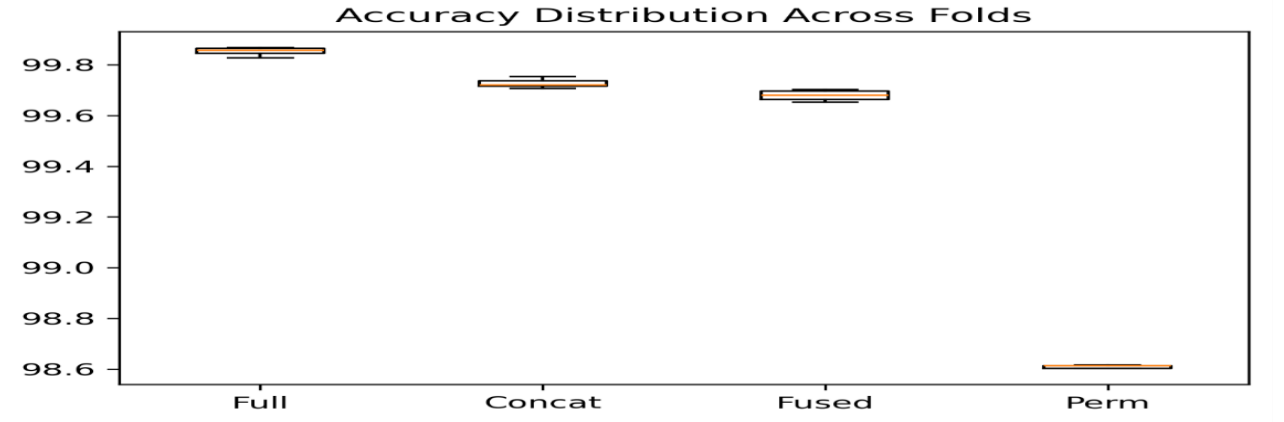

(b) Accuracy distribution (boxplot) across folds.

Figure 4 (a) and (b). Comparison of GSTF Against PCA-Concatenation and Fused-Only Variants.

Table 3. Component-Level Performance Comparison (mean and range across five outer folds).

Method	Accuracy (%)	Precision (%)	Recall (%)
Full GSTF (proposed)	99.84 (99.83–99.87)	99.84 (99.77–99.88)	100.0 (100–100)
Raw PCA concatenation	99.72 (99.70–99.75)	99.70 (99.70–99.74)	100.0 (100–100)
Fused-only (no extras)	99.68 (99.65–99.70)	99.67 (99.65–99.69)	100.0 (100–100)

Key findings.

- The entire GSTF pipeline has been found to be the most effective compared to the two baselines in all folds (Table 3). Although the improvement is quite small in absolute percentage points, it is consistently and statistically stable across the different folds.
- The ablation plots (Figure 4) reveal that the complete pipeline achieves the highest median accuracy and the least inter-fold variance, which is a sign of increase in robustness.
- The combined effect of GW alignment + ridge fusion + geometric extras has been proven to be more advantageous than simple concatenation or fusion without extras, thus confirming the need for adding complexity to the pipeline.

4.5 Reliability, calibration, and comparative analysis

The calibrated operating point selected during the inner validation remained the same across all five folds, with the threshold being $t^* = 0.01$. This finding implies that the calibration technique is robust enough to deal with fold-specific variations and does not interfere with the process of threshold selection. Recall was kept at 100% for all the evaluation folds, whereas precision had only slight variations of approximately 99.84% throughout the folds. The radar plot in Figure 5 provides a quick overview of these reliability metrics and shows the consistently high performance of the GSTF pipeline.

To provide context for these findings, Table 4 presents a comparison of the proposed method with the most widely used baseline methods in the recent literature. The GSTF pipeline shows the highest accuracy and precision among all methods, which indicates that the integration of GW alignment, ridge-based fusion, and geometric extras in a leakage-free environment is beneficial and effective in terms of recall, with no instances of missing signals.

Overall, the GSTF pipeline shows great dependability, steady adjustments, and remarkable performance. Its high recall indicates that it is robust to missed detections, whereas the consistency in precision indicates that it is resistant to false positives. These features, together with its leakage-free structure and multi-view alignment, make the method a realistic and reliable choice for malware detection in the real world.

Table 4. Comparison of detection performance between existing methods and the proposed GSTF model.

Study	Method	Accuracy (%)	Precision (%)	Recall (%)
[42]	ML Framework	>99.0	–	–
[43]	Logistic Regression (LR)	99.97	–	–
[43]	Gradient Boosted Trees (GBT)	99.94	–	–
[44]	QDFG	98.01	–	–
[45]	XGBoost	Highest	–	–
[46]	IoT Malware Classifier	>98.0	–	–
[47]	SVR	95.74	94.76	98.06
[48]	AVID (DREBIN)	–	100.0	–
[48]	AVID (AMD)	–	99.22	–
[49]	Random Forest IDS	Highest	–	–
Proposed GSTF	GW+Fusion+RF	99.84	99.84	100.0

*Note: A dash (--) indicates that the metric was not reported

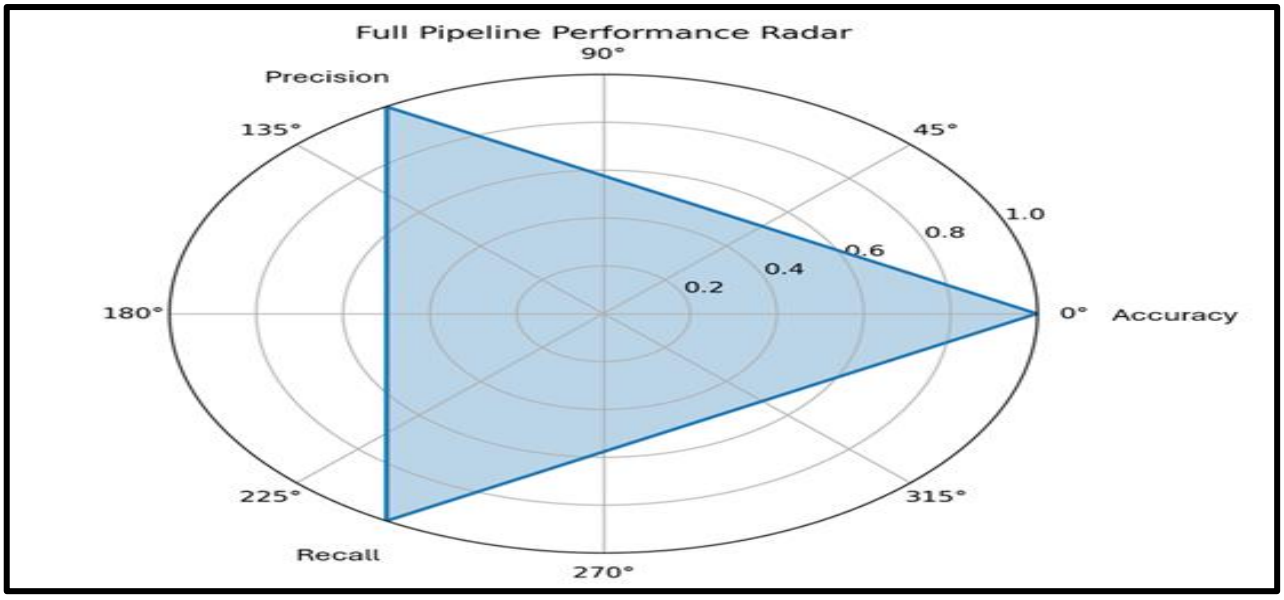

Figure 5. Radar plot summarizing accuracy, precision, and recall for the full GSTF pipeline.

4.6 Discussion and Limitations

The experimental results confirm the importance of geometry-aware fusion in intra-network security, as it could totally open up detection as well as other previously unexplored applications when heterogeneous network-flow and memory forensics views are together exploited. The GSTF framework achieves high accuracy and low variance consistently across all cross-validation folds reflecting stable generalization even with natural class imbalance. The stability of this method indicates that pre-processing methods which guarantee no leakage, along with dimension reduction specific to the modalities, are crucial to eliminate the optimistic bias in the evaluation of multimodal IDS.

It is indeed the case that in the transport of Gromov–Wasserstein distortion, which is the measure of the distance between two measures, the greatest reduction is achieved after the alignment when geometry-consistent transport is used. Unlike the naive concatenation strategies which merely combine the different views, GSTF first does the pairwise alignment of the samples according to their relational similarity—rather than raw feature correspondence—which eventually leads to better representation coherence and upstream classification reliability. This phenomenon gets even more pronounced since ablation study outcomes reveal that the GW alignment removal or the discriminative augmentation withdrawal results in the detection robustness being notably lowered.

From the point of view of operational IDS, calibrated confidence estimation is the main practical benefit that comes up. The calibration analysis shows that rather the opposite of it being the case, GSTF generates poorly-aligned confidence scores, which greatly reduces the chance for overconfident false positives and facilitates the setting of decision thresholds for analyst intervention or abstention that are safer. This kind of system is very appropriate for real-world deployments because alert fatigue and trust erosion have already become the persistent challenges.

In spite of these advantages, there are still some limitations that need to be recognized. To begin with, the existing assessment is limited to a large-scale dataset, a binary classification setting, and a future study will still be needed for the extension of the framework to multi-class or cross-dataset scenarios, which remain an essential direction for research. Secondly, even though the proposed pipeline is computationally efficient compared to deep multimodal architectures, Gromov–Wasserstein alignment still requires a significantly non-trivial overhead in the case of very large or streaming data. Finally, even though the interpretability is improved through example-driven explanations, the inclusion of richer causal or temporal explanations would still further strengthen the analyst usability.

To conclude, the findings demonstrate that GSTF provides a methodical compromise between robustness, interpretability, and reliability, besides revealing future to scalability, online adaptation, and generalization across malware families and deployment areas.

5. Conclusion

This study presents a strong, leak-proof multi-view fusion system for malware detection that combines heterogeneous network flow and memory forensics data perfectly. By means of Gromov–Wasserstein alignment, the GSTF pipeline unifies intrinsic geometric discrepancies among different types of data, resulting in a fused representation that remarkably boosts classification performance. The regression-based cross-view mapping and discriminative augmentation lead to an even larger improvement in class separability, while the balanced training method deals with the extreme class imbalance and maintains the leak-proof guarantees intact at the same time.

The large-scale MAL-NetMem 2025 dataset yielded experimental results representing superb accuracy, precision, and total recall, hence always outperforming the basic fusion and classification techniques. The significant decrease in the Gromov–Wasserstein distortion confirms the efficiency of geometric alignment, and the ablation study indicates the significance of each pipeline module. Moreover, the calibrated decision threshold produces a detection that is both reliable and confidence-aware; thus, it is suitable for operational deployment.

Research in the area of intrusion detection systems has advanced significantly, but there are still important gaps which can be addressed in the course of future research. The current IDS do not feature any geometry-aware fusion methods for combining disparate data sources, which are based on principles; they are usually based on data leaking-prone experimental protocols, and their integration of network telemetry and memory-forensics evidence is very rare and mostly done unification-wise. On the other hand, many previous methods do not consider calibration, interpretability, and robustness under class imbalance and distribution shifts at all, which makes it hard for the analysts to trust the system and consequently, limits real-world adoption. However, bringing these issues to a resolution is a matter of precedence since it would lead to the creation of reliable, trustworthy, and operationally viable intrusion detection systems.

The GSTF framework will be extended in future work to include more modalities, such as static and dynamic behavioral traits, which will further enhance the fused representation. Real-time execution and scalability improvements will be considered as a means to allow production cybersecurity environments to absorb the technology. In addition, incorporating sophisticated uncertainty measurement and interpretability methods will boost analysts' trust and ease the path to gaining insights. Finally, adapting malware detection systems to changing malware landscapes through continuous learning and domain change remains a critical area for ensuring sustainability and robustness.

Acknowledgements
Not applicable.

Funding
This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia (Grant No. KFU260220).

Contributions
V.K; V.J; A.A; G.A: Conceptualization, Methodology, Software, Writing – Original Draft, Visualization, Project Administration. V.J: Data Curation, Writing – Review & Editing, Supervision.

Ethics declarations
This article does not contain any studies with human participants or animals performed by any of the authors.

Consent for publication
Not applicable.

Competing interests
The authors have no conflicts of interest to disclose.

References

[1] Achuthan, K., Ramanathan, S., Srinivas, S., & Raman, R. (2024). Advancing cybersecurity and privacy with artificial intelligence: current trends and future research directions. *Frontiers in Big Data*, 7. https://doi.org/10.3389/fdata.2024.1497535

[2] Aldhafferi, N. (2024). Android malware detection using support vector regression for dynamic feature analysis. *Information*, *15*(10), 658. https://doi.org/10.3390/info15100658

[3] Almseidin, M., Alzubi, M., Kovacs, S., & Alkasassbeh, M. (2017). Evaluation of machine learning algorithms for intrusion detection system. *2017 ieee 15th international symposium on intelligent systems and informatics (sisy)*, 000277-000282.

[4] Alohali, M. A., Al-Wesabi, F. N., Hilal, A. M., Goel, S., Gupta, D., & Khanna, A. (2022). Artificial intelligence enabled intrusion detection systems for cognitive cyber-physical systems in industry 4.0 environment. *Cognitive Neurodynamics*, *16*(5), 1045-1057. https://doi.org/10.1007/s11571-022-09780-8

[5] Al-Qudah, M., Ashi, Z., Alnabhan, M., & Abu Al-Haija, Q. (2023). Effective one-class classifier model for memory dump malware detection. *Journal of Sensor and Actuator Networks*, *12*(1), 5. https://doi.org/10.3390/jsan12010005

[6] Alvarez-Melis, D., & Jaakkola, T. (2018). Gromov-wasserstein alignment of word embedding spaces. *Proceedings of the 2018 conference on empirical methods in natural language processing*, 1881-1890.

[7] Amer, E., El-Sappagh, S., Abuhamad, T., Al-Rimy, B. A. S., & Mohasseb, A. (2026). Graphshield: Advanced dynamic graph-based malware detection using graph neural networks. *Expert Systems with Applications*, 298. https://doi.org/10.1016/j.eswa.2025.129812

[8] Arafah, M., Phillips, I., Adnane, A., Alauthman, M., & Aslam, N. (2025). An enhanced bigan architecture for network intrusion detection. *Knowledge-Based Systems*, *314*. https://doi.org/10.1016/j.knosys.2025.113178

[9] Ashwini, K., & Nagasundara, K. B. (2024). An intelligent ransomware attack detection and classification using dual vision transformer with mantis search split attention network. *Computers and Electrical Engineering*, *119*. https://doi.org/10.1016/j.compeleceng.2024.109509

[10] Bensaoud, A., & Kalita, J. (2025). Optimized detection of cyber-attacks on iot networks via hybrid deep learning models. *Ad Hoc Networks*, *170*. https://doi.org/10.1016/j.adhoc.2025.103770

[11] Dehfouli, Y., & Habibi Lashkari, A. (2025). Vadvit: Vision transformer-driven memory forensics for malicious process detection and explainable threat attribution. *Journal of Information Security and Applications*, *94*. https://doi.org/10.1016/j.jisa.2025.104200

[12] Dener, M., Ok, G., & Orman, A. (2022). Malware detection using memory analysis data in big data environment. *Applied Sciences*, *12*(17), 8604. https://doi.org/10.3390/app12178604

[13] Euh, S., Lee, H., Kim, D., & Hwang, D. (2020). Comparative analysis of low-dimensional features and tree-based ensembles for malware detection systems. *IEEE Access*, *8*, 76796-76808. https://doi.org/10.1109/ACCESS.2020.2986014

[14] Habibi Lashkari, A., Li, B., Lucas Carrier, T., & Kaur, G. (2021). Volmemlyzer: Volatile memory analyzer for malware classification using feature engineering. 2021 reconciling data analytics, automation, privacy, and security: *A big data challenge (rdaaps)*, 1-8.

[15] Habibi Lashkari, A., Shafi, M., Li, Y., Singh, A. P., & Barkworth, A. (2025). Unveiling evasive malware behavior: toward generating a multi-sources benchmark dataset and evasive malware behavior profiling using network traffic and memory analysis. *The Journal of Supercomputing*, *81*(6), 782. https://doi.org/10.1007/s11227-025-07267-x Retrieved from https://doi.org/10.1007/s11227-025-07267-x

[16] Hossain, M.A., & Islam, M. S. (2024). Enhanced detection of obfuscated malware in memory dumps: a machine learning approach for advanced cybersecurity. *Cybersecurity*, *7*(1), 16. https://doi.org/10.1186/s42400-024-00205-z

[17] Lee, H., Kim, S., Baek, D., Kim, D., & Hwang, D. (2023). Robust iot malware detection and classification using opcode category features on machine learning. *IEEE Access*, *11*, 18855-18867. https://doi.org/10.1109/ACCESS.2023.3247344

[18] Li, C., Mills, K., Niu, D., Zhu, R., Zhang, H., & Kinawi, H. (2019). Android malware detection based on factorization machine. *IEEE Access*, *7*, 184008-184019. https://doi.org/10.1109/ACCESS.2019.2958927

[19] Li, R., Zhang, Q., & Shen, H. (2025). Malgea: A malware analysis framework via matrix factorization-based node embedding and graph external attention. *Array*, *27*. https://doi.org/10.1016/j.array.2025.100493

[20] Li, Y., Li, Z., & Li, M. (2025). A comprehensive survey on intrusion detection algorithms. *Computers and Electrical Engineering*, *121*. https://doi.org/10.1016/j.compeleceng.2024.109863

[21] Maniriho, P., Mahmood, A. N., & Chowdhury, M. J. M. (2024). Memaldet: A memory analysis-based malware detection framework using deep autoencoders and stacked ensemble under temporal evaluations. *Computers & Security*, *142*, 103864. https://doi.org/10.1016/j.cose.2024.103864

[22] Merlino, V., & Allegra, D. (2024). Energy-based approach for attack detection in iot devices: A survey. *Internet of Things*, *27*. https://doi.org/10.1016/j.iot.2024.101306

[23] Mohammadian, H., Higgins, G., Ansong, S., Razavi-Far, R., & Ghorbani, A. A. (2025). Explainable malware detection through integrated graph reduction and learning techniques. *Big Data Research*, *41*. https://doi.org/10.1016/j.bdr.2025.100555

[24] Muhammed Shafi, K. P., Vinod, P., & Guerra-Manzanares, A. (2025). Hexnet: Enhancing malware classification through hierarchical cnns and multi-level feature attribution. *Journal of Information Security and Applications*, *94*. https://doi.org/10.1016/j.jisa.2025.104207

[25] Nyholm, H., Monteith, K., Lyles, S., Gallegos, M., DeSantis, M., Donaldson, J., & Taylor, C. (2022). The evolution of volatile memory forensics. *Journal of Cybersecurity and Privacy*, *2*(3), 556-572. https://doi.org/10.3390/jcp2030028

[26] Okoli, U.I., Obi, O.C., Adewusi, A.O., & Abrahams, T.O. (2024). Machine learning in cybersecurity: A review of threat detection and defense mechanisms. *World Journal of Advanced Research and Reviews*, *21*(1), 2286-2295. https://doi.org/10.30574/wjarr.2024.21.1.0315

[27] Ottmann, J., Breitinger, F., & Freiling, F. (2024). An experimental assessment of inconsistencies in memory forensics. *ACM Transactions on Privacy and Security*, *27*(1), 1-29. https://doi.org/10.1145/3628600

[28] Pagani, F., Fedorov, O., & Balzarotti, D. (2019). Introducing the temporal dimension to memory forensics. *ACM Transactions on Privacy and Security*, *22*(2), 1-21. https://doi.org/10.1145/3310355

[29] Pourardebil khah, Y., Hosseini Shirvani, M., & Taheri, J. (2026). A survey study on meta-heuristic-based feature selection approaches of intrusion detection systems in distributed networks. *Computer Standards & Interfaces*, *96*. https://doi.org/10.1016/j.csi.2025.104074

[30] Alrajeh, M., Almaiah, M., & Mamodiya, U. (2026). Cyber Risk Analysis and Security Practices in Industrial Manufacturing: Empirical Evidence and Literature Insights. *International Journal of Cybersecurity Engineering and Innovation*, *2026*(1).

[31] Rao, S. X., Han, Z., Yin, H., Jiang, J., Zhang, Z., Zhao, Y., & Shan, Y. (2025). Fraud detection at ebay. *Emerging Markets Review*, *66*. https://doi.org/10.1016/j.ememar.2025.101277

[32] Ho, M., Ang, S., Huy, S., & Janarthanan, M. (2026). MUMSPI: A Model for Usability Measurement of Single-Platform Interface for Multi-Tasking in Big Data Tools. Jordanian Journal of Informatics and Computing, 2026(1), 1–14. https://doi.org/10.63180/jjic.thestap.2026.1.1

[33] Salles, R., Lange, B., Akbarinia, R., Masseglia, F., Ogasawara, E., & Pacitti, E. (2025). Scalable and accurate online multivariate anomaly detection. *Information Systems*, *131*. https://doi.org/10.1016/j.is.2025.102524

[34] Santhanam, P. K., Vellanki, H. C., Bellapu, S. R. R., & Mithra, K. (2026). Q-flexivit: A quantum-flexible vision transformer optimized by octopus-inspired algorithm for intrusion detection. *Computers and Electrical Engineering*, *129*. https://doi.org/10.1016/j.compeleceng.2025.110793

[35] Al-shareeda, M., & Alrudainy, H. (2026). Sustainable and Secure Energy Optimization Strategies in the Internet of Healthcare Things (IoHT). *International Journal of Cybersecurity Engineering and Innovation*, *2026*(1).

[36] Shelke, P., & Hamalainen, T. (2024). Analysing multidimensional strategies for cyber threat detection in security monitoring. *European Conference on Cyber Warfare and Security*, *23*(1), 780-787. https://doi.org/10.34190/eccws.23.1.2123

[37] Singh, S., Krishnan, D., Vazirani, V., Ravi, V., & Alsuhibany, S. A. (2024). Deep hybrid approach with sequential feature extraction and classification for robust malware detection. *Egyptian Informatics Journal*, *27*. https://doi.org/10.1016/j.eij.2024.100539

[38] Ali, A. (2024). Adaptive and Context-Aware Authentication Framework Using Edge AI and Blockchain in Future Vehicular Networks. STAP Journal of Security Risk Management, 2024(1), 45–56. https://doi.org/10.63180/jsrm.thestap.2024.1.3

[39] Souza, C. H., Pascoal, T., Neto, E. P., Sousa, G. B., Filho, F. S., Batista, D. M., & Dantas Silva, F. S. (2025). Sdn based solutions for malware analysis and detection: State-of-the-art, open issues and research challenges. *Journal of Information Security and Applications*, *93*. https://doi.org/10.1016/j.jisa.2025.104145

[40] Alshinwan, M., Memon, A. G., Ghanem, M. C., & Almaayah, M. (2025). Unsupervised text feature selection approach based on improved Prairie dog algorithm for the text clustering. Jordanian Journal of Informatics and Computing, 2025(1), 27–36. https://doi.org/10.63180/jjic.thestap.2025.1.4

[41] Tuan, T. A., Nguyen, P. S., Van, P. N., Hai, N. D., Trung, P. D., Son, N. T. K., & Long, H. V. (2025). A novel framework for cross-platform malware detection via afsp and adasyn-based balancing. *Computers and Electrical Engineering*, *128*. https://doi.org/10.1016/j.compeleceng.2025.110625

[42] Al-shareeda, M., Musa, H. A., Jaafar, A., Salman, A. A., Tami, Z. J., Hameed, H. M., ... & Bashkh, N. S. (2026). Design and Implementation of a Speech-to-Sign Robotic Arm for Deaf Communication. *International Journal of Cybersecurity Engineering and Innovation*, *2026*(1).

[43] Vayer, T., Chapel, L., Flamary, R., Tavenard, R., & Courty, N. (2020). Fused gromovwasserstein distance for structured objects. *Algorithms*, *13*(9), 212. https://doi.org/10.3390/a13090212

[44] Yassin, A., & Almaiah, M. (2026). Cyber security risk assessment for determining threats and countermeasures for banking systems. *International Journal of Cybersecurity Engineering and Innovation*, *2026*(1).

[45] Wang, X., Zhang, J., Zhang, A., & Ren, J. (2019). Tkrd: Trusted kernel rootkit detection for cybersecurity of vms based on machine learning and memory forensic analysis. *Mathematical Biosciences and Engineering*, *16*(4), 2650-2667. https://doi.org/10.3934/mbe.2019132

[46] Wuchner, T., Ochoa, M., & Pretschner, A. (2015). Robust and effective malware detection through quantitative data flow graph metrics. *Detection of intrusions and malware, and vulnerability assessment*. 98–118.

[47] Yumlembam, R., Issac, B., & Jacob, S. M. (2025). Enhancing decision-making in windows pe malware classification during dataset shifts with uncertainty estimation. *Knowledge-Based Systems*, *331*. https://doi.org/10.1016/j.knosys.2025.114723

[48] Zhang, L., Tang, G., He, X., Qi, K., Su, G., & Zhang, H. (2025). Automatic generation of industrial internet attack graphs with graph neural networks and bayesian models. *Computer Networks*, *272*. https://doi.org/10.1016/j.comnet.2025.111736

[49] Zhang, S., Shan, S., Hu, Z., Shen, Y., Li, C., Zhang, K., & Wei, H. (2025). Outof-distribution fault detection in multi-sensor systems using spatio-temporal dynamic graph neural networks. *Mechanical Systems and Signal Processing*, *241*. https://doi.org/10.1016/j.ymssp.2025.113524

[50] Alsahaim, S., Almaiah, M. A., & Sulaiman, R. B. (2023). Security Threats in Mobile Phones: Challenges, Countermeasures, and the Importance of User Awareness. *International Journal of Cybersecurity Engineering and Innovation*, *2023*(1).

[51] Abu Laila, D., Aljawarneh, M., Al-Na'amneh, Q., & Bin Sulaiman, R. (2025). Optimizing Intrusion Detection Systems through Benchmarking of Ensemble Classifiers on Diverse Network Attacks. STAP Journal of Security Risk Management, 2025(1), 71–84. https://doi.org/10.63180/jsrm.thestap.2025.1.4

Biographies

Vijay H. Kalmani is a Professor in the Department of Computer Science and Engineering at Rajarambapu Institute of Technology, Sangli, India. He holds an M.Tech in CSE from Gogte Institute of Technology, Belagavi, and a Ph.D. in CSE from Suresh Gyan Vihar University, Jaipur. His research interests include Cloud Security, Artificial Intelligence, and Machine Learning. He has successfully guided three Ph.D. scholars. Contact: vijaykalmani@gmail.com

Vedant V. Jadhav is pursuing his Master's in Computer Science and Engineering at Rajarambapu Institute of Technology, Sangli, India. His research focuses on Malware Detection, Intrusion Detection Systems, Memory Forensics, Network Security, and machine-learning-driven cybersecurity. He actively contributes to research on multi-view malware detection frameworks and AI-based threat analytics. Contact: iamvedant07@gmail.com

Amer Al-Qutaish, Strategic Planning and Institutional Identity Administration, King Faisal University, Al-Ahsaa 31982, Saudi Arabia, aalqutish@kfu.edu.sa

Ghada Alradwan, Strategic Planning and Institutional Identity Administration, King Faisal University, Al-Ahsaa 31982, Saudi Arabia.