ASYMGAD: LABEL-FREE LATERAL-MOVEMENT PIVOT RANKING VIA ASYMMETRIC GRAPH ANOMALY DETECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Enterprise authentication activities naturally form graph-structured data that capture how machines authenticate to one another and to administrative services. Correspondingly, graph anomaly detection (GAD) provides a promising paradigm for identifying compromised machines involved in lateral movement, where attackers expand access by reusing legitimate credentials and services. However, detecting the intermediate stepping stones, or pivots, remains challenging because malicious authentication behavior often resembles routine activity, while the underlying interaction patterns continuously evolve over time. To address these challenges, we propose ASYMGAD, a label-free GAD framework for ranking suspicious pivot machines in continuous authentication streams. ASYMGAD first constructs evidence-sufficient observation graphs directly from the event stream, allowing graph boundaries to adapt to changing activity rather than relying on externally specified snapshots. It then introduces directional expansion as an explicit structural prior to capture machines whose outgoing behavior is disproportionately strong relative to their incoming activity, providing an interpretable signal of pivot behavior. Furthermore, ASYMGAD incorporates causal information across consecutive observation graphs to distinguish persistent benign infrastructure from transient suspicious behavior, while a lightweight directed graph neural network captures complementary relational patterns beyond node-level statistics. By integrating adaptive graph construction, asymmetric structural cues, temporal context, and learned graph representations, ASYMGAD enables label-free and interpretable pivot prioritization in large-scale authentication networks. Extensive experiments on a billion-event enterprise authentication stream demonstrate the effectiveness of ASYMGAD for identifying compromised pivot machines at enterprise scale.

1 INTRODUCTION

Lateral movement is the stage of an enterprise compromise in which an attacker, after breaching one machine, reuses legitimate credentials and standard administrative tools to move between machines and reach high-value assets. It is also a central design target of modern zero-trust architectures, which verify every access request and restrict movement between machines Rose et al. (2020); lateral movement nevertheless occurs because the same credentials and tools that support routine administration can be reused by the attacker Ho et al. (2021). The authentication records left by such movement therefore closely resemble routine activity, and identifying the machines an attacker uses as stepping stones—which we call pivots—becomes the core of the detection problem.

Automated detection with static rules is difficult, because attackers reuse legitimate tools and the stream contains billions of events, far too many to inspect by hand. Supervised classifiers face the scarcity, incompleteness and skew of malicious labels Sommer & Paxson (2010), and extreme class imbalance makes event-level accuracy an unreliable objective He & Garcia (2009); Davis & Goadrich (2006). The problem is better formulated as ranking candidate machines, which can be evaluated with Average Precision. Ranking requires a representation of machine behavior that reflects the relational structure of authentication activity Kent et al. (2015).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Graphs provide such a representation, and relational models have shown their value for anomaly detection, as temporal graph networks for anomalous link prediction King & Huang (2023) and graph foundation models for cyber telemetry Qiao et al. (2025a); King et al. (2026) demonstrate. The first limitation is an assumption they share: that the event stream has already been divided into graph snapshots and that the boundaries are given in advance. The assumption breaks down when the event rate varies over time, since a fixed boundary changes the node population, edge density and neighborhood composition of the resulting graph. Windows derived from known attack timelines are especially problematic, because they preselect the relevant period and can overstate detection performance Flood et al. (2024); Bouke & Abdullah (2023); Hutzelmann et al. (2024); Bilot et al. (2025).

The second limitation is that these models reason only within one snapshot. Pivots also differ in the behaviors that expose them: some become visible through directional fanout, which arises when an attacker reuses harvested credentials to probe many destinations He et al. (2024a), while others keep the connection volume low and force a workstation into a role that differs from the role predicted for comparable machines He et al. (2024b). Benign infrastructure produces the first of these patterns as well: a patch server or a backup job repeats a comparable fanout across windows, and a snapshot-local model cannot see that regularity.

We present AsymGAD, a label-free framework that addresses both. The stream is divided by a label-free adaptive window construction (AWC), which closes a window only once explicit evidence thresholds are met, so every observation graph has enough events and enough distinct sources and destinations for downstream learning. Within each graph, a structural asymmetry score measures directional expansion, and a self-supervised structural branch learns, without labels, to estimate where a host stands in the persistence ordering. The two are combined with a weight the window itself determines. A causal cross-window persistence memory then records, for every host, the fraction of past windows in which it ranked at the top of the asymmetry order; hosts at the extreme of that statistic are removed from the candidate list before the final ordering is produced.

In summary, this work makes four contributions.

1. We identify observation-unit boundary selection as a source of implicit bias in stream-based graph anomaly detection, and propose AWC, a label-free, evidence-sufficient construction that forms part of the detector.
2. We introduce a self-supervised structural branch trained end to end on a label-free target, and show by ablation that its gain comes from message passing over the observation graph.
3. We turn the observation that persistently expansion-dominant hosts are benign infrastructure into an explicit causal filter whose constants are fixed before the stream is read, and compute the fusion weight from the window itself.
4. We evaluate AsymGAD on a billion-event enterprise authentication stream against seven public baselines under one shared protocol, with paired per-window statistical testing.

2 Related Work

2.1 Lateral Movement Detection in Enterprise Telemetry

Lateral movement detection has been studied at the event, path and machine levels Smiliotopoulos et al. (2024). Signature- and rule-based systems are inexpensive when attack patterns are known, but attackers evade them by reusing legitimate credentials and tools already present on the host Uetz et al. (2024). Anomaly detectors instead model normal enterprise behavior, although their performance depends strongly on the representation Sommer & Paxson (2010). Graph-based methods reconstruct attack paths or identify unusual traversals from the relations in authentication streams Kent et al. (2015); Ho et al. (2021). Much of this work focuses on retrospective path reconstruction. Ranking machines for immediate analyst review is a distinct operational need Soliman et al. (2024); Liu et al. (2026); Jiang et al. (2025).

Temporal graph models and cyber-oriented foundation models extend this line of work to evolving links. EULER King & Huang (2023) models network-state transitions to detect anomalous links, and graph foundation models such as CybGFM King et al. (2026) learn representations for link-level

scoring. These systems generally produce edge-level or path-level outputs and operate on snapshots defined beforehand. Pivot ranking addresses a different operational need by aggregating relational evidence at active source machines and returning a short queue of candidates.

2.2 Unsupervised Graph Anomaly Detection

Unsupervised graph anomaly detection identifies structural or attributed deviations without ground-truth labels Ma et al. (2023); Qiao et al. (2025b). Reconstruction-based methods assume that common patterns are easier to reconstruct than anomalies, with representative models DOMINANT Ding et al. (2019), AnomalyDAE Fan et al. (2020) and GAD-NR Roy et al. (2024). Contrastive methods learn normality from agreement between nodes and local subgraphs, such as CoLA Liu et al. (2022) and SL-GAD Zheng et al. (2023). One-class methods, such as TAM Qiao & Pang (2023), maximize affinity under a homophily assumption. Other methods model uncertainty in reconstruction error (GEL Wei et al. (2025)), align prototypes across graphs for transfer (AnomalyGFM Qiao et al. (2025a)), or learn masked graph representations for self-supervised APT detection (MAGIC Jia et al. (2024)).

Two limitations matter for enterprise telemetry. These methods are typically evaluated on static or preconstructed benchmark graphs, so their behavior under a different observation unit is rarely reported Jung et al. (2024). The one-class family in particular assumes undirected or homophilous structure Qiao & Pang (2023), which does not match the asymmetric relations in authentication data.

2.3 Streaming Segmentation and Directional Graph Learning

Streaming telemetry is commonly divided into graphs by fixed wall-clock intervals, session timeouts or fixed event counts Yang et al. (2023); Zhang et al. (2025); Wang et al. (2024). Time-based and count-based boundaries behave differently. A wall-clock or session boundary follows the arrival rate, so the volume per graph varies with it; a fixed event count holds the volume and leaves the relational coverage open. Neither guarantees that a graph contains enough distinct sources and destinations for the analysis it is used for. Existing authentication-graph systems usually treat segmentation as an external parameter Kent et al. (2015); King & Huang (2023), even though it changes the difficulty of the detection task.

An event $u \to v$ carries a different operational meaning from $v \to u$, and directed graph architectures model this orientation through complex Laplacians Zhang et al. (2021). These architectures do not directly address how to rank source machines in a dynamic stream. AsymGAD combines evidence-based window construction with directed structural learning for this purpose.

3 Problem Formulation

We observe a chronological stream of authentication events $\mathcal{D} = \{e_1, e_2, \ldots\}$, where each event $e_i = (t_i, s_i, d_i, \mathbf{a}_i)$ carries a timestamp, a source entity, a destination entity, and a vector of categorical service and protocol attributes. The defender has no payload signatures, no incident timeline, and no labeled training examples. Incident labels are used only after the ranking has been produced, to compute evaluation metrics.

The task is to produce, for each observation unit, a ranking of machines by their likelihood of being a pivot.

The choice of observation unit determines what evidence the detector sees. A detector running on a continuous stream without incident labels must therefore construct its own observation units, and the windowing step becomes part of the problem. Appendix A sets out why boundaries taken from a known incident timeline overstate detection performance Flood et al. (2024); Bouke & Abdullah (2023); Hutzelmann et al. (2024); Bilot et al. (2025).

A pivot initiates at least one observed event. We therefore call $\mathcal{C}_k = \{v \in \mathcal{V}_k : d_k^{\text{out}}(v) > 0\}$ the active-source candidate set. To keep the ranked universe identical for every method, the reported evaluation orders every node present in $\mathcal{G}_k$. Nodes with zero out-degree sit below the active sources under $A(v)$, and every known pivot is itself an active source.

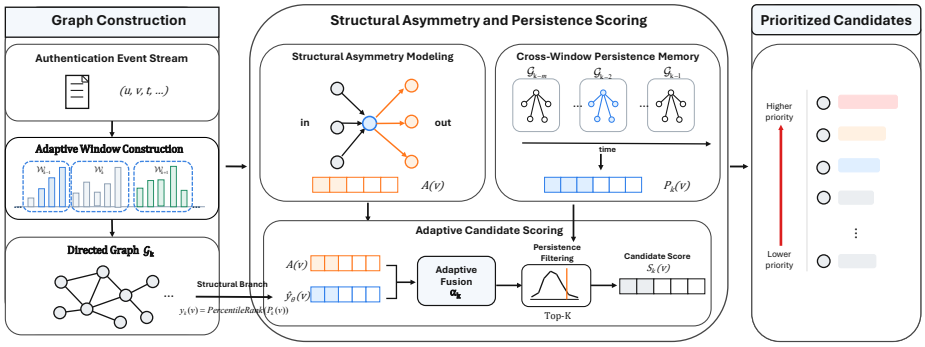


Figure 1: Overview of AsymGAD. AWC converts the variable-rate stream into observation graphs that meet evidence thresholds. Within each graph, a structural asymmetry score and a self-supervised structural branch are computed, and a causal cross-window persistence memory determines which hosts are persistent infrastructure. The asymmetry score and the branch are fused with a weight derived from the window itself, and hosts in the persistence core are removed before the final ordering.

The detector returns $\pi_k = \operatorname{argsort}^{\downarrow}_{v \in \mathcal{V}_k} S_k(v)$, where $S_k(v)$ is a relative anomaly score. For a positive window, the ranking is evaluated with Average Precision, $\mathrm{AP}_k = \sum_n (R_n - R_{n-1}) P_n$, where P_n and R_n are the precision and recall after the nth ranked candidate Davis & Goadrich (2006). AP rewards methods that place the known pivots near the top of the list. Detailed notation is given in Appendix A.

4 The AsymGAD Framework

This section describes AsymGAD, which ranks lateral-movement pivot candidates directly from a raw authentication stream without labels. Figure 1 summarizes the pipeline.

4.1 Design Rationale

Lateral movement does not have a single topological signature, and the two patterns that expose it are different in kind. The first is *directional expansion*: the pivot contacts many destinations over the window while receiving few incoming authentications, so its outgoing behavior is disproportionate to its incoming activity. We score this pattern as structural asymmetry $A(v)$ He et al. (2024a). Benign infrastructure produces a very similar fanout within a single window, because patch servers, backup jobs and scanners also contact many machines and receive little; what separates the two is that the infrastructure host repeats that fanout across windows, which is the connectivity profile of the infrastructure population. A snapshot-local score therefore cannot separate them, which motivates the cross-window persistence memory.

The second pattern is a low-volume intrusion: the attacker keeps its connection count unremarkable and moves through machines that take on roles inconsistent with comparable hosts He et al. (2024b). Expansion alone cannot rank this pattern, so the framework carries a second signal as well: a learned structural branch that estimates the persistence ordering from the structure of the current graph.

4.2 Adaptive Window Construction (AWC)

AWC uses a dynamic boundary rule so that each graph contains enough events and enough distinct sources and destinations for downstream learning. For a candidate window $\mathcal{W}$, let $M(\mathcal{W})$ be the number of raw events, $E(\mathcal{W})$ the number of unique directed pairs, and $V_{\mathrm{src}}(\mathcal{W})$, $V_{\mathrm{dst}}(\mathcal{W})$ the numbers of distinct sources and destinations. The window closes only when all four conditions hold:

$$M(\mathcal{W}) \geq M_{\min}, \qquad E(\mathcal{W}) \geq E_{\min},$$
$$V_{\mathrm{src}}(\mathcal{W}) \geq V_{\mathrm{src,min}}, \qquad V_{\mathrm{dst}}(\mathcal{W}) \geq V_{\mathrm{dst,min}}.$$

The threshold $M_{\min}$ controls total observation volume, $E_{\min}$ prevents repeated events among a small set of machines from satisfying the rule, and the two vertex thresholds require participation from enough sources and destinations. The rule reads only event counts and set cardinalities, so it is

label-free and runs online: Algorithm 1 (Appendix C.1) applies it sequentially and closes a window at the first event that satisfies all four conditions. A high-rate period reaches the thresholds quickly, whereas a quiet period remains open until enough relations have accumulated. Any residual events at the end of the stream are merged into the last complete window, so no undersized graph is emitted.

4.3 Directed Attributed Graph Construction

Each window $\mathcal{W}_k$ is mapped to a directed attributed graph $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k, \mathbf{X}_k, \mathbf{F}_k)$. All authentication events with the same ordered pair (u, v) are aggregated into one edge $u \to v \in \mathcal{E}_k$, with multiplicity $c_{uv} = |\{e_i \in \mathcal{W}_k : s_i = u,\, d_i = v\}|$.

The node feature matrix $\mathbf{X}_k$ holds two per-node features, the log in-degree and the log out-degree. Each edge $u \to v$ carries a vector $\mathbf{f}_{uv} \in \mathbb{R}^{37}$ recording the authentication mechanism and protocol, the service ports observed for the same machine pair, and the log interaction count (Appendix C.1). The reported structural branch uses the directed edge set and the two log-degree node features; no edge attribute is an input to the learned score.

4.4 Self-Supervised Structural Branch

The structural branch is a single directed message-passing layer whose incoming and outgoing aggregations are kept separate, since the quantity of interest is what a machine *initiates*:

$$\mathbf{m}_v^{\text{in}} = \frac{1}{|\mathcal{N}^{\text{in}}(v)|} \sum_{u \to v} \mathbf{W}_{\text{in}} \mathbf{h}_u,$$

$$\mathbf{m}_v^{\text{out}} = \frac{1}{|\mathcal{N}^{\text{out}}(v)|} \sum_{v \to w} \mathbf{W}_{\text{out}} \mathbf{h}_w,$$

$$\mathbf{h}'_v = \text{LayerNorm}\big(\mathbf{h}_v + \text{MLP}[\mathbf{h}_v \,\|\, \mathbf{m}_v^{\text{in}} \,\|\, \mathbf{m}_v^{\text{out}}]\big).$$

The branch is trained end to end on a *label-free* target. Let $\mathcal{H}_k = \{v \in \mathcal{V}_k : N_k(v) \geq 3\}$ be the hosts for which the persistence statistic (§4.5) is defined, and let the target be $y_k(v) = \text{PercentileRank}\big(\{P_k(u) : u \in \mathcal{H}_k\}, P_k(v)\big)$. The branch minimizes $\mathcal{L} = \frac{1}{|\mathcal{H}_k|} \sum_{v \in \mathcal{H}_k} \big(\hat{y}_\theta(v) - y_k(v)\big)^2$. The target is the percentile because $P_k(v)$ is concentrated at zero: hosts that have never reached the top of the asymmetry order sit at the bottom of the scale, while a small set of infrastructure hosts occupies its upper tail. The percentile therefore grades a host by whether it has an established record of top-of-A activity. The two ends of its range — no record at all, and a record in every window — are the cases the filter of §4.6 separates; the branch estimates the position between them from the current window's structure alone.

Both the input graph and the target are computed from the current window and strictly earlier windows, so no label and no future information enters training. At a refit, the branch is trained from scratch on the sixty most recent observation windows of the stream; those fitted parameters are then used for subsequent windows until the next refit, which follows on a fixed cadence of eight observation windows. Neither the buffer nor the cadence refers to which windows were later found to contain an attack.

The target makes the branch persistence-targeted. The persistence percentile cannot be read off a single window, and the directed one-hop neighborhood provides a relational representation of the host's current structural context. The architecture, refit cadence and buffer are specified in Appendix C.2.

4.5 Cross-Window Persistence Memory

Every window k has a within-window asymmetry score $A(v) = \text{PercentileRank}\big(\{A_{\text{raw}}(u) : u \in \mathcal{C}_k\}, A_{\text{raw}}(v)\big)$ with $A_{\text{raw}}(v) = d^{\text{out}}(v)/(d^{\text{in}}(v) + 1)$.

The memory records, for every host v, the number of past windows in which it appeared, $N_k(v)$, and the number of those in which it reached the top K of the asymmetry order, $\kappa_k(v)$, and defines

$$P_k(v) = \kappa_k(v)/N_k(v). \tag{1}$$

We call $\mathcal{T}_k$ the *extreme-asymmetry set* of window k: the K hosts with the largest expansion ratio A_{raw} in that window, with $K = 50$. $N_k(v)$ counts the windows before k in which v was observed, and $\kappa_k(v)$ those in which it entered $\mathcal{T}_k$; a host is *sufficiently observed* once $N_k(v) \geq 3$, and only then is $P_k(v)$ defined.

The memory is updated after every window, whether or not that window contains a recorded attack, and it never reads incident labels. $P_k(v)$ is a history of expansion: it records how often a host has entered the extreme-asymmetry set. The two ends of its range mean different things: a host with no record has never stood out, and a host that enters the set in every window it appears in is the profile of a patch server, a backup job or a scanner, whose A_{raw} is what §4.6 removes. The recorded pivots sit between the two ends: their persistence averages 0.31 and never exceeds 0.56, below the threshold the filter applies, so among the candidates that remain, an established record is the informative case. The statistic is used at two scales for that reason. A_{raw} identifies the core to exclude, and its within-window percentile P_k, read on the surviving candidates, supplies the regression target of §4.4.

4.6 Candidate Scoring and Adaptive Fusion

The final score fuses the two terms and removes the persistence core. Both fused terms are oriented in the same direction, a larger value of either marking a stronger pivot signal; the benign core is excluded by the filter:

$$S_k(v) = \begin{cases} -\infty, & P_k(v) \geq \Theta, \\ \alpha_k \, q(A(v)) + (1 - \alpha_k) \, q(\hat{y}(v)), & \text{otherwise}, \end{cases} \tag{2}$$

where $q(\cdot)$ denotes the within-window percentile rank, taken over all hosts of the window, so that both terms are on a common scale. The two terms differ in what they measure: $A(v)$ is the window's own directional expansion, whereas $\hat{y}(v)$ is the branch's estimate, from the current structure, of where the host stands in the persistence ordering.

A host is treated as persistently expansion-dominant, and removed, once it has entered the extreme-asymmetry set in at least two thirds of its sufficiently observed appearances, $\Theta = 2/3$. The rule is a threshold on the statistic: the number of hosts it excludes varies with the window, and the rule itself is fixed before the stream is read. Appendix E gives the label-free quantities that set Θ, K and the fusion clamp, and when each is fixed.

The weight α_k is derived from the window itself, because how informative the asymmetry order is varies from window to window:

$$\alpha_k = \text{clip}\left(\tfrac{1}{3}\log(1 + s_k),\ \alpha_{\text{lo}},\ \alpha_{\text{hi}}\right), \tag{3}$$

$$s_k = \frac{\text{mean}\left(\text{top } 5\% \text{ of } A_{\text{raw}}\right)}{\text{median}\left(A_{\text{raw}}\right)}. \tag{4}$$

When the top of the asymmetry distribution is sharply separated from its bulk, the asymmetry score is carrying a clear signal and keeps the larger share of the weight; when it is not, the branch receives a larger share of it. The statistic is label-free and is computed from the current window alone.

Sorting candidates by decreasing $S_k(v)$ gives the ranking π_k. The complete per-window procedure, including the two ordering constraints that keep the label-free target causal, is given in Algorithm 2 (Appendix C.1).

5 Experiments

We first describe the setup (§5.1), then report the end-to-end ranking benchmark (§5.2), the per-pivot breakdown (§5.3), the paired per-window tests (§5.4) and the component ablation (§5.5). Supporting detail is carried by the appendix, and the analyses that are only summarized here are listed at the end of this section.

5.1 Setup

Dataset and segmentation. We use the LANL enterprise authentication dataset Kent (2015), which contains 1,051,430,459 events in chronological order across 211 source files. AWC partitions

Table 1: Benchmark results on 54 positive adaptive windows. Strong, Medium and Weak denote AP ≥ 0.01, $0.001 \leq$ AP < 0.01 and AP < 0.001. A unique win is the highest AP in a window without a tie. The ASYMGAD row is a single run of the structural branch; Table 18 gives all three. Per-window values behind this table are in Table 7.

Method	Mean AP	Median AP	Best AP	Worst AP	Strong	Medium	Weak	Unique Wins
CybGFM	0.0299	0.0072	1.0000	0.0002	12	35	7	13
AnomalyGFM	0.0107	0.0049	0.0526	0.0001	17	28	9	8
CoLA	0.0033	0.0010	0.0227	0.0001	5	21	28	3
GEL	0.0024	0.0004	0.0110	0.0001	2	18	34	0
EULER	0.0023	0.0004	0.0270	0.0001	3	15	36	2
TAM	0.0003	0.0001	0.0012	0.0001	0	2	52	0
GAD-NR	0.0003	0.0002	0.0027	0.0001	0	2	52	0
A alone (unlearned)	0.0858	0.0070	0.3333	0.0001	25	4	25	0
ASYMGAD (Ours)	**0.2724**	0.0082	**1.0000**	**0.0001**	**27**	5	22	**28**

the stream into 525 non-overlapping observation windows, a rule Appendix D.1 compares against fixed boundaries. Incident records identify 54 windows with at least one active malicious pivot; the other 471 contain background traffic with no recorded attack. AP is defined only when the candidate set contains a positive instance, so the ranking comparison uses the 54 positive windows. The four known pivots are `C17693` (active in 44 windows), `C19932` (14), `C22409` (3) and `C18025` (1). Table 3 and Table 4 give the stream and window statistics, and Table 5 gives the pivot archetypes (Appendix B.1).

Baselines. We compare ASYMGAD with seven public baselines—GAD-NR Roy et al. (2024), GEL Wei et al. (2025), TAM Qiao & Pang (2023), AnomalyGFM Qiao et al. (2025a), CoLA Liu et al. (2022), EULER King & Huang (2023) and CybGFM King et al. (2026)—described in Appendix B.2. We also report A alone, the unlearned asymmetry score $A(v)$, as an internal diagnostic that isolates the gain from the two learned components. Every method is evaluated on the same 54 adaptive-window graphs, candidate sets, labels and ranking code.

Metrics and statistical protocol. For each positive window a method returns a ranking π_k, evaluated with average precision; the main metric is the macro mean over the 54 windows. Windows are grouped as Strong (AP ≥ 0.01), Medium ($0.001 \leq$ AP < 0.01) and Weak (AP < 0.001), and a method receives a unique win when it has the highest AP in a window without a tie. For each baseline b we report the paired difference $\Delta_{k,b} = \mathrm{AP}_{k,\mathrm{ASYMGAD}} - \mathrm{AP}_{k,b}$ with a 95% percentile bootstrap interval (20,000 window resamples) and a one-sided Wilcoxon signed-rank test with Holm-Bonferroni correction across baselines. The full paired table is Table 11.

Implementation. The structural branch has one message-passing layer with hidden dimension 64, trained with Adam at learning rate 3×10^{-3} for 25 epochs per refit, on a buffer of the sixty most recent observation windows, and is refit every eight observation windows. The evidence thresholds are $M_{\min} = 2{,}000{,}000$, $E_{\min} = 50{,}000$, $V_{\mathrm{src},\min} = 5{,}000$ and $V_{\mathrm{dst},\min} = 3{,}000$; ASYMGAD's remaining constants are listed in Appendix C.2. Incident labels take no part in any of these choices: they are excluded from window construction, graph construction, feature extraction, the regression target, the persistence memory, the refit schedule, the training buffer, the fusion weight and the candidate filter. Every constant above is fixed before the stream is read, and labels are read only after the rankings have been produced, to compute the metrics and to report what the filter removed.

5.2 MAIN BENCHMARK

Table 1 compares ASYMGAD with the seven public baselines and with the unlearned score. ASYMGAD attains the highest AP in 28 of the 54 windows, all as unique wins; no other method does so in more than 13 (CybGFM), and A alone never does. The aggregate follows: 0.2724 mean AP on the reported run, the one behind the per-window tables and figures, and 0.2616 ± 0.0111 over three runs of the structural branch (Table 18), against a range of 0.0003 to 0.0299 for the seven baselines and 0.0858 for the unlearned score. The paired tests separate it from every baseline; the smallest margin is CybGFM, and §5.4 reports the intervals. In analyst-budget terms (Table 8), ASYMGAD places a known pivot among the top 10, 50 and 100 candidates in 23, 25 and 27 of the 54 windows,

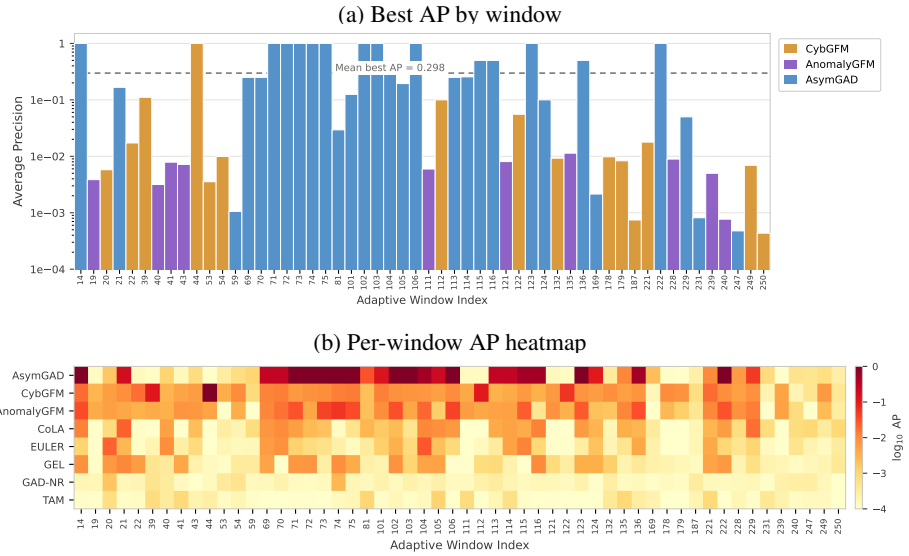


Figure 2: (a) Best AP achieved by ASYMGAD, CybGFM and AnomalyGFM in each attack window, with the bar color marking the method that attains it. ASYMGAD attains the best window-level AP in 29 of 54 windows (53.7%); the dashed line marks the mean best AP of 0.298. (b) Per-window Average Precision of ASYMGAD and the seven public baselines on a $\log_{10}$ scale.

respectively, against 18, 24 and 25 for A alone, and its median rank for the best-ranked pivot falls from 142 to 130; Table 18 recomputes these counts for each run.

Figure 2 shows the best AP attained in each of the 54 attack windows among ASYMGAD, CybGFM and AnomalyGFM. ASYMGAD attains the best window-level AP in 29 windows (53.7%), compared with 15 for CybGFM and 10 for AnomalyGFM in this three-method view; the 28 unique wins in Table 1 use the stricter all-baseline comparison.

5.3 PER-PIVOT RANKING PERFORMANCE

Windows containing `C17693` yield the highest window-level AP for every method; ASYMGAD reaches 0.3336 there, against 0.0365 for CybGFM, and holds the best candidate rank of any method (median 9, best 1). The low-volume pivot `C19932` remains the most difficult: ASYMGAD leads at 0.0324 with a median rank of 1,526, against 0.0093 and 1,794 for A alone. `C18025` appears in a single window and is not recovered by any method, though ASYMGAD improves its rank from 8,518 to 7,927 and more than doubles its AP relative to the unlearned score. Table 9 gives the full breakdown and Table 10 the corresponding candidate ranks.

5.4 PAIRED PER-WINDOW STATISTICAL EVALUATION

Across the 54 positive windows, every bootstrap interval excludes zero and every Holm-adjusted p-value is below 3×10^{-3}, in all three runs of the branch. On the reported run, the mean paired difference relative to A alone is $+0.1866$ (95% CI $[+0.1150, +0.2648]$; Holm $p = 7.6 \times 10^{-9}$). Relative to CybGFM and AnomalyGFM, the differences are $+0.2425$ ($[+0.1309, +0.3577]$; $p = 2.82 \times 10^{-3}$) and $+0.2617$ ($[+0.1634, +0.3668]$; $p = 2.77 \times 10^{-3}$), respectively. Table 11 reports the full table, including the baselines whose intervals are wider.

5.5 COMPONENT ABLATION

Table 2 separates the components. A count-matched random deletion of the same number of hosts leaves the unlearned score at 0.0862, so the gain of the persistence filter is attributable to which hosts it removes (Appendix D.7). The filter takes the unlearned score from 0.0858 to 0.2402, and the structural branch adds a further 0.0322 on top of it (Appendix D.5): the branch ranks candidates that the filter has not removed, and the two carry different information, because the memory reads

Table 2: Component ablation, mean Average Precision over the 54 positive windows. The rows separate the unlearned score, the learned branch, the persistence filter and the full method. The random control removes as many candidates as the filter does in each window but chooses them at random, which isolates the filter's gain from the number of removals. The full sweep is in Table 12.

Configuration	Mean AP
A alone (unlearned)	0.0858
R alone (learned branch)	0.1294
$A + P$ filter	0.2402
ASYMGAD (adaptive α_k)	0.2724
Random deletion (count-matched)	0.0862

the history of the asymmetry order while the branch reads the structure of the current graph. Scored on its own, the branch reaches 0.1294 in the reported run and 0.1069 ± 0.0220 across the three runs, against 0.0858 for the unlearned score (Table 18). The fusion weight is not chosen by hand. Eq. 3 computes it per window from a label-free statistic of that window, and the pipeline lands within 0.001 of the highest value in the hand-set sweep (Appendix D.6).

Additional Analyses. The remaining analyses are in the appendix. The full per-pivot and paired tables are in Appendix D.3, together with the signed rank-change figure; the full fusion-weight sweep, including the mechanism of Eq. 3 and the alternative fusion rules we tried, is in Appendix D.6; the persistence filter is isolated against a count-matched random-deletion control in Appendix D.7; and the behavior of the unfiltered asymmetry order on the 471 attack-unlabeled windows is in Appendix D.8. A full per-window benchmark table is provided in Appendix D.2 so that every aggregate statistic reported here can be checked window by window.

6 CONCLUSION

We presented ASYMGAD, which combines the adaptive window construction, a within-window asymmetry score, a causal cross-window persistence memory, and a self-supervised structural branch. On LANL, with a label-free pipeline whose constants are all fixed before the stream is read, ASYMGAD leads every public baseline in all three runs of its structural branch, and attains the highest Average Precision in 28 of the 54 attack windows, all as unique wins.

AI USE STATEMENT

We employed large language models (LLMs) to assist in drafting and refining this manuscript; specifically, they helped optimize language expression, improve readability and clarity, restructure sentences, check grammar, and enhance overall fluency. LLMs were not involved in conceptualization, research design, or experimental planning; all research concepts, ideas, and analyses were independently developed and implemented by the authors. Their contribution was limited to enhancing the manuscript's linguistic quality and did not involve scientific content or data analysis. The authors take full responsibility for the integrity of the manuscript.

ETHICS STATEMENT

We adhere to the ICLR Code of Ethics and have taken every measure to ensure that our research complies with the ethical guidelines set forth by ICLR. Our work does not involve human subjects, nor does it raise any ethical concerns related to privacy or data usage. The dataset used in our experiments is publicly released for research use, and we ensure that its release and use adhere to proper data-sharing policies. We have carefully selected the evaluation metrics to ensure a fair and transparent comparison in our findings. No potential conflicts of interest, sponsorship biases, or ethical violations have been identified in our study. We commit to maintaining the highest standards of research integrity, and we are open to addressing any ethical concerns that may arise during the review process.

REPRODUCIBILITY STATEMENT

The stream is the public LANL enterprise authentication dataset Kent (2015). The four evidence thresholds, the branch hyperparameters, the persistence settings and the fusion clamp are listed in Appendix C.2, and the algorithms are given in Appendix C.1. The paired bootstrap uses 20,000 window resamples, and the baselines are scored from a frozen per-window record produced by the same segmentation, candidate sets and ranking code as ASYMGAD. Incident labels only after the rankings have been produced.

REFERENCES

Tristan Bilot, Baoxiang Jiang, Zefeng Li, Nour El Madhoun, Khaldoun Al Agha, Anis Zouaoui, and Thomas Pasquier. Sometimes simpler is better: A comprehensive analysis of state-of-the-art provenance-based intrusion detection systems. In *Proceedings of the 34th USENIX Security Symposium*. USENIX Association, 2025.

Mohamed Aly Bouke and Azizol Abdullah. An empirical study of pattern leakage impact during data preprocessing on machine learning-based intrusion detection models reliability. *Expert Systems with Applications*, 230:120715, 2023. doi: 10.1016/j.eswa.2023.120715.

Jesse Davis and Mark Goadrich. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd International Conference on Machine Learning*, pp. 233–240, 2006. doi: 10.1145/1143844.1143874.

Kaize Ding, Jundong Li, Rohit Bhanushali, and Huan Liu. Deep anomaly detection on attributed networks. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pp. 594–602, 2019. doi: 10.1137/1.9781611975673.67.

Haoyi Fan, Fengbin Zhang, and Zuoyong Li. Anomalydae: Dual autoencoder for anomaly detection on attributed networks. In *ICASSP 2020*, pp. 5685–5689, 2020. doi: 10.1109/ICASSP40776.2020.9053387.

Robert Flood, Gints Engelen, David Aspinall, and Lieven Desmet. Bad design smells in benchmark NIDS datasets. In *2024 IEEE 9th European Symposium on Security and Privacy*, pp. 658–675, 2024. doi: 10.1109/EuroSP60621.2024.00042.

Daojing He, Hongjie Gu, Shanshan Zhu, Sammy Chan, and Mohsen Guizani. A comprehensive detection method for the lateral movement stage of APT attacks. *IEEE Internet of Things Journal*, 11(5):8440–8447, 2024a. doi: 10.1109/JIOT.2023.3322412.

Daojing He, Xin Lv, Xueqian Xu, Sammy Chan, and Kim-Kwang Raymond Choo. Double-layer detection of internal threat in enterprise systems based on deep learning. *IEEE Transactions on Information Forensics and Security*, 19:4741–4751, 2024b. doi: 10.1109/TIFS.2024.3372771.

Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. doi: 10.1109/TKDE.2008.239.

Grant Ho, Mayank Dhiman, Devdatta Akhawe, Vern Paxson, Stefan Savage, Geoffrey M. Voelker, and David Wagner. Hopper: Modeling and detecting lateral movement. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 3093–3110. USENIX Association, 2021.

Thomas Hutzelmann, Dominik Mauksch, Ana Petrovska, and Alexander Pretschner. Generation of tailored and confined datasets for IDS evaluation in cyber-physical systems. *IEEE Transactions on Dependable and Secure Computing*, 21(4):3948–3962, 2024. doi: 10.1109/TDSC.2023.3341211.

Zian Jia, Yun Xiong, Yuhong Nan, Yao Zhang, Jinjing Zhao, and Mi Wen. MAGIC: Detecting advanced persistent threats via masked graph representation learning. In *Proceedings of the 33rd USENIX Security Symposium*. USENIX Association, 2024.

Baoxiang Jiang, Tristan Bilot, Nour El Madhoun, Khaldoun Al Agha, Anis Zouaoui, Shahrear Iqbal, Xueyuan Han, and Thomas Pasquier. ORTHRUS: Achieving high quality of attribution in provenance-based intrusion detection systems. In *Proceedings of the 34th USENIX Security Symposium*. USENIX Association, 2025.

Seokweon Jung, DongHwa Shin, Hyeon Jeon, Kiroong Choe, and Jinwook Seo. MoNetExplorer: A visual analytics system for analyzing dynamic networks with temporal network motifs. *IEEE Transactions on Visualization and Computer Graphics*, 30(10):6725–6739, 2024. doi: 10.1109/TVCG.2023.3337396.

Alexander D. Kent. *Cyber Security Data Sources for Dynamic Network Research*, chapter 2, pp. 37–65. World Scientific, 2015. doi: 10.1142/9781786340757_0002.

Alexander D. Kent, Lorie M. Liebrock, and Joshua C. Neil. Authentication graphs: Analyzing user behavior within an enterprise network. *Computers & Security*, 48:150–166, 2015. doi: 10.1016/j.cose.2014.09.001.

Isaiah J. King and H. Howie Huang. Euler: Detecting network lateral movement via scalable temporal link prediction. *ACM Transactions on Privacy and Security*, 26(3):1–36, 2023. doi: 10.1145/3588771.

Isaiah J. King, Bernardo Trindade, Benjamin Bowman, and H. Howie Huang. Cybergfm: Graph foundation models for lateral movement detection in enterprise networks, 2026. URL `https://arxiv.org/abs/2601.05988`.

Qi Liu, Kaibin Bao, Wajih Ul Hassan, and Veit Hagenmeyer. HADES: Detecting and investigating active directory attacks via whole network provenance analytics. *IEEE Transactions on Dependable and Secure Computing*, 23(1):936–953, 2026. doi: 10.1109/TDSC.2025.3611866.

Yixin Liu, Zhao Li, Shirui Pan, Chen Gong, Chuan Zhou, and George Karypis. Anomaly detection on attributed networks via contrastive self-supervised learning. *IEEE Transactions on Neural Networks and Learning Systems*, 33(6):2378–2392, 2022. doi: 10.1109/TNNLS.2021.3068344.

Xiaoxiao Ma, Jia Wu, Shan Xue, Jian Yang, Chuan Zhou, Quan Z. Sheng, Hui Xiong, and Leman Akoglu. A comprehensive survey on graph anomaly detection with deep learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12012–12038, 2023. doi: 10.1109/TKDE.2021.3118815.

Hezhe Qiao and Guansong Pang. Truncated affinity maximization: One-class homophily modeling for graph anomaly detection. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Hezhe Qiao, Chaoxi Niu, Ling Chen, and Guansong Pang. Anomalygfm: Graph foundation model for zero/few-shot anomaly detection. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2326–2337, 2025a. doi: 10.1145/3711896.3736843.

Hezhe Qiao, Hanghang Tong, Bo An, Irwin King, Charu Aggarwal, and Guansong Pang. Deep graph anomaly detection: A survey and new perspectives. *IEEE Transactions on Knowledge and Data Engineering*, 37(9):5106–5126, 2025b. doi: 10.1109/TKDE.2025.3581578.

Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. Zero trust architecture. Technical Report NIST Special Publication 800-207, National Institute of Standards and Technology, 2020.

Amit Roy, Juan Shu, Jia Li, Carl Yang, Olivier Elshocht, Jeroen Smeets, and Pan Li. Gad-nr: Graph anomaly detection via neighborhood reconstruction. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pp. 576–585, 2024. doi: 10.1145/3616855.3635767.

Christos Smiliotopoulos, Georgios Kambourakis, and Constantinos Kolias. Detecting lateral movement: A systematic survey. *Heliyon*, 10(4):e26317, 2024. doi: 10.1016/j.heliyon.2024.e26317.

Hazem M. Soliman, Dušan Sovilj, Geoff Salmon, Mohan Rao, and Niranjan Mayya. RANK: AI-assisted end-to-end architecture for detecting persistent attacks in enterprise networks. *IEEE Transactions on Dependable and Secure Computing*, 21(4):3834–3850, 2024. doi: 10.1109/TDSC.2023.3338136.

Robin Sommer and Vern Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *2010 IEEE Symposium on Security and Privacy*, pp. 305–316, 2010. doi: 10.1109/SP.2010.25.

Rafael Uetz, Marco Herzog, Louis Hackländer, Simon Schwarz, and Martin Henze. You cannot escape me: Detecting evasions of SIEM rules in enterprise networks. In *Proceedings of the 33rd USENIX Security Symposium*. USENIX Association, 2024.

Songlei Wang, Yifeng Zheng, and Xiaohua Jia. GraphGuard: Private time-constrained pattern detection over streaming graphs in the cloud. In *Proceedings of the 33rd USENIX Security Symposium*. USENIX Association, 2024.

Chunyu Wei, Wenji Hu, Xingjia Hao, Yunhai Wang, Yueguo Chen, Bing Bai, and Fei Wang. Graph evidential learning for anomaly detection. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3122–3133, 2025. doi: 10.1145/3711896.3736989.

Fan Yang, Jiacen Xu, Chunlin Xiong, Zhou Li, and Kehuan Zhang. PROGRAPHER: An anomaly detection system based on provenance graph embedding. In *Proceedings of the 32nd USENIX Security Symposium*. USENIX Association, 2023.

Bo Zhang, Yansong Gao, Changlong Yu, Boyu Kuang, Zhi Zhang, Hyoungshick Kim, and Anmin Fu. TAPAS: An efficient online APT detection with task-guided process provenance graph segmentation and analysis. In *Proceedings of the 34th USENIX Security Symposium*. USENIX Association, 2025.

Xitong Zhang, Yixuan He, Nathan Brugnone, Michael Perlmutter, and Matthew Hirn. Magnet: A neural network for directed graphs. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27003–27015, 2021.

Yu Zheng, Ming Jin, Yixin Liu, Lianhua Chi, Khoa T. Phan, and Yi-Ping Phoebe Chen. Generative and contrastive self-supervised learning for graph anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12220–12233, 2023. doi: 10.1109/TKDE.2021.3119326.

A Notation and Detailed Problem Formulation

Notation. The stream is $\mathcal{D} = \{e_1, e_2, \ldots\}$ with $e_i = (t_i, s_i, d_i, \mathbf{a}_i)$; $\mathcal{W}_k$ is the kth observation window and $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k, \mathbf{X}_k, \mathbf{F}_k)$ the directed attributed graph built from it. For a node $v \in \mathcal{V}_k$ we write $d^{\text{in}}(v)$ and $d^{\text{out}}(v)$ for its in- and out-degree within the window. $\mathcal{C}_k = \{v \in \mathcal{V}_k : d_k^{\text{out}}(v) > 0\}$ is the active-source candidate set and $A_{\text{raw}}(v) = d^{\text{out}}(v)/(d^{\text{in}}(v) + 1)$ is the raw directional expansion ratio, with $A(v)$ its within-window percentile rank over $\mathcal{C}_k$. $N_k(v)$ counts the windows before k in which v appeared, $\kappa_k(v)$ those in which v reached the top K of the asymmetry order, and $P_k(v) = \kappa_k(v)/N_k(v)$ is the persistence statistic. $\hat{y}_\theta(v)$ is the output of the learned structural branch for the window (written R in the tables), $y_k(v)$ its label-free regression target, and α_k the fusion weight. $\mathcal{T}_k$ is the extreme-asymmetry set of window k, $\mathcal{H}_k$ the hosts v with $N_k(v) \geq 3$, Θ the persistence threshold, $S_k(v)$ the final score of Eq. 2, s_k the spread statistic of Eq. 4, and $q(\cdot)$ the within-window percentile rank over all hosts. π_k is the produced ranking and AP_k its Average Precision.

Why the observation unit is part of the problem. A detector that receives a window boundary implicitly receives a decision about which events are compared with which. Windows taken from a known incident timeline suppress the surrounding background traffic, so the resulting graph is denser in evidence than a graph the detector would have built on its own, and a score evaluated only on such graphs is optimistic. This is the leakage mode studied under the name of bad design and tailored evaluation Flood et al. (2024); Bouke & Abdullah (2023); Hutzelmann et al. (2024); Bilot et al. (2025). Our formulation therefore fixes the boundary rule in advance of any label access and lets the same rule run over the whole stream; AWC is specified in Appendix C.1.

Ranking objective. Average Precision is used because the positive class contains at most four machines inside a graph with roughly ten thousand nodes (Table 4), and the operational output is a short queue read from the top. A method that places one pivot at rank 1 and the rest at the bottom scores above a method that places every pivot in the middle of the list, which is the ordering the

analyst workflow actually consumes. The analyst-budget view in Table 8 reports the same behavior in the form of top-k counts.

B Datasets and Baselines

B.1 Dataset and Adaptive Window Statistics

Table 3 summarizes the stream and its partitioning. The stream contains $1{,}051{,}430{,}459$ authentication events; AWC emits 525 non-overlapping windows, of which 54 contain at least one active recorded pivot.

Table 3: Stream partitioning of the LANL enterprise authentication dataset. These are properties of the stream and of AWC; the window statistics on which the detector operates are given in Table 4, and the last row here is the population a window ranks a host within.

Statistic	Value
Total Authentication Events	1,051,430,459
Source Parquet Log Files	211
Total Emitted Adaptive Windows	525
Attack-Containing (Positive) Windows	54
Background (Unlabeled) Windows	471
Attack Window Ratio	10.3%
Total Active Nodes per Window (Median)	9,780

Table 4 gives the distribution of graph scale and elapsed duration. The raw event count is tightly concentrated at the $M_{\min} = 2{,}000{,}000$ threshold, as the rule intends, while elapsed duration spans 1.2 to 7.1 hours with a median of 2.7 hours. The rule therefore holds observation volume roughly constant and lets the wall-clock span of a window vary with the arrival rate.

Table 4: Graph scale and duration across the 525 adaptive windows. Raw event counts are bounded below by the $M_{\min}$ threshold and cluster on it, whereas elapsed duration varies by almost an order of magnitude, which is the behavior the boundary rule is designed to produce.

Metric	Min	Q1	Median	Q3	Max	Mean
Elapsed Duration (h)	1.2	2.0	2.7	3.2	7.1	2.7
Raw Event Count	2,000,000	2,000,000	2,000,000	2,000,000	2,779,791	2,002,725
Directed Edges ($\lvert\mathcal{E}_k\rvert$)	50,000	61,560	69,957	95,194	114,553	76,885
Unique Sources (V_{src})	5,264	8,338	9,485	10,748	11,304	9,485
Unique Destinations (V_{dst})	5,383	8,355	9,449	10,613	11,166	9,430
Total Active Nodes ($\lvert\mathcal{V}_k\rvert$)	5,508	8,603	9,780	11,110	11,727	9,791

Table 5 lists the four pivot identities, their occurrence counts and the behavioral archetype each represents. `C17693` is the high-volume case, `C19932` the low-volume case, `C22409` mixes moderate volume with unusual service roles, and `C18025` appears once.

Table 5: Known malicious pivot identities, the number of the 54 positive windows in which each is active, and the behavioral archetype the per-pivot breakdown in Table 9 tracks.

Pivot Host	Occurrences	Behavioral Archetype
`C17693`	44	Expansion-dominant, high-fanout subnet scanner
`C19932`	14	Low-volume targeted administrative traversal; limited directional fanout
`C22409`	3	Structural role anomaly, context-inconsistent service usage
`C18025`	1	Rare strike-team intrusion, highly constrained footprint

Graph coverage. Every one of the 525 emitted windows is materialized as a graph, and the memory-based stage of the pipeline — the asymmetry score and the persistence statistic — advances over the whole stream. That is what the attack-unlabeled analysis of Appendix D.8 reads. Average Precision is defined only on the windows that contain a recorded pivot, and the structural branch is scored there, so the ranking comparison of §5 covers all 54 positive windows.

B.2 Baseline Descriptions

We compare against seven public methods. Five are general graph anomaly detection models that differ in what they treat as evidence of normality. Two are temporal or cyber-oriented models built for link prediction on evolving enterprise networks; both report edges, and both are reduced to host scores here as described below. Every method is evaluated on the same 54 adaptive-window graphs, the same candidate sets, the same labels and the same ranking code, and none receives a segmentation of its own.

B.2.1 General graph anomaly detection

- **GAD-NR** Roy et al. (2024) reconstructs a node's entire one-hop neighborhood: its local structure, the node's own attributes, and the attributes of its neighbors. The score is a weighted sum of the three reconstruction losses, so it measures how hard that neighborhood is to reproduce. The method is unsupervised.
- **GEL** Wei et al. (2025) is a graph autoencoder that reconstructs node features and the adjacency matrix, and it fits a distribution to each reconstruction, so that a poor reconstruction and an uncertain one can be told apart. A node is scored by that uncertainty together with the reconstruction error, and the method is unsupervised. GEL's encoder starts from node features, and the observation graphs of this paper give each host two degree features and nothing more, so this entry adds Laplacian positional encodings, which give the encoder the neighborhood structure it would otherwise have to read from node attributes.
- **TAM** Qiao & Pang (2023) assumes homophily: normal nodes hold strong affinity with their neighbors, and that affinity is weaker around anomalies. It maximizes the local affinity of a node to its neighbors and takes the negative of that affinity as the anomaly score, after truncating edges that violate the assumption. TAM is unsupervised.
- **CoLA** Liu et al. (2022) forms contrastive instance pairs of a target node against a sampled local subgraph and trains a binary classifier over them. A node is scored by the negated discriminator output on its positive pair. Both this form and the difference between the negative and positive pair scores that the paper reports are present in the reference implementation. No category or abnormality labels enter training.
- **AnomalyGFM** Qiao et al. (2025a) is a graph foundation model for zero- and few-shot detection. It aligns learnable normal and abnormal prototypes with the residual of a node's representation relative to its neighbors, and scores a node by its similarity to those prototypes. The zero-shot setting trains on an auxiliary graph and applies the model to new graphs without further tuning; the model is evaluated here in that setting.

B.2.2 Temporal and cyber-oriented

- **EULER** King & Huang (2023) models lateral movement as temporal link prediction: it learns normal behavior from earlier snapshots and assigns a likelihood to each edge that appears later, so an edge with low likelihood is an anomalous connection. A host score is obtained by averaging one minus the edge probability over that host's outgoing edges.
- **CybGFM** King et al. (2026) is a graph foundation model for enterprise security telemetry. It is pretrained on benign network activity by masked token prediction over random walks, fine-tuned for link prediction, and used as an edge-level anomaly detector; the host score averages one minus the edge probability over the host's outgoing edges. The pretraining stage is not reproduced here: it runs over large graph corpora, and the reference implementation does not include the pretrained checkpoints. What this entry evaluates is the random-walk link-prediction mechanism at single-process scale.

Algorithm 1 AWC

Require: Chronological stream $\mathcal{D}$ and thresholds $\{M_{\min}, E_{\min}, V_{\mathrm{src},\min}, V_{\mathrm{dst},\min}\}$
Ensure: Sequence of non-overlapping observation windows
current $\leftarrow \emptyset$; windows $\leftarrow [\,]$
for each event $e \in \mathcal{D}$ in chronological order **do**
  Append e to current and update $\{M, E, V_{\mathrm{src}}, V_{\mathrm{dst}}\}$
  if $M \geq M_{\min}$ and $E \geq E_{\min}$ and $V_{\mathrm{src}} \geq V_{\mathrm{src},\min}$ and $V_{\mathrm{dst}} \geq V_{\mathrm{dst},\min}$ **then**
    Append current to windows; current $\leftarrow \emptyset$
  end if
end for
if current $\neq \emptyset$ **then**
  Merge residual events into the last window
end if
return windows

A **alone** is an internal diagnostic: it is the first stage of ASYMGAD on its own, so it isolates what the two learned components add. §5.1 reports it alongside the baselines.

C Algorithm and Implementation Details

C.1 Graph Construction, Features, and Algorithms

Node and edge features. Each node carries two features: the log in-degree $\log(1 + d^{\mathrm{in}}(v))$ and the log out-degree $\log(1 + d^{\mathrm{out}}(v))$. Each aggregated edge $u \to v$ carries a 37-dimensional attribute vector $\mathbf{f}_{uv}$: dimensions 0–5 encode authentication mechanisms and protocol types, dimensions 6–35 encode service-port indicators, and dimension 36 stores the log interaction count $\log(1 + c_{uv})$. The authentication records of this dataset do not carry a port field; the port indicators are taken from the network flow logs released with the same dataset and are attached to an authentication edge when a flow between the same ordered machine pair falls inside the window's time span, so they record which services the two machines were seen exchanging during the window. The reported structural branch uses the directed edge set and the two log-degree node features; none of the 37 edge-attribute dimensions is an input to the reported learned score, and no result in this paper depends on them.

Window construction. Algorithm 1 applies the four evidence thresholds sequentially over the stream. The rule reads only event counts and the cardinalities of the source and destination sets, so it is label-free and can run online: a window is closed at the first event that satisfies all four conditions. Residual events at the end of the stream are merged into the last complete window.

Ranking. Algorithm 2 gives the per-window procedure. Two ordering constraints matter for label isolation: the persistence state used at window k is the state strictly before k, and the memory is updated only after π_k has been formed. No step of the procedure reads an incident label.

C.2 Structural Branch Training Protocol and Label Isolation

Architecture and optimization. The structural branch is a single directed message-passing layer with hidden dimension 64, trained with Adam at learning rate 3×10^{-3} for 25 epochs per refit. At a refit the branch is trained from scratch on a buffer holding the sixty most recent observation windows of the stream; the fitted parameters are then reused for subsequent windows until the next refit. The buffer is a window of the stream: every observation window enters it as the stream advances, whether or not it contains an attack, and the oldest leaves.

Refit cadence. Refits occur on a fixed cadence of eight observation windows. The counter is a position in the stream, fixed before any window is scored, so refits fall where the stream says they fall. The schedule therefore carries over unchanged to a live stream.

Algorithm 2 Window-Level Ranking

Require: Observation graph $\mathcal{G}_k$, persistence memory, buffer $\mathcal{B}$ of past windows, counter of observation windows since the last fit
Ensure: Ranked candidates π_k and updated persistence memory
if the counter has reached eight observation windows **then**
  Refit the structural branch on $\mathcal{B}$ (label-free target) and reset the counter
else
  Reuse the parameters fitted at the previous refit
end if
Compute $\hat{y}_\theta(v)$ for $v \in \mathcal{V}_k$
Compute $P_k(v)$ from the persistence state strictly before window k (Eq. 1)
Compute $A(v)$ and the window statistic s_k (Eq. 4)
Compute α_k from s_k (Eq. 3)
$S_k(v) \leftarrow -\infty$ for every v with $P_k(v) \geq \Theta$; $S_k(v) \leftarrow \alpha_k\, q(A(v)) + (1-\alpha_k)\, q(\hat{y}(v))$ otherwise
$\pi_k \leftarrow$ the candidates sorted by decreasing $S_k(v)$
Update $N_k(v)$ and $\kappa_k(v)$ with window k
return π_k

Persistence and fusion settings. Hosts enter the memory after three appearances. The extreme-asymmetry set holds the $K = 50$ hosts with the largest expansion ratio, the persistence threshold is $\Theta = 2/3$, and the fusion clamp is $[\alpha_{\mathrm{lo}}, \alpha_{\mathrm{hi}}] = [0.2, 0.8]$. The evidence thresholds are $M_{\min} = 2{,}000{,}000$, $E_{\min} = 50{,}000$, $V_{\mathrm{src},\min} = 5{,}000$ and $V_{\mathrm{dst},\min} = 3{,}000$.

Label isolation. Incident labels are excluded from window construction, graph construction, feature extraction, the regression target, the persistence memory, the refit schedule, the training buffer, the fusion weight and the candidate filter. They are read only after the rankings π_k have been produced, to compute the evaluation metrics and to report what the filter removed. Two properties of the design make this possible. The regression target is the within-window percentile of $P_k(v)$, a function of strictly earlier windows, so the branch cannot copy its own input features into the target. And every constant is fixed before the stream is read. The persistence threshold follows from the definition of the quantity it thresholds: a host is persistently expansion-dominant once it has entered the extreme-asymmetry set in two thirds of its sufficiently observed appearances.

D Additional Experimental Results

D.1 Observation-Window Scheme Comparison

The observation unit is part of the method (§3), so we isolate its effect by holding the whole scoring pipeline fixed and varying only where the window boundaries fall. Every scheme in Table 6 sees exactly the same events and is scored by the same code: events are aggregated into a de-duplicated directed edge set, each present host is scored with the asymmetry score A, the persistence filter of §4.6 is applied, and Average Precision is evaluated on the positive windows that the scheme itself produces.

Two readings follow. First, the boundary rule alone moves the reported number a long way: across segmentations the filtered score of an identical method ranges from 0.1925 to 0.2402, and the unlearned score from 0.0739 to 0.0885. The observation unit is therefore not a detail of the experimental setup but a quantity that the setup has to report.

Second, at matched observation volume the adaptive rule is the better one. The adaptive rule holds its windows at 2.0×10^6 events, and at that volume it exceeds a fixed event-count grid by 12.8% (0.2402 against 0.2130) and a fixed wall-clock grid by 14.6% (0.2402 against 0.2096). The advantage is not the amount of data per window: grids of comparable average volume score differently according to what they contain. The three-hour grid averages 2.27×10^6 events per window and scores below the two-hour grid at 1.51×10^6, and the three-million-event grid nearly doubles the volume of the two-hour grid while gaining 0.0007 of AP (0.2338 against 0.2331). A wall-clock boundary follows the arrival rate and lets the observation volume follow it too, which is what the evidence thresholds

Table 6: Observation-window schemes under a fixed scoring pipeline. Every row sees the same events, the whole stream, and is scored by the same code, so the only difference is where the window boundaries fall; *Positive* is the number of windows in which the scheme places at least one recorded pivot. The adaptive row is the configuration used throughout the paper, so its scores are those reported in Table 1.

Segmentation	Windows	Positive	Events/window	A alone	$A + P$ filter
AWC	525	54	2,002,725	0.0858	0.2402
Wall clock, 1 h	1392	84	755,338	0.0739	0.1925
Wall clock, 2 h	696	56	1,510,676	0.0885	0.2331
Wall clock, 3 h	464	45	2,266,014	0.0752	0.2096
Event count, 2×10^6	526	55	1,998,917	0.0781	0.2130
Event count, 3×10^6	351	45	2,995,528	0.0746	0.2338

exist to prevent: the three grids average 0.76, 1.51 and 2.27 million events per window, a factor of three, whereas the adaptive rule holds its event count to a standard deviation of 2.2% of its mean.

The unlearned column's ordering across segmentations is not systematic — the two-hour grid exceeds the adaptive window's own raw score — whereas the filtered column, which is the regime the method operates in, places the adaptive rule first. This comparison holds the scoring pipeline fixed and therefore excludes the structural branch, whose contribution is quantified under the segmentation used throughout the paper in Appendix D.4.

D.2 Full Benchmark Results

Table 7 gives every per-window Average Precision behind Table 1, for the seven public baselines, A alone and AsymGAD, so that the aggregate statistics and the paired tests can be checked window by window.

Table 8 restates the comparison of §5.2 in analyst-budget terms: how often a known pivot is reachable within a queue of 10, 50 or 100 candidates, and the median rank of the best-ranked known pivot. The count-matched random control is identical to A alone on all four statistics.

D.3 Per-Pivot and Per-Window Analyses

Table 9 breaks the benchmark down by pivot identity. AsymGAD leads on three of the four pivots; the exception is the single-window `C18025`, where the best score of any method is 0.0039 (AnomalyGFM) and AsymGAD itself reaches 0.0001. We therefore do not claim `C18025` as a recovery, and note that a single window cannot support a per-pivot conclusion in either direction.

The gaps on the remaining pivots differ greatly in size. On the dominant `C17693` the margin over CybGFM is $9.1\times$, and on the low-volume `C19932` it is $10\times$. What differs is the absolute level: on `C19932` the best AP of any method is 0.0324, so the relative lead there rests on small numbers.

Table 10 gives the median and best rank of each pivot within the windows in which it is active. The rank improvement is broad: on `C19932` the median rank moves from 1,794 under A alone to 1,526 under AsymGAD, and the single `C18025` window improves from rank 8,518 to 7,927. The latter is an improvement that still leaves the pivot far below any practical analyst budget.

Figure 3 shows the signed per-window rank change of the best-ranked known pivot relative to A alone. The rank improves in 51 of the 54 positive windows and worsens in 3, and the median rank of the best-ranked pivot falls from 142 to 130.

Table 11 collects the paired per-window statistics for every method.

D.4 Component Ablation

Table 12 gives the full component ablation; §5.5 reads the first four rows. The rows this table adds show what the filter is worth on its own: fusing the two components without the filter reaches

Table 7: Per-window Average Precision on all 54 positive adaptive windows.

Window	A alone	GAD-NR	GEL	TAM	AnomalyGFM	EULER	CybGFM	CoLA	AsymGAD
14	0.0833	0.0002	0.0086	0.0001	0.0333	0.0011	0.0182	0.0120	1.0000
19	0.0001	0.0002	0.0001	0.0001	0.0039	0.0001	0.0027	0.0001	0.0001
20	0.0005	0.0018	0.0069	0.0007	0.0056	0.0217	0.0058	0.0009	0.0017
21	0.0270	0.0002	0.0110	0.0001	0.0043	0.0026	0.0091	0.0227	0.1667
22	0.0001	0.0002	0.0060	0.0001	0.0031	0.0001	0.0172	0.0001	0.0001
39	0.0001	0.0005	0.0005	0.0008	0.0040	0.0002	0.1111	0.0001	0.0001
40	0.0007	0.0003	0.0004	0.0004	0.0032	0.0094	0.0029	0.0053	0.0011
41	0.0001	0.0001	0.0045	0.0009	0.0079	0.0003	0.0046	0.0001	0.0002
43	0.0003	0.0001	0.0002	0.0002	0.0072	0.0008	0.0058	0.0022	0.0008
44	0.0001	0.0001	0.0001	0.0003	0.0048	0.0001	1.0000	0.0001	0.0001
53	0.0002	0.0004	0.0001	0.0001	0.0001	0.0002	0.0035	0.0003	0.0003
54	0.0005	0.0005	0.0004	0.0002	0.0036	0.0002	0.0099	0.0002	0.0007
59	0.0009	0.0001	0.0004	0.0001	0.0001	0.0002	0.0008	0.0008	0.0011
69	0.0833	0.0002	0.0045	0.0001	0.0115	0.0062	0.0116	0.0073	0.2500
70	0.1111	0.0002	0.0001	0.0001	0.0161	0.0086	0.0130	0.0098	0.2500
71	0.3333	0.0001	0.0070	0.0001	0.0286	0.0012	0.0069	0.0054	1.0000
72	0.3333	0.0001	0.0074	0.0001	0.0044	0.0008	0.0084	0.0025	1.0000
73	0.3333	0.0001	0.0003	0.0001	0.0400	0.0006	0.0096	0.0019	1.0000
74	0.3333	0.0027	0.0093	0.0001	0.0526	0.0003	0.0132	0.0001	1.0000
75	0.3333	0.0002	0.0037	0.0002	0.0333	0.0001	0.0135	0.0005	1.0000
81	0.0169	0.0002	0.0002	0.0011	0.0021	0.0003	0.0021	0.0028	0.0294
101	0.0345	0.0002	0.0002	0.0001	0.0093	0.0013	0.0128	0.0015	0.1250
102	0.3333	0.0001	0.0083	0.0002	0.0312	0.0033	0.0085	0.0024	1.0000
103	0.3333	0.0002	0.0011	0.0001	0.0044	0.0008	0.0076	0.0005	1.0000
104	0.1250	0.0002	0.0048	0.0001	0.0204	0.0270	0.0074	0.0111	0.5000
105	0.0509	0.0003	0.0052	0.0012	0.0013	0.0021	0.0055	0.0092	0.1943
106	0.3333	0.0001	0.0001	0.0001	0.0400	0.0009	0.0097	0.0066	1.0000
111	0.0001	0.0001	0.0001	0.0009	0.0060	0.0001	0.0041	0.0001	0.0001
112	0.0001	0.0002	0.0008	0.0001	0.0043	0.0001	0.1000	0.0001	0.0001
113	0.0717	0.0003	0.0009	0.0002	0.0051	0.0003	0.0031	0.0039	0.2502
114	0.0874	0.0002	0.0002	0.0006	0.0054	0.0051	0.0039	0.0050	0.2555
115	0.2500	0.0001	0.0001	0.0001	0.0244	0.0119	0.0069	0.0037	0.5000
116	0.2000	0.0002	0.0100	0.0001	0.0005	0.0010	0.0032	0.0123	0.5000
121	0.0001	0.0002	0.0009	0.0001	0.0081	0.0001	0.0068	0.0001	0.0002
122	0.0001	0.0001	0.0003	0.0001	0.0032	0.0001	0.0556	0.0001	0.0001
123	0.2500	0.0002	0.0034	0.0001	0.0250	0.0003	0.0080	0.0063	1.0000
124	0.0526	0.0001	0.0034	0.0001	0.0020	0.0019	0.0061	0.0028	0.1000
132	0.0003	0.0003	0.0013	0.0004	0.0017	0.0002	0.0093	0.0002	0.0004
135	0.0074	0.0001	0.0002	0.0010	0.0114	0.0005	0.0056	0.0035	0.0114
136	0.2500	0.0002	0.0004	0.0003	0.0312	0.0043	0.0076	0.0028	0.5000
169	0.0019	0.0001	0.0002	0.0002	0.0001	0.0002	0.0002	0.0004	0.0021
178	0.0001	0.0001	0.0002	0.0001	0.0026	0.0001	0.0098	0.0001	0.0001
179	0.0001	0.0001	0.0002	0.0001	0.0013	0.0001	0.0083	0.0001	0.0001
187	0.0003	0.0003	0.0001	0.0002	0.0004	0.0002	0.0007	0.0002	0.0003
221	0.0069	0.0002	0.0047	0.0009	0.0111	0.0041	0.0179	0.0028	0.0102
222	0.2000	0.0001	0.0068	0.0001	0.0294	0.0010	0.0094	0.0058	1.0000
228	0.0072	0.0002	0.0002	0.0002	0.0089	0.0005	0.0053	0.0007	0.0062
229	0.0435	0.0002	0.0007	0.0001	0.0132	0.0011	0.0089	0.0175	0.0500
231	0.0007	0.0001	0.0002	0.0004	0.0001	0.0001	0.0002	0.0008	0.0008
239	0.0001	0.0002	0.0002	0.0002	0.0050	0.0001	0.0045	0.0001	0.0001
240	0.0005	0.0003	0.0001	0.0001	0.0008	0.0005	0.0002	0.0008	0.0005
247	0.0004	0.0003	0.0002	0.0001	0.0001	0.0002	0.0005	0.0010	0.0005
249	0.0005	0.0002	0.0002	0.0001	0.0024	0.0001	0.0069	0.0003	0.0007
250	0.0003	0.0001	0.0004	0.0001	0.0003	0.0003	0.0004	0.0010	0.0003

Table 8: Analyst-budget view: number of the 54 positive windows in which a known pivot appears within the top k, and the median rank of the best-ranked known pivot.

Method	Top-10	Top-50	Top-100	Median rank
A alone	18	24	25	142
$A + P$ filter	21	25	25	124
AsymGAD	23	25	27	130
Random (count-matched)	18	24	25	142

0.1723, below the filtered score alone. The count-matched random deletion control is isolated in Appendix D.7.

Why the components are complementary. The persistence memory uses only the history of the asymmetry order, whereas the structural branch uses the structure of the current graph. Their roles are complementary: the memory removes hosts that are persistently extreme, and the branch supplies

Table 9: Per-pivot mean Average Precision. Each column averages over the positive windows in which that pivot is active; the count is given in parentheses.

Method	C17693 (44)	C18025 (1)	C19932 (14)	C22409 (3)
GAD-NR	0.0003	0.0002	0.0003	0.0002
GEL	0.0029	0.0001	0.0007	0.0019
TAM	0.0003	0.0001	0.0003	0.0010
AnomalyGFM	0.0130	0.0039	0.0014	0.0029
EULER	0.0028	0.0001	0.0010	0.0025
CybGFM	0.0365	0.0027	0.0032	0.0038
CoLA	0.0039	0.0001	0.0017	0.0057
A alone	0.1048	0.0001	0.0093	0.0518
$A + P$ filter	0.2941	0.0001	0.0223	0.1115
AsymGAD	0.3336	0.0001	0.0324	0.1597

Table 10: Median and best rank of the pivot named in each column, over the windows in which that pivot is active.

Method	C17693	C18025	C19932	C22409
A alone	26/3	8518/8518	1794/359	59/45
$A + P$ filter	12/1	8499/8499	1778/339	39/28
AsymGAD	9/1	7927/7927	1526/263	34/28

an ordering signal on the remaining hosts that a single window's asymmetry order does not separate. This is consistent with the per-pivot behavior, where the branch contributes most on the low-volume pivot that the asymmetry order ranks poorly.

D.5 Graph Structure Ablation

The component ablation of Appendix D.4 attributes an increment of 0.0322 to the structural branch without separating the graph from the model that reads it. This subsection isolates the graph by training the same model, target and schedule with neighbour aggregation removed, so that each layer sees only the host's own features; every other quantity in the pipeline is unchanged, including the candidate filter, the fusion weight and the windows. The two configurations are called the *graph arm* and the *graph-free arm*. Table 13 reports both arms.

On the reported seed the graph arm raises the filtered score by 0.0322 and the graph-free arm by 0.0040: message passing supplies most of the branch's contribution, and it is what carries the learned score past what a host's own degree already implies. The graph-free arm leads on at most 10 of the 54 windows in any seed.

The graph's contribution is concentrated where the degree signal is ambiguous. Where the top of the asymmetry order is not sharply separated from its bulk — $s_k \leq 10$, 40 of the 54 windows — the graph-free arm adds nothing to the filtered score, and it is the graph arm that carries the branch's contribution on these windows. Where the order is already sharply peaked the two arms differ by under 0.01 (14 windows), so once a host's own degree settles the question the neighbourhood adds little. On rank-based measures the two arms do not separate consistently across seeds.

D.6 Fusion Weight

Table 14 is the fusion-weight analysis. The weight the method uses is the one Eq. 3 computes per window, and it is the only row not chosen by hand; the other rows hold the weight at hand-set constants and show what the weight contributes to the full pipeline. The mechanism reaches 0.2724, within 0.001 of the highest value in the hand-set part of the sweep; since the method never sees the evaluation, that figure serves as a reference point.

Because s_k is computed from the raw ratio, the mechanism produces a distinct weight in 42 of the 54 windows, spanning 0.678 to 0.800; it consistently reads the asymmetry order as the dominant signal and assigns the branch between 20% and 32% of the weight. The hand-set part is one-sided:

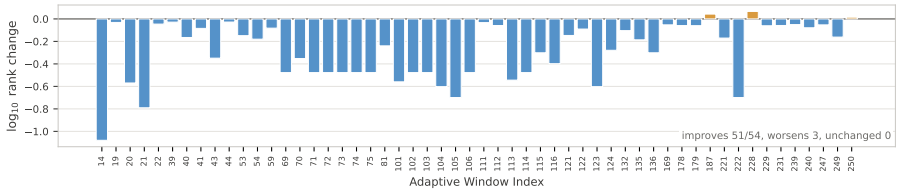


Figure 3: Signed change in the rank of the best-ranked known pivot, ASYMGAD relative to A alone, per positive window on a $\log_{10}$ scale. Negative bars are improvements.

Table 11: Paired per-window comparison of ASYMGAD against each baseline over the 54 positive windows: mean paired difference, 95% percentile bootstrap interval (20,000 window resamples), one-sided Wilcoxon signed-rank p-value, and the Holm-Bonferroni adjustment across the seven public baselines and the unlearned score.

Baseline	Mean AP	Mean Δ	95% CI	p	Holm p
A alone	0.0858	$+0.1866$	$[+0.1150, +0.2648]$	$8.44e-10$	$7.60e-09$
GAD-NR	0.0003	$+0.2722$	$[+0.1706, +0.3805]$	$5.04e-08$	$2.52e-07$
GEL	0.0024	$+0.2701$	$[+0.1711, +0.3787]$	$6.96e-07$	$2.79e-06$
TAM	0.0003	$+0.2722$	$[+0.1725, +0.3803]$	$1.07e-08$	$8.53e-08$
AnomalyGFM	0.0107	$+0.2617$	$[+0.1634, +0.3668]$	$1.39e-03$	$2.77e-03$
EULER	0.0023	$+0.2701$	$[+0.1686, +0.3794]$	$2.98e-08$	$1.79e-07$
CybGFM	0.0299	$+0.2425$	$[+0.1309, +0.3577]$	$2.82e-03$	$2.82e-03$
CoLA	0.0033	$+0.2691$	$[+0.1678, +0.3776]$	$2.45e-08$	$1.72e-07$

the branch is most useful when the expanded score keeps most of the weight, and giving the branch more than half costs 0.05 or more ($\alpha = 0.3$ gives 0.2159 against 0.2729 at the far end).

The clamp is the one constant in Eq. 3 that is not computed from the window, and it was fixed with the rest of the design, before the stream was read. Read after the fact, the lower bound cannot bind on any input: the mean of the top 5% is at least the median, so $s_k \geq 1$ and the raw weight is at least $\frac{1}{3}\log 2 = 0.231$. The upper bound binds in 13 of the 54 windows, and widening the clamp to $[0, 1]$ moves the mean from 0.2724 to 0.2726, an eighth of the spread across the plateau of Table 16.

D.7 PERSISTENCE FILTER CONTROLS

Table 15 isolates the persistence filter against a count-matched random-deletion control that removes the same number of candidates as the threshold does, drawn at random. The control is indistinguishable from A alone on the mean, the median and the median rank, while the persistence filter raises the mean from 0.0858 to 0.2402. The gain therefore comes from the identity of the removed hosts. The removed set is also coherent: over the whole stream its members have a mean in-degree of 0.8, and 99.9% of the distinct hosts they exchange events with in a window are contacts they initiate themselves. What places them at the top of A_{raw} is therefore the near-zero denominator: they contact a median of 22 distinct hosts in a window, and a few of them reach into the thousands. That is the profile of a scanner, a backup job or a patch server.

The separation the filter relies on is visible in the statistic itself. The pivot rate rises with P over the range the filter keeps: it is $0.3\times$ the base rate below 0.05 and $42\times$ between 0.20 and 0.25, and above 0.30 it reaches $181\times$ (33 of 1,802 host-rows, against a base rate of 0.01%). The host-rows the threshold removes contain no recorded pivot at all, and no pivot exceeds $P = 0.56$. The relation between P and the label is therefore not monotone: the pivots occupy the middle of the range, and that is the range on which the regression target of §4.4 is read.

The threshold is fixed by definition: $\Theta = 2/3$ is what the method takes "persistently expansion-dominant" to mean, and it was set before any window was scored. Table 16 is the post-hoc check on that choice, not the procedure that produced it. The curve is flat over $[0.55, 0.70]$, where the mean AP varies by 0.002 while the number of hosts the gate removes falls from 23.2 to 19.2 per window. The band has a lower edge, and the sweep locates it: at $\Theta = 0.50$ the gate removes 19 pivot host-rows and the mean collapses to 0.1974, while from $\Theta = 0.60$ upward it removes none. The best value in the sweep is the same on the first and the second half of the positive windows taken

Table 12: Full component ablation, mean Average Precision over the 54 positive windows. Only the mean is reported; per-window values for A alone and for ASYMGAD are in Table 7.

Configuration	Mean AP
A alone (unlearned)	0.0858
$A + P$ filter	0.2402
R alone (learned branch)	0.1294
$A + R$ (no filter, $\alpha = 0.5$)	0.1723
$A + P + R$ ($\alpha = 0.9$, best fixed)	0.2729
$A + P + R$ (adaptive α_k)	0.2724
Random deletion (count-matched)	0.0862

Table 13: Graph structure ablation, mean Average Precision over the 54 positive windows. The *graph arm* trains the structural branch on the observation graphs; the *graph-free arm* runs the same model, target and schedule with neighbour aggregation removed. Each row states its own quantity: mean AP, the difference between the two arms, or an increment over the filtered score. The last three rows split the 54 windows by the spread statistic s_k of Eq. 4.

Quantity	seed 0	seed 1	seed 2
Filtered score $A + P$ (mAP)	0.2402	0.2402	0.2402
Graph arm (mAP)	0.2724	0.2621	0.2502
Graph-free arm (mAP)	0.2443	0.2409	0.2359
Difference in mAP	$+0.0282$	$+0.0212$	$+0.0144$
Increment, graph arm, $s_k \leq 10$	$+0.029$	$+0.027$	$+0.000$
Increment, graph-free arm, $s_k \leq 10$	-0.008	$+0.001$	-0.016
Difference, $s_k > 10$	$+0.003$	$+0.007$	$+0.008$

separately, which shows the plateau does not rest on a few windows; both halves are drawn from the same labeled set, so this is a robustness check and not a held-out one. Read after the fact, $\Theta = 2/3$ sits inside the plateau, 0.0007 below the best point in the sweep, and it removes no recorded pivot — a property observed afterward, not the reason it was chosen.

D.8 ATTACK-UNLABELED WINDOW STABILITY

The 471 windows with no recorded attack carry no label, so no ranking metric is defined on them. They are still informative about the quantity the persistence filter targets: whether a small set of hosts dominates the top of the unfiltered asymmetry order. Table 17 reports the behavior on all 471 of them.

A single host, `C1798`, is the top-ranked host in all 471 attack-unlabeled windows and the only host to hold that position in any of them. The filter behaves there as it does on the rest of the stream: its 7,926 firings come from 41 distinct hosts, 16.8 per window, and removing them leaves 33.2 of the top-K candidates on average. This is the signature the filter is designed for: an infrastructure host whose fanout is extreme in every window is exactly the host that a snapshot-local detector would place at the top of every queue.

D.9 RUN-TO-RUN ROBUSTNESS

The structural branch is the only stochastic component of the pipeline: the asymmetry score, the persistence statistic, the filter threshold, the fusion weight and the paired test are all deterministic given the stream, and each baseline is scored from a frozen per-window record. Re-training the branch is therefore the only source of run-to-run variation in a reported number. We ran it three times, with seeds 0, 1 and 2, and scored each run through the identical pipeline.

The ordering the paper reports is stable across the three runs: ASYMGAD leads every baseline in all of them, the paired bootstrap intervals exclude zero and the Holm-adjusted p-values stay below 3×10^{-3}. The absolute level moves with the seed, and the main text reports seed 0, which is the most favorable of the three on the fused score. The seed affects nothing else: window construction, the asymmetry score, the persistence memory, the candidate filter, the fusion weight and the baselines are all fixed by the stream.

Table 14: Fusion-weight analysis. The hand-set rows are a sensitivity grid: each holds the weight at the stated constant, so the table shows what the weight contributes to the full pipeline. The last row is the weight the method uses, computed per window by Eq. 3; it is the only one not chosen by hand. "Win/Loss" counts windows in which the configuration beats or loses to the filtered score without the branch.

Fusion	Mean AP	Win/Loss vs. filter
fixed $\alpha = 0.0$	0.2191	34/17
fixed $\alpha = 0.1$	0.2283	34/17
fixed $\alpha = 0.2$	0.2197	36/15
fixed $\alpha = 0.3$	0.2159	35/15
fixed $\alpha = 0.4$	0.2470	35/11
fixed $\alpha = 0.5$	0.2601	36/8
fixed $\alpha = 0.6$	0.2666	37/8
fixed $\alpha = 0.7$	0.2700	37/7
fixed $\alpha = 0.8$	0.2722	37/7
fixed $\alpha = 0.9$	0.2729	36/4
Adaptive α_k (Eq. 3)	0.2724	37/7

Table 15: Count-matched random-deletion control. In each window the control removes as many candidates as the persistence filter removes in that window, but chooses them at random, so any gain from $A + P$ beyond this row is attributable to *which* hosts are removed.

Configuration	Mean AP	Median AP	Median rank
A alone	0.0858	0.0070	142
$A + P$ filter	0.2402	0.0081	124
Random deletion (count-matched)	0.0862	0.0071	142

The baselines are single frozen runs, so Table 18 compares three runs of our branch against one run of each baseline.

E LIMITATIONS

Data. The evaluation rests on LANL Kent (2015), the only public enterprise authentication stream of comparable scale that we are aware of. Its recorded attacks are sparse: four pivot identities, one of them active in 44 of the 54 positive windows and another in a single window. The per-pivot conclusions therefore rest on unequal samples, and we report them per pivot.

Assumed adversary. The method assumes an adversary that reuses credentials and moves between machines with directional fanout, and whose use of a given machine is transient relative to the lifetime of the infrastructure that dominates the persistence order. An adversary that holds a machine over many windows, or whose fanout is no higher than that of comparable hosts, is not separated by either signal.

Setting the constants. $K = 50$, $\Theta = 2/3$ and the fusion clamp $[0.2, 0.8]$ are fixed on the stream, without reading a label and before the first window they apply to is scored; Appendix D.7 reports how sensitive the result is to where they sit.

Table 16: The persistence threshold. Each row is the full pipeline of Eq. 2 with Θ set to the stated value and nothing else changed; Θ enters only the gate, so the sweep covers the whole pipeline. The three Average Precision columns are all 54 positive windows and the first and second halves of them in stream order. "Removed" is the mean number of hosts the gate excludes from a positive window; the last column counts how many of those host-rows are recorded pivots, which is why the band has a lower edge. The curve is flat from 0.55 to 0.70, the best value is the same on all 54 windows and on each half separately, and the threshold the paper uses removes no pivot at all.

Θ	Mean AP			Removed	Pivots
	all 54	first half	second half	per window	removed
0.50	0.1974	0.3554	0.0394	25.7	19
0.55	0.2731	0.3910	0.1552	23.2	2
0.60	0.2731	0.3909	0.1552	21.9	0
$2/3$	0.2724	0.3897	0.1552	20.2	0
0.70	0.2714	0.3876	0.1551	19.2	0
0.75	0.2678	0.3871	0.1486	18.3	0
0.80	0.2537	0.3658	0.1416	16.1	0
0.85	0.2525	0.3635	0.1416	15.0	0
0.90	0.2430	0.3445	0.1415	14.1	0
0.95	0.2391	0.3382	0.1399	13.1	0
1.01	0.1381	0.2164	0.0599	0.0	0

Table 17: Ranking behavior on the attack-unlabeled windows. These windows contain no recorded attack, so no metric is defined on them; the statistics characterize the unfiltered asymmetry order on these windows and what the filter removes from it.

Statistic	Value
Attack-unlabeled windows analyzed	471 of 471
Most frequent top-ranked host	`C1798` (471 of 471 windows)
Mean Jaccard overlap of top-K sets	0.289
Distinct hosts the filter removes	41
Filter firings per window	16.8
Top-K queue after removing them	33.2

Table 18: The structural branch re-run under three seeds, everything else held fixed. The branch is the only stochastic component of the pipeline: the asymmetry score, the persistence statistic, the filter threshold and the fusion weight do not depend on the branch's initialisation, and the baselines are scored from a frozen per-window record. Rows are one run each; the last row is their mean and standard deviation. *Hit@10/50/100*, *median rank* and *wins* are the analyst-budget and unique-win statistics of Tables 8 and 1, recomputed for each run.

Seed	Fused mAP	R alone	$A + P + R$, $\alpha = 0.9$	hit@10/50/100	Med. rank	Wins
0	0.2724	0.1294	0.2729	23/25/27	130	28
1	0.2621	0.0854	0.2581	22/25/28	94	29
2	0.2502	0.1061	0.2644	23/25/26	118	28
mean $\pm$ std	0.2616 ± 0.0111	0.1069 ± 0.0220	0.2651 ± 0.0074	—	—	28.3 ± 0.6