

Article

Prototype-Based Transferability Analysis for Few-Shot Domain Adaptation in Cross-Domain Intrusion Detection

Kangseok Kim

Department of Cyber Security, College of Computing and Informatics in Ajou University,
Suwon 16499, Republic of Korea; kangskim@ajou.ac.kr

Featured Application

The proposed framework can be integrated into enterprise security architecture for automated, cross-platform intrusion detection between Windows and Linux environments without log relabeling. Additionally, cybersecurity vendors can utilize the multi-dimensional criteria to pre-assess model transferability, significantly reducing false-alarm rates and operational monitoring costs in new deployment environments.

Abstract

Few-shot domain adaptation (FSDA) has become an important approach for cross-domain intrusion detection by enabling models to leverage limited labeled target data under distribution shifts. Although numerous adaptation methods have been proposed, their effectiveness often varies considerably across transfer scenarios, leading to inconsistent performance across domains. This study investigates how prototype-based transferability can be used to characterize source–target compatibility prior to adaptation. To this end, a transferability-aware perspective on FSDA is presented by distinguishing between intra-domain separability, which characterizes the internal class structure of a domain, and cross-domain transferability, which reflects how effectively source-derived representations generalize to a target domain. Based on this distinction, a set of asymmetric transferability metrics is introduced to characterize prototype-based transferability from complementary perspectives. Across the evaluated transfer scenarios, stronger intra-domain separability did not necessarily coincide with better cross-domain transferability. Furthermore, the empirical results indicate that different adaptation strategies exhibit different behaviors across the examined transfer directions. Across the two evaluated transfer directions, the prototype-based transferability analysis and the few-shot adaptation results exhibit different behaviors. While target-only few-shot learning performs competitively in one transfer direction, prototype-based alignment methods provide larger improvements in the other, illustrating that adaptation performance reflects not only source–target compatibility but also intrinsic target-domain separability and the adaptation process itself. These observations provide an empirical perspective for interpreting the varying performance of existing FSDA methods and highlight the importance of considering transfer conditions when selecting adaptation strategies for cross-domain intrusion detection.



Academic Editors: Yirui Wu and
Shaohua Wan

Received: 18 June 2026

Revised: 4 July 2026

Accepted: 9 July 2026

Published: 10 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article
distributed under the terms and

conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: few-shot domain adaptation; cross-domain intrusion detection; transferability; prototype alignment; uncertainty-aware refinement

1. Introduction

1.1. Background and Motivation

Cross-domain intrusion detection remains a fundamental challenge due to distribution shifts between source and target environments. While few-shot domain adaptation (FSDA) has emerged as a practical solution by leveraging limited labeled target data, existing research has primarily focused on developing increasingly sophisticated adaptation mechanisms. However, a critical question remains largely unexplored: When and why do these adaptation strategies actually work? In practice, the effectiveness of FSDA methods varies significantly across different transfer scenarios, often leading to inconsistent and difficult-to-explain results.

This observation suggests that domain adaptation should not be treated as a uniform problem. Instead, the adaptation efficacy depends on the underlying transfer conditions between domains. In this work, it is argued that an overlooked factor—prototype-based transferability—provides useful insights into understanding adaptation behavior across heterogeneous transfer scenarios. Instead of merely proposing another adaptation method, this study aims to establish a foundational understanding of cross-domain transfer within few-shot settings.

To this end, a transferability-aware perspective on FSDA is proposed. This approach differentiates between two related but fundamentally distinct concepts: intra-domain separability and cross-domain transferability. The former characterizes how well classes are separated within a domain, while the latter measures how effectively source-derived representations generalize to a different domain. While previous studies often implicitly assume a connection between these two properties, this study demonstrates that they can diverge significantly in practice.

Recent work in few-shot learning has focused on metric-based approaches such as prototypical networks, which assume that training and testing data share the same distribution. Domain adaptation methods, on the other hand, aim to reduce distribution discrepancies through techniques such as adversarial alignment and representation learning. Few-shot domain adaptation combines these two paradigms, requiring both generalization under limited supervision and robustness to domain shift.

Despite these advances, existing FSDA approaches typically focus on improving alignment or representation quality without explicitly analyzing the conditions under which such mechanisms are necessary. As a result, their performance varies across different transfer scenarios, and their behavior remains difficult to interpret.

In contrast, this work analyzes FSDA from the perspective of prototype-based transferability. Asymmetric prototype-based transferability metrics are proposed to directly characterize how well source-derived representations generalize to the target domain prior to adaptation. Through extensive experiments, the proposed transferability analysis and the evaluated few-shot adaptation methods exhibit distinct behaviors across the two transfer directions. These findings provide an empirical perspective for interpreting the behavior of existing FSDA methods under heterogeneous transfer conditions and suggest that source–target transferability should be considered when analyzing cross-domain adaptation.

1.2. Related Works

Few-shot domain adaptation aims to transfer knowledge from a labeled source domain to a target domain given only a limited number of labeled target samples. This problem lies at the intersection of few-shot learning and domain adaptation, attracting increasing attention due to its practical relevance in data-scarce security analytics and logging scenarios.

1.2.1. Cross-Domain Security Analytics and Log Transfer

Recent studies have increasingly explored representation transfer and cross-domain generalization in security analytics and log anomaly detection. Log-based anomaly detection systems have evolved from traditional statistical methods toward semantic representation learning using transformer architectures and large language models. Several studies investigate cross-system log transfer by learning domain-invariant semantic representations across heterogeneous software environments. For example, CroSysLog [1] investigates cross-system log anomaly detection through transferable semantic embeddings, while LogAction [2] introduces active domain adaptation strategies for maintaining stable detection performance across dynamic logging systems. More recent approaches such as ADALog [3] and FusionLog [4] further integrate transformer-based architectures and representation fusion strategies to improve cross-domain robustness in log analytics.

Despite these advances, most existing studies focus primarily on software-system logs with relatively similar structures and semantic distributions. In contrast, cross-platform transfer between fundamentally heterogeneous operating system environments, such as Linux system-call traces and Windows security events, remains largely unexplored. Furthermore, prior work rarely analyzes how transfer difficulty, representation geometry, and class separability jointly influence few-shot adaptation behavior under asymmetric domain shifts.

1.2.2. Few-Shot Learning and Few-Shot Domain Adaptation

Few-shot learning (FSL) focuses on learning novel classes from a limited number of labeled samples, typically utilizing metric-based or optimization-based paradigms. The metric-based approaches, including Matching Networks [5], Prototypical Networks [6], and Relation Networks [7], aim to learn transferable similarity structures or a domain-invariant distance space to enable classification with minimal target supervision. Conversely, optimization-based meta-learning methods, most notably Model-Agnostic Meta-Learning (MAML) [8], seek to acquire a highly adaptable parameter initialization. However, both paradigms assume that the training and testing data are drawn from an identical distribution, a condition that rarely holds in real-world deployments. Consequently, their classification performance degrades significantly when directly applied across heterogeneous, unaligned domains.

Domain adaptation (DA) aims to mitigate distribution shifts between source and target domains by learning domain-invariant representations. Traditional approaches often rely on explicit statistical alignment, such as minimizing Maximum Mean Discrepancy in Deep Adaptation Networks (DAN) [9] or matching second-order covariances via Deep CORAL [10], to reduce global domain discrepancies. Building upon these statistical methods, adversarial domain adaptation methodologies align latent feature distributions by encouraging domain indistinguishability, enabling robust transfer even in the complete absence of labeled target data [11]. More recent domain adaptation approaches incorporate adversarial learning frameworks [12], which have further evolved into conditional adversarial learning [13] and discrepancy-aware optimization [14,15]. Notably, Balanced Distribution Adaptation (BDA) [16] adapts domain gaps by dynamically balancing the importance of marginal and conditional distribution differences, thereby better preserving class-level decision boundaries and discriminability during alignment.

Few-shot domain adaptation combines the challenges of both FSL and DA, requiring models to generalize across distinct domain distributions while effectively learning from highly constrained target samples. Early work on domain-adaptive few-shot learning (DA-FSL) highlighted a critical fundamental challenge: aligning global domain distributions may inherently conflict with preserving fine-grained, class-level discriminability [17]. To address

this bottleneck, recent approaches focus on learning shared embedding spaces that strategically balance domain alignment and class discrimination. Prototype-based methods are particularly prominent, as they enable class-level semantic alignment using limited samples. For instance, Transferrable Prototypical Networks (TPN) [18] align cross-domain distributions by mapping each target example to its corresponding prototype in the source domain, whereas task-level augmentation [19] and causal mechanism transfer [20] have been proposed to adapt to complex distributional shifts. Recently, this paradigm has been extended to log-based anomaly detection to bridge systemic gaps across heterogeneous environments. For example, LogTAD [21] adopts adversarial domain adaptation to minimize the distribution gap between diverse system logs. To further leverage minimal target supervision, LogMeta [22] introduces an optimization-based framework using Model-Agnostic Meta-Learning (MAML) for rapid target-domain adaptation, while LogSynergy [23] employs Large Language Models (LLMs) to disentangle system-specific features from system-unified features. Unlike conventional domain adaptation, FSDA operates under extremely sparse target supervision, making reliable estimation of target-domain geometry substantially more challenging.

1.2.3. Prototype-Based, Contrastive, and Representation-Centric FSDA

More recent studies incorporate contrastive learning paradigms alongside prototype alignment to fundamentally improve representation quality. Prototypical contrastive frameworks align cross-domain semantic structures by integrating intra-domain prototypical contrastive learning with cross-domain prototypical self-supervision [24]. Prototype alignment methods are particularly attractive in FSDA because they preserve class-level semantic structures while remaining robust under sparse target supervision. To further enhance robustness, advanced refinement strategies incorporate uncertainty estimation to mitigate noisy prototype updates caused by unstable pseudo-labels under low-shot conditions.

Parallel to these computer vision frameworks, recent state-of-the-art log analysis methods, such as LogDLR [25] and LogICL [26], have attempted to leverage self-supervised or LLM-distilled representations to bridge the semantic gap across different systems. However, even though these methods are tailored for cross-domain configurations, they frequently suffer from severe overfitting or insufficient adaptation when the available target samples are restricted to extreme few-shot scenarios (e.g., 1-shot or 5-shot configurations). This vulnerability stems from their inherent reliance on dense source data distributions to fine-tune representations, which inadvertently induces a source-domain bias given the strictly constrained target supervision characteristic of FSDA.

1.2.4. Transfer Difficulty and Motivation of This Work

Despite these collective advances, existing FSDA methods typically assume a uniform adaptation setting and do not explicitly consider the impact of directional transfer difficulty. In particular, conventional research is heavily directed toward method-centric optimization without analyzing the critical boundary conditions: when target supervision alone is sufficient, or when explicit representation alignment becomes mandatory. Furthermore, pseudo-label-based refinement methods, while effective in large-scale settings, remain notoriously unreliable in few-shot scenarios due to label noise and unstable prediction confidence estimation. Consequently, their empirical performance gains are often highly inconsistent when labeled target data is extremely limited and heavily entangled. Importantly, most prior FSDA studies implicitly assume symmetric transfer behavior and evaluate adaptation under a single fixed domain-pair configuration. As a result, the influence of directional transfer difficulty and asymmetric domain geometry remains poorly understood.

In contrast to prior method-centric studies, this work systematically analyzes few-shot domain adaptation under varying transfer difficulties. The experimental results reveal that adap-

tation behavior differs fundamentally between Low-Difficulty Transfer and High-Difficulty Transfer settings. In particular, the relative importance of target supervision, prototype alignment, and uncertainty-aware refinement is shown to depend strongly on the underlying domain geometry and class separability. By establishing a unified perspective for understanding asymmetric cross-domain intrusion detection, these insights provide a principled explanation for the inconsistent behaviors frequently observed in the existing FSDA literature.

As shown in Table 1, most existing domain adaptation methods are designed for conventional unsupervised or supervised domain adaptation settings and are not specifically intended to analyze source–target transferability prior to adaptation. Prototype-based approaches improve feature alignment under limited supervision, yet they generally assume that a single adaptation strategy can be applied across different transfer scenarios. In contrast, the present study focuses on characterizing prototype-based transferability prior to few-shot adaptation and empirically investigates how different adaptation strategies behave under heterogeneous transfer conditions.

Table 1. Comparison of Representative Cross-Domain Intrusion Detection Methods. Transferability-aware indicates whether a method explicitly evaluates or characterizes source–target transferability prior to domain adaptation.

Method	Cross-Domain	Few-Shot Adaptation	Target Labels	Prototype Alignment	Transfer-Aware
DANN [12]	O	X	Unlabeled	X	X
CDAN [13]	O	X	Unlabeled	X	X
BDA [16]	O	Limited	Unlabeled (pseudo-labels)	X	X
PCS [24]	O	O	Few-shot	O	X
Proposed FSDA	O	O	Few-shot	O	O

1.3. Contributions

The main contributions of this work are as follows:

1. Transferability-Aware Perspective.

This work introduces a transferability-aware perspective on few-shot domain adaptation, shifting the focus from designing better adaptation methods to understanding source–target transferability under heterogeneous cross-domain conditions.

2. Separation of Intra-Domain Structure and Transferability.

A clear distinction is made between intra-domain separability and cross-domain transferability, demonstrating that strong internal class structure does not necessarily imply favorable cross-domain transferability.

3. Asymmetric Prototype-Based Transferability Metrics.

A set of asymmetric metrics is proposed to characterize prototype-based transferability by evaluating how well source-derived representations generalize to the target domain, providing a principled way to characterize transfer behavior.

4. Empirical Analysis and Practical Insights.

Through systematic experiments across the two evaluated transfer directions, the proposed transferability analysis and the evaluated few-shot adaptation methods are shown to exhibit distinct behaviors, providing an empirical perspective for interpreting cross-domain adaptation under heterogeneous transfer conditions.

2. Materials and Methods

2.1. Problem Setup

Few-shot domain adaptation (FSDA) aims to transfer knowledge from a labeled source domain to a target domain when only a very limited number of labeled target samples are available. This study considers cross-domain intrusion detection between heterogeneous operating system log environments.

Let the source domain be defined as $D_s = \{(X_i^s, y_i^s)\}_{i=1}^{N_s}$, where X_i^s denotes a source-domain log representation and $y_i^s \in \{0, 1\}$ indicates the corresponding label (normal or attack). Similarly, the target domain is defined as $D_t = \{(X_j^t, y_j^t)\}_{j=1}^{N_t}$, where only a small subset of labeled target samples is available during adaptation.

Unlike conventional few-shot learning, FSDA simultaneously faces two challenges:

- It has extremely sparse target supervision;
- A substantial domain discrepancy between source and target distributions.

To systematically analyze adaptation behavior under asymmetric transfer conditions, this study explicitly distinguishes between intra-domain separability and cross-domain transferability. The intra-domain separability measures how well normal and attack classes are separated within a single domain, and the cross-domain transferability measures how effectively representations learned from one domain generalize to another domain. Based on these concepts, this study further defines transferability as an asymmetric property of cross-domain adaptation.

2.2. Intra-Domain Separability

To quantify the structural separability of normal and attack samples within a domain, a dimension-wise generalized Fisher score [27] is employed. For each feature dimension k , the Fisher discriminant ratio (or J-index) is computed as:

$$J_k = \frac{(\mu_{atk,k} - \mu_{norm,k})^2}{\sigma_{atk,k}^2 + \sigma_{norm,k}^2} \quad (1)$$

where $\mu_{atk,k}$ and $\mu_{norm,k}$ denote the mean values of the attack and normal classes for dimension k , respectively, while $\sigma_{atk,k}^2$ and $\sigma_{norm,k}^2$ represent their corresponding variances. To capture the overall structural separability across the entire feature space, the dimension-wise Fisher discriminant ratios are averaged over all dimensions. The final intra-domain separability score, denoted as J_{domain} , is computed as follows:

$$J_{domain} = \frac{1}{d} \sum_{k=1}^d J_k \quad (2)$$

where d represents the total number of feature dimensions, and J_k denotes the Fisher discriminant ratio computed for the k -th dimension.

A larger J_{domain} indicates stronger class separability, whereas smaller values imply substantial overlap between normal and attack distributions. The intra-domain separability measures only the internal class structure of a domain and does not directly indicate cross-domain transferability.

2.3. Cross-Domain Transferability

To evaluate how effectively knowledge transfers across domains, this study measures the performance of classifiers trained on the source domain and evaluated directly on the target domain without additional adaptation. Transferability is quantified using PR-AUC, ROC-AUC, and cross-domain J-index consistency.

Unlike intra-domain separability, these metrics evaluate whether source-derived decision structures remain meaningful under target-domain distributions. A high transferability score indicates that the source-domain representation captures behavioral semantics that generalize across heterogeneous domains.

2.4. Transfer Difficulty

Transfer difficulty is defined as the degree to which source-domain representations fail to generalize reliably to the target domain. Importantly, transfer difficulty is inherently asymmetric because the structural geometry and class separability of the two domains may differ substantially. To quantify asymmetric transfer characteristics, the following difficulty measures are derived from the proposed prototype-based transferability analysis.

- ROC-based difficulty: $D_{ROC}(A \rightarrow B) = 1 - ROC - AUC(A \rightarrow B)$
- PR-based difficulty: $D_{PR}(A \rightarrow B) = 1 - PR - AUC(A \rightarrow B)$
- Score-based difficulty: $D_J(A \rightarrow B) = 1 / J - index(A \rightarrow B)$

where each metric is computed from the prototype-based transferability scores. This formulation captures the intuition that source–target compatibility decreases when source-domain decision structures cannot be preserved under target-domain distributions. Across the evaluated transfer directions, these measures provide complementary empirical characterizations of prototype-based transferability. They should not be interpreted as independent predictors of the final few-shot adaptation performance, which is also influenced by intrinsic target-domain separability and the adaptation process itself.

2.5. Relationship Between Separability and Transferability

A central hypothesis of this work is that intra-domain separability and cross-domain transferability are fundamentally distinct properties. A domain may exhibit strong internal class separation while still remaining difficult to transfer across domains. Conversely, a structurally entangled domain may still contain behavioral patterns that transfer effectively when appropriate semantic representations are used. Consequently, high separability does not inherently imply favorable cross-domain transferability, nor does low separability necessarily imply unfavorable cross-domain transferability. This distinction is particularly important in cross-platform intrusion detection, where heterogeneous logging structures introduce substantial semantic and geometric discrepancies between domains. Based on the proposed prototype-based transferability analysis, the relationship between source–target transferability and the behavior of different few-shot adaptation strategies is empirically investigated under asymmetric cross-domain transfer conditions.

2.6. Few-Shot Domain Adaptation Frameworks

To analyze adaptation behavior under varying transfer difficulty conditions, this study evaluates a diverse set of few-shot domain adaptation (FSDA) frameworks ranging from simple baselines to explicit prototype-based alignment methods. These approaches represent different assumptions regarding domain alignment, adaptation dynamics, and representation transferability.

2.6.1. Zero-Shot Transfer

This baseline evaluates the transferability of source-domain representations without any adaptation. The classifier is trained exclusively on the labeled source-domain data and is directly applied to the target domain without using any target-domain samples during training. The resulting performance reflects the intrinsic generalization capability of the learned representation under cross-domain distribution shift and serves as a reference point for assessing the effectiveness of subsequent few-shot adaptation strategies.

2.6.2. Target-Only Few-Shot Learning

This baseline trains a classifier using only the limited labeled target-domain samples without leveraging any source-domain information. It represents the lower bound of adaptation performance under extremely sparse supervision. Specifically, a binary classifier consisting of a feature extractor and a linear classification head is trained exclusively on the labeled target samples using the standard binary cross-entropy loss.

2.6.3. Naive Joint Training

The entire labeled source-domain dataset and the limited labeled target-domain samples are merged into a single training set without any explicit alignment or adaptation mechanism. A binary classifier is then jointly optimized on the combined dataset using the standard binary cross-entropy loss. This approach relies entirely on implicit representation learning and assumes that a shared feature space can be learned despite the presence of domain discrepancy.

2.6.4. Static Few-Shot Domain Adaptation (Static FSDA)

Static FSDA employs a fixed adaptation objective throughout training, where the relative importance of classification and alignment losses remains unchanged. The model is optimized using the following objective function:

$$L_{static} = \lambda_{src}L_{src} + \lambda_{tgt}L_{tgt} + \lambda_{align}L_{align} + \lambda_{proto}L_{proto} + \lambda_{contra}L_{contra} \quad (3)$$

where L_{src} and L_{tgt} denote the standard binary cross-entropy losses for the source and target domains, respectively. To regularize the latent feature space, several widely adopted domain adaptation components are incorporated. Specifically, L_{align} corresponds to the Correlation Alignment (CORAL) loss [10], which reduces domain discrepancy by aligning the covariance structures of source and target feature representations. In addition, L_{proto} and L_{contra} promote class-level consistency across domains through prototype-based semantic alignment [6] and contrastive regularization [28]. Since all loss weights remain fixed throughout training, the resulting adaptation behavior is static and does not explicitly account for transfer difficulty or domain-specific adaptation dynamics.

2.6.5. Dynamic FSDA

To address the limitations of fixed loss weighting, Dynamic FSDA adaptively adjusts the relative importance of source supervision and domain alignment during training. This framework employs the same objective function as Static FSDA but replaces the fixed loss coefficients with dynamically scheduled weights. Let $p = t/(T - 1)$ denote the normalized training progress, where t and T represent the current and total training epochs, respectively. To progressively shift the training focus from source-domain supervision toward target-domain adaptation, FSDA employs a dynamic weighting mechanism inspired by domain adaptation and self-training frameworks [28]. The coefficients are defined as:

$$\lambda_{src}(p) = 1.0 - 0.5p, \quad \lambda_{tgt}(p) = 0.5 + 0.5p, \quad \lambda_{align}(p) = \gamma \cdot p \quad (4)$$

where γ represents the base alignment coefficient. At the beginning of training ($p \rightarrow 0$), the framework places greater emphasis on source-domain supervision to establish stable feature representations. As training progresses ($p \rightarrow 1$), the influence of target-domain supervision and alignment is gradually increased. This progressive shift encourages the model to transfer knowledge from the source domain while simultaneously adapting to target-domain structures, thereby improving robustness under cross-domain distribution shifts. Unlike Static FSDA, Dynamic FSDA explicitly accounts for the changing adaptation requirements throughout training by continuously rebalancing classification and alignment objectives.

2.6.6. Prototypical Contrastive Alignment (PCA-FSDA)

Building upon prototype-based adaptation methods in few-shot learning and domain adaptation literature [6,18], PCA-FSDA aligns source and target domains through class-level prototype matching and contrastive representation learning. For each class $k \in \{0, 1\}$ (representing normal and attack classes, respectively), class prototypes $C_{src}^{(k)}$ and $C_{tgt}^{(k)}$ are computed as the mean of the extracted features within the current mini-batch:

$$C_{src}^{(k)} = \frac{1}{|X_{src}^{(k)}|} \sum_{x \in X_{src}^{(k)}} f_{\Theta}(x), C_{tgt}^{(k)} = \frac{1}{|X_{tgt}^{(k)}|} \sum_{x \in X_{tgt}^{(k)}} f_{\Theta}(x) \quad (5)$$

where $f_{\Theta}(\cdot)$ denotes the feature extractor.

To encourage semantic consistency across domains, a prototype alignment loss is first defined as:

$$L_{proto} = \frac{1}{2} \sum_{k=0}^1 \|C_{src}^{(k)} - C_{tgt}^{(k)}\|_2^2 \quad (6)$$

which reduces the discrepancy between corresponding source and target class prototypes.

While prototype alignment effectively minimizes intra-class domain discrepancy, it does not explicitly preserve inter-class separability. Consequently, prototypes belonging to different classes may remain insufficiently discriminative in the shared embedding space. To address this limitation, PCA-FSDA incorporates a prototype-level contrastive learning objective L_{contra} , defined in Equation (3). Inspired by recent advances in prototypical contrastive learning [24,29,30], class prototypes are treated as semantic anchors, and cosine similarity-based contrastive learning is employed to simultaneously pull corresponding cross-domain prototypes closer while pushing prototypes of different classes apart. This dual objective promotes both cross-domain semantic consistency and inter-class discrimination, resulting in a more transferable and discriminative representation under sparse target supervision.

2.6.7. Uncertainty-Aware Prototype Refinement (UAPR)

Building upon recent prototype refinement frameworks [31], Uncertainty-Aware Prototype Refinement (UAPR) introduces an uncertainty-guided weighting mechanism to improve prototype reliability under sparse target supervision. This mechanism aims to reduce the influence of noisy or ambiguous target samples and enhances alignment robustness under highly entangled transfer conditions. Under extreme data scarcity, standard prototypes computed using simple empirical means (Equation (6)) can be strongly affected by outlier feature representations, resulting in unstable prototype estimation and degraded domain alignment. To mitigate this issue, UAPR employs the normalized distance between a target sample and its corresponding source-domain prototype as a proxy uncertainty estimate. Specifically, target samples that are located farther from the source prototype are regarded as less reliable, whereas samples closer to the source prototype are assumed to provide more representative class information. For a target sample $x_i \in X_{tgt}^{(k)}$, the uncertainty score is defined as the normalized Euclidean distance to the corresponding source-domain prototype:

$$u_i = \frac{\|f_{\Theta}(x_i) - C_{src}^{(k)}\|_2}{\frac{1}{|X_{tgt}^{(k)}|} \sum_{j \in X_{tgt}^{(k)}} \|f_{\Theta}(x_j) - C_{src}^{(k)}\|_2} \quad (7)$$

where $f_{\Theta}(\cdot)$ denotes the feature extractor and $C_{src}^{(k)}$ represents the source-domain prototype of class (k) .

Based on this uncertainty estimate, an uncertainty-guided weight is assigned to each target sample as $w_i = \exp(-\beta u_i)$, where β is a scaling factor controlling the sensitivity of uncertainty weighting. Consequently, samples with larger prototype distances receive

smaller weights, while samples closer to the source prototype contribute more strongly to prototype estimation. The refined target prototype $\tilde{C}_{tgt}^{(k)}$ is then computed as an uncertainty-weighted mean:

$$\tilde{C}_{tgt}^{(k)} = \frac{\sum_{i \in X_{tgt}^{(k)}} w_i f_{\theta}(x_i)}{\sum_{i \in X_{tgt}^{(k)}} w_i} \quad (8)$$

As a result, uncertain or potentially noisy target samples exert less influence on prototype construction, whereas reliable samples contribute more substantially to the alignment process. This refinement encourages the learned prototypes to better capture representative target-domain structures, thereby improving adaptation robustness under sparse target supervision.

Collectively, these frameworks represent progressively stronger adaptation assumptions, ranging from pure transfer without adaptation to prototype-based alignment and uncertainty-aware refinement. By evaluating these methods under different transfer scenarios, this study empirically analyzes the relationship between prototype-based transferability and few-shot adaptation behavior. The experimental results provide an empirical perspective on how different adaptation methods behave under heterogeneous source–target transfer conditions.

3. Results

This section presents a comprehensive evaluation of cross-platform log anomaly detection. To establish a consistent baseline for this evaluation, the preprocessing pipeline utilized to align and normalize the heterogeneous log sources is first detailed. The subsequent analysis is structured as follows. First, the intrinsic class structures are examined via intra-domain analysis to understand each dataset's characteristics. Second, the cross-domain transferability is evaluated using prototype-based distance scores to characterize the directional asymmetry of source–target transferability. Third, transferability-derived difficulty measures are analyzed to provide complementary perspectives on prototype-based transferability. Fourth, the performance of the proposed framework is evaluated under few-shot adaptation settings to empirically investigate the relationship between prototype-based transferability and few-shot adaptation behavior under heterogeneous transfer conditions. Finally, sensitivity analysis is conducted using PCA-FSDA + UAPR and the training time and inference latency of the evaluated methods are presented.

3.1. Data Preprocessing and Representation

To enable cross-platform few-shot domain adaptation, both ADFA-LD [32] and LMD-2023 [33] were converted into a unified sequence representation prior to feature extraction. ADFA-LD consists of Linux system-call traces, whereas LMD-2023 contains Windows Sysmon events with heterogeneous logging formats. Since the two datasets originate from fundamentally different operating systems, all preprocessing was designed to construct platform-independent behavioral representations while preserving the original temporal structure of each execution trace.

For ADFA-LD, each original execution trace was regarded as one behavioral sequence. Raw numeric system-call identifiers were first converted into their corresponding syscall names to preserve semantic information. A sliding-window strategy with a window size of 64 and stride of 32 was then applied within each trace to generate local behavioral contexts. The resulting window embeddings were aggregated into a single fixed-length representation for the original trace using a sequence aggregation operator. Consequently, each original ADFA-LD trace produced exactly one sequence

embedding, and all windows originating from the same trace remained within the same sample throughout the experiments.

For LMD-2023, events were first grouped according to their *ProcessGuid*, where each *ProcessGuid* represents an individual process execution. Events belonging to the same process were chronologically ordered according to their timestamps, preserving the original execution sequence. Each process was therefore treated as one behavioral sequence. Sliding windows with the same window size (64) and stride (32) were applied only within each *ProcessGuid*, followed by sequence aggregation to produce one embedding per process. No windows were generated across different *ProcessGuid* values, preventing temporal mixing between independent process executions.

Each log event was then embedded into a high-dimensional semantic vector space using the pre-trained BGE-M3 (BAAI General Embedding) model [34] implemented through the SentenceTransformer framework. Unlike template-based log parsers such as Drain [35], which are primarily designed for template extraction within homogeneous logging systems, the objective of this study is to learn transferable semantic representations across heterogeneous platforms. Since Linux system calls and Windows Sysmon logs exhibit substantially different logging formats and event structures, template identifiers alone cannot establish semantic correspondence across domains. Therefore, a pre-trained semantic embedding model was employed to encode heterogeneous log events into a shared semantic embedding space, enabling subsequent sequence aggregation and cross-platform representation alignment.

Three aggregation strategies, namely mean pooling, max pooling, and Top-k pooling, were evaluated during the representation analysis. Based on the preliminary experiments, the aggregation method yielding the best overall performance for each dataset was selected to generate the final sequence embeddings used in the subsequent few-shot domain adaptation experiments.

To avoid information leakage during adaptation, the target-domain few-shot support samples and target-domain evaluation samples were constructed at the sequence level rather than the sliding-window level. Since each sequence corresponds to one original execution trace in ADFA-LD or one *ProcessGuid* in LMD-2023, all windows derived from the same original sequence remain exclusively within either the support subset or the evaluation subset. Consequently, no overlapping windows from the same trace or process appear simultaneously in both subsets.

For the subsequent domain adaptation network, the feature extractor consists of a linear layer projecting the input embedding into a 256-dimensional hidden representation, followed by a ReLU activation and a dropout rate of 0.2. All experiments were conducted on a workstation equipped with an NVIDIA RTX 3500 Ada Generation GPU (12 GB VRAM). The detailed statistics of the preprocessed datasets are summarized in Table 2.

Table 2. Summary and Statistics of the Preprocessed Datasets.

Dataset	Log Source	Granularity	Label Type	Number of Sequences	Normal/Attack Ratio
ADFA-LD	Linux System call traces	Sequence	Binary (Normal vs. Attack)	3852 sequences	3122/730 (4.28:1)
LMD-2023	Windows Sysmon logs	Event → Sequence	Binary (Normal vs. Attack)	16,421 sequences	14,492/1929 (7.51:1)

3.2. Intra-Domain Analysis

As outlined in Section 2.2, the evaluation begins with the intrinsic class structure of each domain. As shown in Table 3, LMD exhibits significantly higher J_{domain} than ADFA,

indicating stronger intrinsic class separability, where J_{domain} is defined as the feature-space dimension-wise Fisher J-index measuring intra-domain class separability. This observation is consistent with the visualization in Figure 1, where LMD samples form more distinct clusters, while ADFA samples exhibit substantial overlap between normal and attack classes. These results confirm that the two domains differ not only in distribution but also in internal structure. Importantly, this analysis characterizes only intra-domain properties and does not directly reflect cross-domain transferability.

Table 3. Intra-domain class separability measured by J_{domain} .

Domain	Number of Normals	Number of Attacks	J_{domain}
ADFA-LD	3122	730	0.059
LMD-2023	14,492	1929	0.903




Figure 1. Cross-Domain Feature Distribution of ADFA and LMD via t-SNE. The visualization demonstrates the distribution shift between the two domains, where LMD exhibits distinct cluster structures whereas ADFA shows substantial overlap between normal and attack classes.

3.3. Cross-Domain Transferability

This section evaluates source-induced transferability using prototype-based distance scores, not classifier-based zero-shot prediction. The purpose of this analysis is to quantify the geometric compatibility between source-domain prototypes and target-domain representations. Let $\mu_{\text{src}, \text{atk}}$ and $\mu_{\text{src}, \text{norm}}$ denote the source-domain attack and normal prototypes, respectively. Following the distance-based assignment principle of Prototypical Networks, a prototype-based transferability score is assigned to each target sample x :

$$s_{\text{proto}}(x) = \frac{d(x, \mu_{\text{src}, \text{norm}})}{d(x, \mu_{\text{src}, \text{atk}}) + d(x, \mu_{\text{src}, \text{norm}})} \quad (9)$$

where $d(\cdot, \cdot)$ denotes the distance metric in the embedding space.

The score $s_{\text{proto}}(x)$ represents the relative position of a target sample with respect to the source-domain attack and normal prototypes. Larger values indicate greater proximity to the source-domain attack prototype, whereas smaller values indicate greater proximity to the source-domain normal prototype. Throughout this study, the attack class (label = 1) is treated as the positive class for ROC-AUC and PR-AUC. These prototype-based metrics quantify how effectively source-derived representations separate target-domain normal and attack samples without performing classifier-based zero-shot prediction or explicit domain adaptation.

Table 4 summarizes the resulting cross-domain transferability. The results reveal a clear directional asymmetry. The transfer direction from LMD→ADFA achieves substantially higher ROC-AUC (0.399) and PR-AUC (0.154) than the reverse direction ADFA→LMD (ROC-AUC = 0.135, PR-AUC = 0.066), indicating better ranking quality and anomaly retrieval performance under the prototype-induced transferability scores. In contrast, the corresponding J-index exhibits the opposite tendency, with ADFA→LMD achieving a substantially larger value (1.285) than LMD→ADFA (0.060). This discrepancy indicates that prototype-based transferability cannot be adequately characterized by a single evaluation metric. While ROC-AUC and PR-AUC evaluate ranking quality and anomaly retrieval performance, respectively, the J-index reflects score-space separability induced by the source-domain prototypes. These complementary observations further support the view that source–target transferability is inherently multi-dimensional. Figure 2 visualizes the resulting score distributions. The two transfer directions generate noticeably different score geometries, suggesting that source-domain structures interact differently with the target-domain data distribution. In the ADFA→LMD direction, the target LMD samples organize into two distinct peaks, confirming that the target behaviors remain inherently distinguishable. In contrast, in the LMD→ADFA direction, the target ADFA samples display a complex multi-modal distribution where the classes are widely scattered into fragmented sub-clusters. This fragmented distribution confirms that the ADFA domain is heavily entangled. These observations motivate a more explicit formulation of transfer difficulty in the following section.

Table 4. Prototype-based source-induced transferability.

Transfer Direction	ROC-AUC (↑)	PR-AUC (↑)	J-index (↑)
ADFA→LMD	0.135	0.066	1.285
LMD→ADFA	0.399	0.154	0.060




Figure 2. Distribution of prototype-based transferability scores for (A) ADFA→LMD and (B) LMD→ADFA. ADFA→LMD exhibits a clear bimodal distribution, whereas LMD→ADFA shows a more entangled multimodal distribution, indicating differences in source-induced cross-domain separability.

3.4. Transferability Characterization

While prior domain adaptation studies have characterized transferability through domain discrepancy or transferability assessment metrics [36,37], these approaches typically summarize transfer behavior using a single discrepancy measure. Such a formulation is often insufficient for intrusion detection, where transferability exhibits strong directional

asymmetry and different evaluation metrics capture distinct aspects of adaptation difficulty. Moreover, intrusion detection is inherently characterized by severe class imbalance, making anomaly retrieval performance as important as overall ranking quality.

Rather than solely quantifying domain discrepancy, this study focuses on characterizing prototype-based transferability under heterogeneous security-log environments. The asymmetric transferability observed in Table 4 further indicates that source–target transferability cannot be adequately described by a single metric. Therefore, three complementary difficulty measures are derived from the proposed prototype-based transferability analysis.

- ROC-based difficulty: $D_{ROC} = 1 - ROC - AUC$;
- PR-based difficulty: $D_{PR} = 1 - PR - AUC$;
- Score-based difficulty: $D_J = 1/J_{transfer}$.

where $J_{transfer}$ denotes the source-induced target score-space separability, quantified by the one-dimensional Fisher J-index computed from the prototype-based target scores. D_{ROC} , D_{PR} , and D_J are computed from the prototype-based transferability scores $S_{proto}(x)$, where the attack class is treated as the positive class.

Each metric captures a different aspect of transfer behavior. D_{ROC} reflects ranking-based difficulty, D_{PR} measures anomaly retrieval difficulty under class imbalance, and D_J quantifies the inverse of score-space separability. Larger values indicate greater difficulty with respect to the corresponding transferability metric. Consequently, the three measures do not necessarily produce identical assessments of transfer difficulty. As shown in Table 5, the two transfer directions exhibit substantially different transfer characteristics. The ROC-based and PR-based difficulties indicate that the two transfer directions exhibit different ranking- and retrieval-based transfer characteristics. In contrast, the score-space separability metric (D_J) highlights a different aspect of prototype-based transferability, reflecting geometric separation in the induced score space.

Table 5. Transferability-derived characteristics. Larger values indicate greater transfer difficulty with respect to the corresponding prototype-based transferability metric.

Transfer Direction	D_{ROC}	D_{PR}	D_J
ADFA→LMD	0.865	0.934	0.778
LMD→ADFA	0.601	0.846	16.709

These observations indicate that prototype-based transferability should be interpreted through multiple complementary perspectives rather than a single scalar quantity. Rather than relying on a single scalar quantity, source–target transferability should be analyzed through multiple complementary perspectives that capture ranking quality, anomaly retrieval capability, and score-space geometry. These complementary metrics characterize different aspects of prototype-based transferability and should not be interpreted as independent predictors of the final few-shot adaptation performance.

3.5. Empirical Analysis of Transfer Characteristics

To examine the practical implications of the proposed prototype-based transferability analysis, the two evaluated transfer directions are compared from both prototype-based compatibility and few-shot adaptation perspectives. Figure 3 compares the prototype-based PR difficulty (D_{PR}) with the corresponding few-shot PR-AUC achieved by the three evaluated adaptation methods (Target-only FS, PCA-FSDA, and PCA-FSDA + UAPR) under the two transfer directions. Because intrusion detection is highly imbalanced, the PR-based transferability metric (D_{PR}) is treated as the primary indicator throughout the subsequent analysis, whereas D_{ROC} and D_J are reported as comple-

mentary transferability measures. Although the adaptation methods exhibit different performance characteristics across the evaluated transfer scenarios, these differences should not be interpreted as evidence that prototype-based transferability independently determines adaptation performance.




Figure 3. Prototype-based PR difficulty versus few-shot PR-AUC under the two evaluated transfer directions.

The two transfer directions exhibit distinct prototype-based transferability characteristics as well as different adaptation behaviors. The prototype-based metrics quantify the geometric compatibility between source-domain prototypes and target-domain representations prior to adaptation, whereas the few-shot results reflect the combined effects of source–target compatibility, intrinsic target-domain separability, and the adaptation process itself. Consequently, differences observed in the prototype-based transferability metrics do not necessarily translate directly into differences in final few-shot detection performance.

These observations provide an empirical perspective for interpreting cross-domain transfer characteristics under the two evaluated transfer scenarios. However, since the current study considers only two cross-platform transfer directions, the results should be interpreted as a two-direction empirical case study rather than evidence of a causal relationship between the proposed prototype-based transferability metrics and few-shot adaptation performance. Further validation using additional source–target combinations will be valuable for assessing the generality of the proposed transferability perspective.

3.6. Few-Shot Performance Under Different Transfer Conditions

Based on the transferability analysis presented in Sections 3.3–3.5, the variation in few-shot adaptation performance under the two evaluated transfer directions is subsequently investigated. Because intrusion detection is characterized by severe class imbalance, the PR-based difficulty metric D_{PR} is used as a descriptive indicator of the observed transfer conditions. Accordingly, the ADFA→LMD direction is referred to as the relatively favorable transfer condition, whereas LMD→ADFA corresponds to the more challenging transfer condition within the scope of the evaluated datasets. Few-shot adaptation performance is evaluated using varying numbers of labeled target samples (1, 5, 10, 20, and 50 shots per class). The complete results are provided in Appendix A, while the performance trajectories are visualized in Figure 4. Specifically, Table 6 summarizes the final performance attained at 50 shots per class.




Figure 4. Few-shot adaptation under asymmetric transfer difficulty. Few-shot adaptation performance (PR-AUC) across varying numbers of labeled target samples (1, 5, 10, 20, and 50 shots per class) under asymmetric transfer conditions. In the low-difficulty ADFA→LMD setting (A), Target-only FS rapidly approaches its maximum achievable performance with minimal target supervision. In contrast, under the high-difficulty LMD→ADFA setting (B), explicit domain adaptation provides substantial and persistent gains across all shot configurations.

Table 6. Few-shot adaptation performance under different transfer conditions (50-shot setting). Paired Wilcoxon signed-rank tests across five random seeds indicated that the observed performance differences between the compared methods were not statistically significant (all $p > 0.05$).

Transfer	Difficulty	Method	PR-AUC	F1	ROC-AUC
ADFA→LMD	Low	Target-only FS	0.965 ± 0.003	0.928 ± 0.022	0.997 ± 0.000
ADFA→LMD	Low	PCA-FSDA	0.937 ± 0.023	0.887 ± 0.042	0.995 ± 0.001
ADFA→LMD	Low	PCA-FSDA + UAPR	0.960 ± 0.002	0.904 ± 0.040	0.996 ± 0.000
LMD→ADFA	High	Target-only FS	0.442 ± 0.040	0.464 ± 0.055	0.777 ± 0.018
LMD→ADFA	High	PCA-FSDA	0.579 ± 0.064	0.643 ± 0.039	0.882 ± 0.025
LMD→ADFA	High	PCA-FSDA + UAPR	0.544 ± 0.066	0.620 ± 0.028	0.867 ± 0.027

In the ADFA→LMD transfer direction, Target-only FS already achieves strong performance, reaching a PR-AUC of 0.965 at 50 shots. As illustrated in Figure 4, performance improves rapidly even with a small number of labeled target samples, indicating that the target-domain representations can be effectively utilized with limited supervision.

In contrast, the LMD→ADFA transfer direction exhibits considerably lower adaptation performance. As shown in Figure 4, Target-only FS remains limited throughout the evaluated shot range. Although performance improves with additional labeled target samples, the resulting PR-AUC remains relatively low (0.442 at 50 shots). Under this transfer direction, PCA-FSDA consistently outperforms the Target-only baseline across all evaluated shot settings. By incorporating transferable source-domain information together with prototype-level alignment, PCA-FSDA achieves a PR-AUC of 0.579 and an F1-score of 0.643 at 50 shots. PCA-FSDA and PCA-FSDA + UAPR exhibited comparable performance in the LMD→ADFA transfer direction.

Overall, the two evaluated transfer directions exhibit different adaptation behaviors. In the relatively lower-difficulty transfer setting, Target-only FS achieves competitive performance with limited target supervision, whereas prototype-based alignment methods provide larger improvements under the relatively higher-difficulty transfer setting. These

observations are consistent with the proposed transferability perspective; however, because the present evaluation is limited to two cross-platform transfer directions, the results should be interpreted as empirical observations rather than evidence that transfer difficulty alone determines adaptation performance. Additional validation using a broader range of source–target combinations will be valuable for assessing the generality of these observations.

3.7. Ablation Study for Smoothing Coefficient and Uncertainty Scaling Factor β

The sensitivity analysis was conducted using PCA-FSDA+UAPR at 50 shots per class because the uncertainty scaling factor and smoothing coefficient (i.e., the minimum uncertainty weight) are specific to the UAPR module. Here, the minimum uncertainty weight acts as a smoothing coefficient that prevents the weights of highly uncertain target samples from dropping completely to zero. Table 7 presents the sensitivity analysis of the uncertainty scaling factor β and the minimum uncertainty weight. Across both transfer directions, the proposed framework remains relatively stable over a wide range of β values, indicating that UAPR is not highly sensitive to the scaling coefficient. In contrast, setting the minimum uncertainty weight to zero degrades performance, suggesting that completely suppressing uncertain target samples removes useful adaptation information. Moderate nonzero minimum weights produce stable and competitive performance across both transfer directions.

Table 7. Sensitivity analysis of UAPR using PR-AUC (50-shot, mean $\pm$ std). PR-AUC is reported because it is the primary evaluation metric for highly imbalanced intrusion detection tasks.

Parameter	Value	ADFA→LMD	LMD→ADFA
β	1	0.9475 $\pm$ 0.0063	0.5403 $\pm$ 0.0369
	2	0.9521 $\pm$ 0.0086	0.5558 $\pm$ 0.0329
	4	0.9603 $\pm$ 0.0046	0.5526 $\pm$ 0.0423
	8	0.9601 $\pm$ 0.0050	0.5483 $\pm$ 0.0510
	16	0.9601 $\pm$ 0.0050	0.5480 $\pm$ 0.0495
smoothing	0.00	0.9220 $\pm$ 0.0049	0.3731 $\pm$ 0.0805
	0.05	0.9598 $\pm$ 0.0044	0.5516 $\pm$ 0.0591
	0.10	0.9601 $\pm$ 0.0050	0.5483 $\pm$ 0.0510
	0.20	0.9567 $\pm$ 0.0059	0.5596 $\pm$ 0.0347
	0.30	0.9466 $\pm$ 0.0103	0.5859 $\pm$ 0.0301

3.8. Time Complexity

Table 8 summarizes the computational cost of the evaluated methods in terms of training time and inference latency. The computational overhead was measured at 50 shots per class, which corresponds to the main experimental setting used for method comparison. Inference latency was measured using batched inference over the target test set and is reported as the average latency per sample.

As expected, prototype alignment and uncertainty-aware refinement increase the training cost because additional prototype computation and sample-wise uncertainty weighting are performed during optimization. Nevertheless, the increase is confined to the training stage. During inference, all methods exhibit nearly identical latency, as the proposed prototype refinement and uncertainty estimation are not required after training. These results indicate that the proposed framework introduces only a modest training overhead while maintaining efficient inference, making it suitable for real-time cross-domain intrusion detection once the adaptation process has been completed.

Table 8. Training time and inference latency under the two evaluated transfer directions (50 shots per class).

Method	ADFA→LMD		LMD→ADFA	
	Training Time (s)	Inference Latency (ms/Sample)	Training Time (s)	Inference Latency (ms/Sample)
Zero-shot	5.490 ± 0.553	0.0011 ± 0.0001	22.363 ± 0.598	0.0012 ± 0.0001
Target-only FS	0.563 ± 0.014	0.0010 ± 0.0001	0.450 ± 0.156	0.0012 ± 0.0001
Naive Joint	5.677 ± 0.816	0.0010 ± 0.0000	22.595 ± 1.016	0.0012 ± 0.0002
Static FSDA	14.077 ± 0.358	0.0012 ± 0.0004	59.517 ± 0.900	0.0019 ± 0.0016
Dynamic FSDA	14.047 ± 0.321	0.0010 ± 0.0000	59.138 ± 0.736	0.0012 ± 0.0001
PCA-FSDA	45.513 ± 1.671	0.0009 ± 0.0001	180.793 ± 7.503	0.0015 ± 0.0005
PCA-FSDA + UAPR	82.463 ± 0.403	0.0010 ± 0.0002	336.788 ± 8.128	0.0012 ± 0.0001

Since the LMD dataset contains substantially more source samples than ADFA-LD, the training time naturally increases for the LMD→ADFA transfer direction due to the additional prototype computation and optimization over a larger source dataset. However, inference latency remains nearly unchanged because the proposed prototype alignment and uncertainty-aware refinement are only applied during training and do not introduce additional computations during deployment.

4. Discussion

In this work, few-shot domain adaptation (FSDA) is revisited from a transfer difficulty-aware perspective to analyze how different adaptation strategies behave under two cross-domain intrusion detection scenarios. A key observation is that intra-domain separability and cross-domain transferability represent different characteristics of a dataset. Although the LMD dataset exhibits substantially higher intrinsic class separability than ADFA, this property alone does not necessarily correspond to improved cross-domain adaptation performance. The experimental results indicate that a well-separated target-domain representation may still exhibit limited compatibility with source-domain representations, resulting in reduced cross-domain performance. These observations suggest that evaluating only the internal structure of a dataset is insufficient for characterizing cross-domain adaptation behavior. To complement conventional domain discrepancy analysis, asymmetric transferability metrics are introduced to quantify directional differences between the evaluated transfer settings. Unlike traditional discrepancy measures, these metrics explicitly consider the asymmetric nature of source-to-target transfer by measuring how well source-derived decision structures generalize to the target domain. The experimental analysis further demonstrates that the two evaluated transfer directions exhibit substantially different transfer characteristics despite involving the same pair of datasets, highlighting the inherently asymmetric nature of cross-domain transfer.

The experimental results also demonstrate that different few-shot adaptation methods exhibit different behaviors across the two evaluated transfer directions. Under the relatively more favorable transfer condition (ADFA→LMD), Target-only FS already achieves competitive performance, suggesting that limited labeled target samples can effectively support adaptation. In contrast, under the relatively challenging transfer condition (LMD→ADFA), methods that incorporate transferable source-domain information, particularly prototype-alignment strategies such as PCA-FSDA, provide larger performance improvements over the Target-only baseline. These observations indicate that the relative effectiveness of

adaptation strategies may vary depending on the transfer characteristics of a particular source–target pair.

Although slight performance differences were observed between adaptation strategies under certain transfer configurations, paired Wilcoxon signed-rank tests indicated that these differences were not statistically significant (all $p > 0.05$). These findings further support the conclusion that adaptation effectiveness depends primarily on source–target transferability rather than exhibiting consistent superiority across all adaptation strategies.

The proposed UAPR module introduces only a lightweight computational overhead because uncertainty estimation requires only distance computations between target features and the corresponding source prototypes. This operation scales linearly with the number of target samples and does not affect the inference procedure, where only the trained feature extractor and classifier are used. Consequently, the additional computational cost is incurred only during training and remains modest in practice.

Although UAPR generally is designed to improve prototype reliability by suppressing uncertain target samples, its effectiveness depends on the quality of the source-domain prototypes used as references. Under the highly challenging transfer conditions, the source and target distributions become substantially separated, causing a large proportion of target samples to receive low confidence weights. As a result, fewer target samples contribute effectively to prototype refinement, slightly reducing adaptation effectiveness. Nevertheless, the observed performance difference between PCA-FSDA and PCA-FSDA + UAPR remains relatively small compared to the variation across random seeds, suggesting that the degradation is limited rather than systematic.

Across both transfer directions, UAPR exhibits relatively low sensitivity to the uncertainty scaling factor β , with performance remaining stable over a broad range of values. In contrast, setting the minimum uncertainty weight to zero consistently degrades performance in both transfer scenarios, indicating that completely suppressing uncertain target samples removes useful adaptation information. These results suggest that UAPR is relatively insensitive to the choice of β while requiring only a moderate nonzero lower bound on uncertainty weights.

From a practical perspective, these findings suggest that transfer characteristics may provide useful information when selecting adaptation strategies for cross-domain intrusion detection. Simpler few-shot learning approaches may be sufficient when the source and target domains exhibit relatively favorable transfer characteristics, whereas alignment-based approaches appear to offer greater benefits under more challenging transfer conditions. However, because the present study evaluates only two cross-platform transfer directions, these observations should be interpreted as empirical findings rather than general evidence that transfer difficulty alone determines adaptation performance. Additional studies involving a broader range of source–target combinations will be necessary to further validate the generality of this perspective and to investigate whether transferability characteristics can support adaptive strategy selection.

Although the experiments were conducted on two heterogeneous operating system log datasets to emphasize cross-platform transfer, evaluating the proposed framework on additional cloud-native, distributed-system, and enterprise-scale log datasets remains an important direction for future research.

5. Conclusions

In this paper, a transferability-aware analysis of few-shot domain adaptation is presented for cross-domain intrusion detection. Rather than solely introducing a new adaptation algorithm, this study investigates how different few-shot adaptation strategies behave under heterogeneous cross-platform transfer scenarios. The experimental analysis shows that intra-

domain separability and cross-domain transferability represent different properties of a dataset, and that strong intra-domain class separation does not necessarily correspond to favorable cross-domain adaptation performance. To characterize these differences, asymmetric prototype-based transferability metrics are introduced to characterize directional transfer characteristics based on cross-domain generalization and J-index scaling.

Experiments on two cross-platform transfer directions demonstrate that the evaluated transfer settings exhibit substantially different adaptation behavior. Across the two evaluated transfer directions, the examined few-shot adaptation methods exhibit different adaptation behaviors. While Target-only few-shot learning achieves competitive performance in the ADFA→LMD transfer direction, alignment-based approaches provide larger improvements in the LMD→ADFA transfer direction. These observations suggest that prototype-based transferability analysis may provide useful insights for interpreting adaptation behavior under heterogeneous cross-domain transfer scenarios.

Because the present study evaluates two heterogeneous operating system datasets and two transfer directions, the reported findings should be interpreted as empirical observations rather than general evidence that prototype-based transferability should not be interpreted as the sole factor governing adaptation performance. Extending the analysis to additional operating systems, cloud-native environments, enterprise log repositories, and other source–target combinations will be an important direction for future work. Such studies will help further assess the generality of the proposed transferability-aware perspective and its potential usefulness for practical cross-domain intrusion detection.

As future work, the proposed framework may be extended to adaptive security systems that estimate transfer characteristics in real time and investigate how such information can assist adaptation strategy selection in operational cross-platform environments.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT: Ministry of Science and ICT) under Grant NRF-2019R1F1A1059036.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Publicly available datasets were analyzed in this study. These data can be found in reference [31,32]. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Extended Classifier-Based Few-Shot Adaptation Results

Values represent the Mean $\pm$ Standard Deviation computed over multiple random seeds.

Table A1. Detailed Few-shot Performance under Low-Difficulty Transfer (ADFA $\rightarrow$ LMD).

Method	PR_AUC (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.300 $\pm$ 0.025	0.300 $\pm$ 0.025	0.300 $\pm$ 0.025	0.299 $\pm$ 0.025	0.296 $\pm$ 0.025
Target-only few-shot	0.814 $\pm$ 0.199	0.954 $\pm$ 0.018	0.957 $\pm$ 0.014	0.959 $\pm$ 0.011	0.965 $\pm$ 0.003
Naive joint	0.548 $\pm$ 0.152	0.890 $\pm$ 0.053	0.938 $\pm$ 0.011	0.953 $\pm$ 0.006	0.953 $\pm$ 0.011
Static FSDA	0.782 $\pm$ 0.225	0.924 $\pm$ 0.047	0.944 $\pm$ 0.017	0.946 $\pm$ 0.015	0.951 $\pm$ 0.017
Dynamic FSDA	0.779 $\pm$ 0.235	0.925 $\pm$ 0.047	0.945 $\pm$ 0.017	0.945 $\pm$ 0.015	0.951 $\pm$ 0.017
PCA-FSDA	0.773 $\pm$ 0.233	0.918 $\pm$ 0.052	0.933 $\pm$ 0.011	0.935 $\pm$ 0.012	0.937 $\pm$ 0.023
PCA-FSDA + UAPR	0.768 $\pm$ 0.227	0.923 $\pm$ 0.043	0.949 $\pm$ 0.010	0.955 $\pm$ 0.005	0.960 $\pm$ 0.002

Table A1. Cont.

Method	F1 (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.316 ± 0.088	0.335 ± 0.072	0.305 ± 0.082	0.302 ± 0.074	0.344 ± 0.028
Target-only few-shot	0.212 ± 0.002	0.612 ± 0.149	0.821 ± 0.078	0.834 ± 0.068	0.928 ± 0.022
Naive joint	0.326 ± 0.068	0.734 ± 0.094	0.831 ± 0.059	0.833 ± 0.063	0.881 ± 0.045
Static FSDA	0.681 ± 0.039	0.796 ± 0.073	0.844 ± 0.070	0.838 ± 0.067	0.894 ± 0.046
Dynamic FSDA	0.729 ± 0.014	0.796 ± 0.073	0.839 ± 0.073	0.834 ± 0.068	0.894 ± 0.045
PCA-FSDA	0.442 ± 0.153	0.747 ± 0.072	0.819 ± 0.072	0.809 ± 0.055	0.887 ± 0.042
PCA-FSDA + UAPR	0.283 ± 0.069	0.717 ± 0.051	0.814 ± 0.066	0.817 ± 0.062	0.904 ± 0.040
Method	ROC_AUC (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.782 ± 0.022	0.782 ± 0.023	0.782 ± 0.023	0.782 ± 0.023	0.782 ± 0.022
Target-only few-shot	0.977 ± 0.026	0.995 ± 0.003	0.995 ± 0.003	0.995 ± 0.002	0.997 ± 0.000
Naive joint	0.911 ± 0.042	0.986 ± 0.007	0.992 ± 0.002	0.995 ± 0.001	0.996 ± 0.001
Static FSDA	0.972 ± 0.029	0.991 ± 0.008	0.994 ± 0.002	0.995 ± 0.002	0.996 ± 0.001
Dynamic FSDA	0.971 ± 0.031	0.992 ± 0.007	0.994 ± 0.002	0.995 ± 0.002	0.996 ± 0.001
PCA-FSDA	0.969 ± 0.031	0.990 ± 0.010	0.993 ± 0.003	0.994 ± 0.002	0.995 ± 0.001
PCA-FSDA + UAPR	0.969 ± 0.030	0.991 ± 0.006	0.994 ± 0.003	0.995 ± 0.001	0.996 ± 0.000

Table A2. Detailed Few-shot Performance under High-Difficulty Transfer (LMD → ADFA).

Method	PR_AUC (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.152 ± 0.003	0.152 ± 0.002	0.151 ± 0.002	0.149 ± 0.002	0.146 ± 0.003
Target-only few-shot	0.334 ± 0.047	0.365 ± 0.082	0.409 ± 0.042	0.432 ± 0.023	0.442 ± 0.040
Naive joint	0.224 ± 0.046	0.319 ± 0.072	0.404 ± 0.030	0.460 ± 0.038	0.489 ± 0.038
Static FSDA	0.264 ± 0.063	0.312 ± 0.079	0.356 ± 0.041	0.445 ± 0.046	0.521 ± 0.047
Dynamic FSDA	0.268 ± 0.062	0.317 ± 0.082	0.362 ± 0.045	0.444 ± 0.052	0.527 ± 0.042
PCA-FSDA	0.265 ± 0.061	0.306 ± 0.087	0.393 ± 0.025	0.467 ± 0.051	0.579 ± 0.064
PCA-FSDA + UAPR	0.249 ± 0.052	0.294 ± 0.096	0.365 ± 0.048	0.467 ± 0.035	0.544 ± 0.066
Method	F1 (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.258 ± 0.056	0.317 ± 0.000	0.316 ± 0.000	0.313 ± 0.003	0.302 ± 0.009
Target-only few-shot	0.311 ± 0.010	0.340 ± 0.055	0.415 ± 0.026	0.464 ± 0.038	0.464 ± 0.055
Naive joint	0.263 ± 0.030	0.333 ± 0.046	0.394 ± 0.057	0.498 ± 0.039	0.569 ± 0.036
Static FSDA	0.279 ± 0.047	0.364 ± 0.059	0.433 ± 0.034	0.513 ± 0.050	0.619 ± 0.032
Dynamic FSDA	0.288 ± 0.039	0.372 ± 0.062	0.436 ± 0.033	0.515 ± 0.051	0.620 ± 0.031
PCA-FSDA	0.291 ± 0.020	0.364 ± 0.051	0.442 ± 0.033	0.531 ± 0.048	0.643 ± 0.039
PCA-FSDA + UAPR	0.299 ± 0.016	0.367 ± 0.064	0.433 ± 0.028	0.513 ± 0.041	0.620 ± 0.028

Table A2. Cont.

Method	ROC_AUC (shots_per_class)				
	1	5	10	20	50
Zero-shot	0.364 ± 0.003	0.364 ± 0.002	0.364 ± 0.002	0.363 ± 0.003	0.364 ± 0.004
Target-only few-shot	0.629 ± 0.029	0.622 ± 0.033	0.685 ± 0.042	0.725 ± 0.038	0.777 ± 0.018
Naive joint	0.425 ± 0.061	0.567 ± 0.068	0.704 ± 0.037	0.790 ± 0.012	0.841 ± 0.015
Static FSDA	0.508 ± 0.060	0.574 ± 0.058	0.717 ± 0.019	0.805 ± 0.025	0.867 ± 0.026
Dynamic FSDA	0.518 ± 0.058	0.574 ± 0.058	0.720 ± 0.018	0.805 ± 0.026	0.869 ± 0.024
PCA-FSDA	0.534 ± 0.052	0.579 ± 0.070	0.731 ± 0.018	0.819 ± 0.023	0.882 ± 0.025
PCA-FSDA + UAPR	0.519 ± 0.039	0.572 ± 0.077	0.720 ± 0.020	0.812 ± 0.016	0.867 ± 0.027

References

- Wang, Y.; Mäntylä, M.V.; Nyyssölä, J.; Ping, K.; Wang, L. Cross-System Software Log-based Anomaly Detection Using Meta-Learning. In Proceedings of the IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER), Montreal, QC, Canada, 4–7 March 2025; pp. 454–464. [\[CrossRef\]](#)
- Duan, C.; He, M.; Xiao, P.; Jia, T.; Zhang, X.; Zhong, Z.; Luo, X.; Niu, Y.; Zhang, L.; Yu, S.; et al. LogAction: Consistent Cross-system Anomaly Detection through Logs via Active Domain Adaptation. In Proceedings of the 40th IEEE/ACM International Conference on Automated Software Engineering (ASE), Seoul, Republic of Korea, 16–20 November 2025. [\[CrossRef\]](#)
- Pospieszny, P.; Mormul, W.; Szyndler, K.; Kumar, S. ADALog: Adaptive unsupervised anomaly detection in logs with self-attention masked language model. In Proceedings of the 10th International Conference on Machine Learning Technologies (ICMLT), Helsinki, Finland, 23–25 May 2025; pp. 248–256. [\[CrossRef\]](#)
- Zhao, X.; Jia, T.; He, M.; Yang, X.; Li, Y. FusionLog: Cross-System Log-based Anomaly Detection via Fusion of General and Proprietary Knowledge. *arXiv* **2025**, arXiv:2511.05878. [\[CrossRef\]](#)
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Kavukcuoglu, K.; Wierstra, D. Matching networks for one shot learning. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Barcelona, Spain, 5–10 December 2016; Volume 29, pp. 3630–3638.
- Snell, J.; Swersky, K.; Zemel, R.S. Prototypical networks for few-shot learning. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS), Long Beach, CA, USA, 4–9 December 2017; pp. 4080–4090.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P.H.S.; Hospedales, T.M. Learning to compare: Relation network for few-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–22 June 2018; pp. 1199–1208. [\[CrossRef\]](#)
- Finn, C.; Abbeel, P.; Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In Proceedings of the International Conference on Machine Learning (ICML), Sydney, Australia, 6–11 August 2017; pp. 1126–1135.
- Long, M.; Cao, Y.; Wang, J.; Jordan, M.I. Learning transferable features with deep adaptation networks. In Proceedings of the International Conference on Machine Learning (ICML), Lille, France, 6–11 July 2015; pp. 97–105.
- Sun, B.; Saenko, K. Deep CORAL: Correlation alignment for deep domain adaptation. In Proceedings of the European Conference on Computer Vision (ECCV), Amsterdam, The Netherlands, 11–14 October 2016; pp. 443–450. [\[CrossRef\]](#)
- Ganin, Y.; Lempitsky, V. Unsupervised domain adaptation by backpropagation. In Proceedings of the International Conference on Machine Learning (ICML), Lille, France, 6–11 July 2015; pp. 1180–1189.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; March, M.; Lempitsky, V. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* **2016**, *17*, 1–35.
- Long, M.; Cao, Z.; Wang, J.; Jordan, M.I. Conditional adversarial domain adaptation. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Montréal, QC, Canada, 3–8 December 2018; Volume 31, pp. 1647–1657.
- Saito, K.; Watanabe, K.; Ushiku, Y.; Harada, T. Maximum classifier discrepancy for unsupervised domain adaptation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–22 June 2018; pp. 3723–3732. [\[CrossRef\]](#)
- Zhang, Y.; Liu, T.; Long, M.; Jordan, M.I. Bridging theory and algorithm for domain adaptation. In Proceedings of the International Conference on Machine Learning (ICML), Long Beach, CA, USA, 9–15 June 2019; pp. 7404–7413.
- Wang, J.; Chen, Y.; Hao, S.; Feng, W.; Shen, Z. Balanced distribution adaptation for transfer learning. In Proceedings of the 2017 IEEE International Conference on Data Mining (ICDM), New Orleans, LA, USA, 18–21 November 2017; pp. 1129–1134. [\[CrossRef\]](#)
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; Vaughan, J.W. A theory of learning from different domains. *Mach. Learn.* **2010**, *79*, 151–175. [\[CrossRef\]](#)

18. Pan, Y.; Yao, T.; Li, Y.; Wang, Y.; Ngo, C.-W.; Mei, T. Transferrable prototypical networks for unsupervised domain adaptation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16–20 June 2019; pp. 2234–2242.
19. Wang, H.; Deng, Z.-H. Cross-domain few-shot classification via adversarial task augmentation. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Virtually, 19–26 August 2021; pp. 1075–1081. [[CrossRef](#)] [[PubMed](#)]
20. Teshima, T.; Sato, I.; Sugiyama, M. Few-shot domain adaptation by causal mechanism transfer. In Proceedings of the International Conference on Machine Learning (ICML), Vienna, Austria, 12–18 July 2020; pp. 9458–9469.
21. Han, X.; Yuan, S. Unsupervised cross-system log anomaly detection via domain adaptation. In Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM), Queensland, Australia, 1–5 November 2021; pp. 3068–3072. [[CrossRef](#)]
22. Sun, Y.; Keung, J.W.; Yu, H.K.; Luo, W. LogMeta: A few-shot model-agnostic meta-learning framework for robust and adaptive log anomaly detection. *J. Syst. Softw.* **2026**, *235*, 112781. [[CrossRef](#)]
23. Sui, Y.; Wang, X.; Cui, T.; Xiao, T.; He, C.; Zhang, S.; Zhang, Y.; Yang, X.; Sun, Y.; Pei, D. Bridging the Gap: LLM-Powered transfer learning for log anomaly detection in new software systems. In Proceedings of the 41st IEEE International Conference on Data Engineering (ICDE), Hong Kong, China, 19–23 May 2025; pp. 4414–4427. [[CrossRef](#)]
24. Yue, X.; Zheng, Z.; Zhang, S.; Gao, Y.; Darrell, T.; Keutzer, K.; Vincentelli, A.S. Prototypical cross-domain self-supervised learning for few-shot unsupervised domain adaptation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Virtually, 19–25 June 2021; pp. 13829–13839. [[CrossRef](#)]
25. Zhou, J.; Ying, S.; Wang, S.; Zhao, D.; Xiang, J.; Liang, K.; Liu, P. LogDLR: Unsupervised cross-system log anomaly detection through domain-invariant latent representation. In *IEEE Transactions on Dependable and Secure Computing*; IEEE: Piscataway, NJ, USA, 2025; Volume 22, pp. 4456–4471. [[CrossRef](#)]
26. Ye, J.; Wang, Z.; Su, C.; Yang, J.; Ding, J.; Liu, C.; Chu, G. LogICL: Distilling LLM reasoning to bridge the semantic gap in cross-domain log anomaly detection. *arXiv* **2025**, arXiv:2512.09627.
27. Gu, Q.; Li, Z.; Han, J. Generalized Fisher score for feature selection. In Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI), Barcelona, Spain, 14–17 July 2011; pp. 266–273. [[CrossRef](#)]
28. Chen, T.; Kornblith, S.; Norouzi, M.; Hinton, G. A simple framework for contrastive learning of visual representations. In Proceedings of the 37th International Conference on Machine Learning (ICML), Vienna, Austria, 12–18 July 2020; pp. 1597–1607.
29. Jiang, Z.; Li, Y.; Yang, C.; Gao, P.; Wang, Y.; Tai, Y.; Wang, C. Prototypical contrast adaptation for domain adaptive semantic segmentation. In Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel, 23–27 October 2022; pp. 36–54. [[CrossRef](#)]
30. Wang, K.; Liew, J.H.; Zou, Y.; Zhou, D.; Feng, J. PANet: Few-shot image semantic segmentation with prototype alignment. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Republic of Korea, 29 October–1 November 2019; pp. 9197–9206.
31. Wei, L.; Lang, C.; Chen, Z.; Wang, T.; Li, Y.; Liu, J. Generated and pseudo content guided prototype refinement for few-shot point cloud segmentation. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Montreal, QC, Canada, 8–13 December 2024; Volume 37, pp. 31103–31123. [[CrossRef](#)]
32. Creech, G.; Hu, J. Generation of a new IDS test dataset: Time to retire the KDD collection. In Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), Shanghai, China, 7–10 April 2013; pp. 4487–4492. [[CrossRef](#)]
33. Smiliotopoulos, C.; Kambourakis, G.; Barbatsalou, K. On the detection of lateral movement through supervised machine learning and an open-source tool to create turnkey datasets from Sysmon logs. *Int. J. Inf. Secur.* **2023**, *22*, 1893–1919. [[CrossRef](#)]
34. Chen, J.; Xiao, S.; Zhang, P.; Luo, K.; Lian, D.; Liu, Z. BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation. *arXiv* **2024**, arXiv:2402.03216. [[CrossRef](#)]
35. He, P.; Zhu, J.; Zheng, Z.; Lyu, M.R. Drain: An online log parsing approach with fixed depth tree. In Proceedings of the 2017 IEEE International Conference on Web Services (ICWS), Honolulu, HI, USA, 25–30 June 2017; pp. 33–40. [[CrossRef](#)]
36. Ben-David, S.; Blitzer, J.; Crammer, K.; Pereira, F. Analysis of representations for domain adaptation. In *Advances in Neural Information Processing Systems (NIPS): Proceedings of the 2006 Conference, Cambridge, MA, USA*; MIT Press: Cambridge, MA, USA, 2007; pp. 137–144.
37. Wang, Z.; Dai, Z.; Póczos, B.; Carbonell, J. Characterizing and avoiding negative transfer. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 16–20 June 2019; pp. 11285–11294. [[CrossRef](#)]

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