

# 000 DOMAIN-AWARE GRADIENT REUSE FOR ANOMALY 001 DETECTION 002

003 **Anonymous authors**

004 Paper under double-blind review

## 005 ABSTRACT

006 Anomaly detection relies on recognizing patterns that diverge from normal be-  
007 havior, yet practical deployment is hampered by the inherent scarcity and hetero-  
008 geneity of anomalous instances. These challenges prevent the training set from  
009 faithfully characterizing the underlying anomaly distribution, thereby fundamen-  
010 tally constraining the development of effective discriminative models for anomaly  
011 detection. Inspired by the observed consistency of gradient distributions across  
012 related domains during training, Domain-Aware Gradient Reuse (DAGR) is in-  
013 troduced as a transfer-learning framework that leverages this property. DAGR  
014 first learns an adaptive transformation by aligning source and target normal gra-  
015 dients, thereby neutralizing domain-specific effects. The same map then pushes  
016 forward the source anomalous gradients to computing estimated target anomalous  
017 gradients, which are combined with the true target normal gradients to guide the  
018 target-domain detector without labeled anomalies. This paper establishes a rig-  
019 orous convergence proof that reinforces the framework’s theoretical foundation.  
020 Comprehensive experiments on image and audio datasets demonstrate that the  
021 proposed method achieves state-of-the-art performance.

## 022 1 INTRODUCTION

023 Anomaly detection flags observations that deviate from the normal data manifold, underpinning ap-  
024 plications such as automated fraud mitigation, early medical diagnosis, and industrial fault predic-  
025 tion. The surge in data volume and complexity therefore demands models with high representational  
026 capacity and robustness to distribution shifts, rendering deep neural networks the prevailing solution.

027 Although deep learning has advanced rapidly, anomaly detection is still impeded by two factors.  
028 First, the scarcity of anomalous samples leads to severe class imbalance. Second, the heterogeneity  
029 of anomalies ensures that any finite dataset represents only a small fraction of the anomaly space.  
030 Together, these limitations prevent the training data from accurately representing the underlying  
031 anomaly distribution and, in turn, hinder the convergence of discriminative models.

032 Prior work in anomaly detection spans supervised and unsupervised paradigms. Supervised meth-  
033 ods address class imbalance and anomaly sparsity/heterogeneity via reweighting, augmentation, or  
034 generative synthesis; however, synthetic anomalies cover limited modes and promote overfitting,  
035 yielding poor open-set generalisation to previously unseen anomaly types. Unsupervised methods  
036 model the normal manifold and detect deviations, yet the absence of anomalous supervision ham-  
037 pers calibration and discriminability—especially for subtle anomalies or under distributional drift.  
038 This scarcity–diversity dilemma motivates exploring transfer learning when the dataset under study  
039 lacks anomaly labels, leveraging related datasets that provide labelled anomalies.

040 Partial Domain Adaptation (PDA) is a natural option in this setting: it aligns the source–target  
041 distributions of normal features, after which source anomaly labels supervise learning in the shared  
042 representation. From an optimisation perspective, mini-batch updates decompose into normal and  
043 anomalous gradient components. Supervised training on the source induces anomalous-gradient  
044 directions tailored to the source distribution; under domain shift, these directions need not benefit  
045 the target to the same extent. Moreover, in anomaly-detection deployments where the target provides  
046 only normal data (one-class condition), the anomalous component is missing, yielding an incomplete  
047 update signal. This motivates estimating the missing component in the target gradient space via  
048 a learned transport map from the source. Fig. 1 examines feasibility: across epochs, per-sample

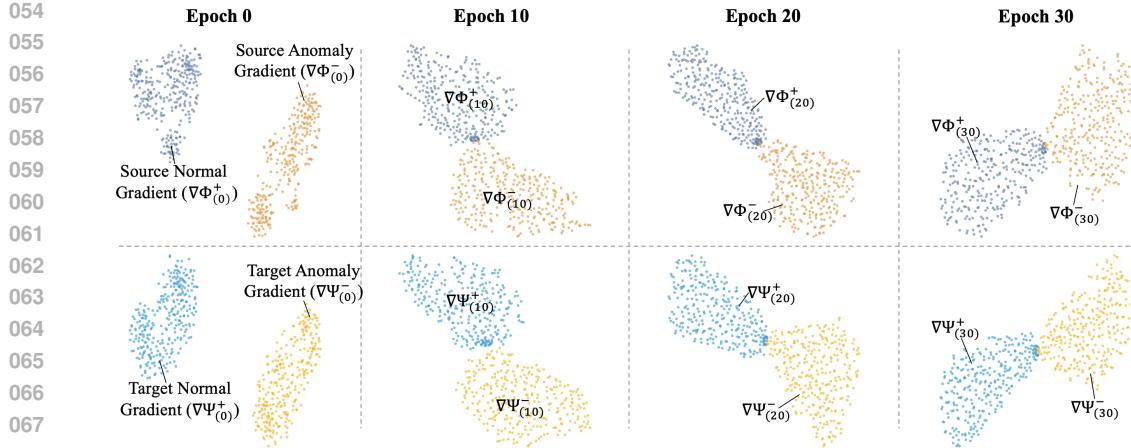


Figure 1: **Cross-Domain Gradient Distribution Consistency.** Each panel visualises, via t-SNE, the per-sample *gradient vectors* obtained at four training epochs (0, 10, 20, 30; left → right). **Top row:** source domain (Fan) showing normal ( $\nabla\Phi^+$ , blue) and anomalous ( $\nabla\Phi^-$ , orange) gradients. **Bottom row:** target domain (Pump) presenting normal gradients ( $\nabla\Psi^+$ , sky-blue) together with anomalous gradients estimated by the proposed mapping ( $\nabla\hat{\Psi}^- = \mathcal{F}(\nabla\Phi^-)$ , yellow). Across all epochs, the spatial arrangement of normal and anomalous manifolds in the two domains remains highly congruent, empirically supporting the assumption  $\mathcal{P}(\nabla\Phi^+, \nabla\Phi^-) \approx \mathcal{P}(\nabla\Psi^+, \nabla\Psi^-)$  and thereby motivating the cross-domain gradient-reuse strategy.

gradients from two proximal domains form normal–anomalous manifolds with highly congruent (near-isometric) geometry, indicating that such a transport is learnable.

Building on this observation, Domain-Aware Gradient Reuse (DAGR) is introduced. DAGR first learns an adaptive transport map by aligning source and target normal gradients, thereby attenuating domain-specific components. The same map is then reused to project source anomalous gradients into the target space, producing estimated target anomalous gradients. In controlled evaluations where target anomalies are available for assessment, Fig. 2 shows that the mapped gradients closely overlap with the empirical target anomalous-gradient distribution across training epochs. By augmenting the target updates with this estimated anomalous component, DAGR guides the convergence of the target-domain detector without labelled anomalies.

This work proposes DAGR, a transfer-learning framework that remains effective even when the target domain contains no anomalous samples. Extensive experiments on image and audio benchmarks demonstrate state-of-the-art performance, while ablation studies isolate the contribution of each module. The appendix provides a complete convergence proof under stated assumptions, thereby giving the method a rigorous theoretical foundation.

## 2 RELATED WORK

Prior work on anomaly detection is grouped into three strands: augmentation-based supervised methods, unsupervised one-class modeling, and transfer learning.

### 2.1 SUPERVISED METHODS

Supervised anomaly detection commonly mitigates class imbalance via augmentation. Input-level transformations—geometric/photometric operations for images, time-warping and jittering for sequences, and graph-topology perturbations—expand the training distribution while preserving labels (Mumuni et al., 2024; Dang et al., 2023; Zhu et al., 2021). Generative augmentation employs GAN- or diffusion-based models to synthesise harder anomalous instances, enriching the minority class (Goodfellow et al., 2020; Cao et al., 2024). At the representation level, contrastive augmentation applies random masking or multi-view transformations with contrastive objectives to improve in-

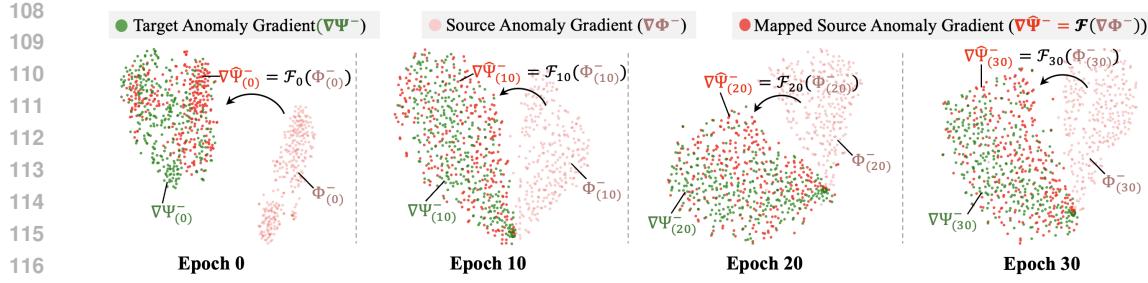


Figure 2: **Epoch-wise estimation of target anomalous gradients via cross-domain mapping.** Each panel shows a t-SNE embedding of *anomalous-sample* gradient vectors obtained at four training epochs (0, 10, 20, 30; panels arranged left → right). Green points represent the *true* target anomalous gradients  $\nabla \Psi^-$ ; light-pink points depict the *source* anomalous gradients  $\nabla \Phi^-$ ; red points are their epoch-specific images in the target space,  $\nabla \hat{\Psi}^- = \mathcal{F}_{(e)}(\nabla \Phi^-)$ . At each epoch, the mapped source-anomaly gradients nearly coincide with the true target-anomaly gradients, indicating that  $\mathcal{F}_{(e)}$  provides an accurate *per-epoch* estimate of the target anomalous-gradient distribution.

variance and discrimination (Han et al., 2023; Zhou et al., 2022). Together, these strategies increase sample diversity without modifying ground-truth labels.

**Limitation.** However, anomalous samples in training rarely reflect the true anomaly space. Even advanced augmentations generate limited variants and cannot bridge the semantic gap between observed and unseen anomalies, making generalisation beyond the augmentation manifold difficult.

## 2.2 UNSUPERVISED METHODS

Unsupervised anomaly detection learns normality from unlabelled data, using signals ranging from reconstruction fidelity to predictive objectives, representation discrimination, and density modelling. Reconstruction-based methods—Auto-Encoders, VAEs, GAN restorers, and diffusion decoders—identify anomalies by large residuals (Chen et al., 2018; An & Cho, 2015; Hussein et al., 2020; Wu et al., 2024). Self-supervised tasks such as future prediction, masked-signal recovery, and transformation discrimination extract intrinsic dynamics without labels (Venkatraman et al., 2015; Xie et al., 2023; Swarna et al., 2022). Contrastive learning compacts the normal manifold by attracting genuine instances and repelling perturbed views (Liang et al., 2022). Probabilistic density estimators—normalising flows and energy-based models—learn likelihoods so that low-density samples can be flagged (Garcia Satorras et al., 2021; Qin et al., 2022). These directions jointly approximate the normal manifold via reconstruction error, embedding compactness, and likelihood.

**Limitation.** Because no anomalous instances participate in training, the learned boundary is inferred solely from normal data, often yielding overly broad decision regions and reduced precision on subtle or high-variance anomalies.

## 2.3 TRANSFER LEARNING FOR ANOMALY DETECTION

Transfer learning is a viable strategy when the target dataset lacks anomalous samples, because a source domain enriched with labelled outliers can furnish the discriminative information that the target model requires. Transfer-based anomaly-detection research can be grouped into three lines of work. Partial Domain Adaptation (PDA) is the most widely adopted paradigm because it matches the practical setting where the target domain contains only normal data. PDA studies align cross-domain normal representations while suppressing source-only anomalies through class-importance weighting (Zhang et al., 2018), instance-level selection (Nguyen et al., 2023), or classifier-consistency (Jeong & Shin, 2020) strategies. Domain Adaptation (DA) assumes identical label spaces and exploits unlabeled target data to mitigate domain shift, typically using adversarial feature alignment (Chen et al., 2019), (Du et al., 2021), or self-supervised reconstruction constraints (Zhou et al., 2024). Domain Generalization (DG) trains without target data and pursues

162 domain-invariant features by applying meta-learning (Khoe et al., 2024) across multiple source do-  
 163 mains, style or feature perturbations (Liu et al., 2024), or gradient-based regularization (Tang et al.,  
 164 2021), thereby improving robustness to unseen environments.  
 165

166 **Limitation.** Although the alignment of normal representations mitigates interdomain shift, the op-  
 167 timisation signals derived from source-domain anomaly supervision remain domain-specific. Gra-  
 168 dient directions that accelerate convergence in the source model do not necessarily align with the target  
 169 optimisation landscape. Consequently, the transferred supervision provides limited discriminative  
 170 guidance for the target detector and constrains further gains in anomaly-detection performance.  
 171

172 In summary, data augmentation improves class balance but does not address the unknown anomaly  
 173 distribution. Unsupervised approaches avoid the need for anomaly labels, yet their precision re-  
 174 mains limited because they lack anomalous guidance. Transfer-based methods also exhibit a critical  
 175 shortcoming: the update directions induced by source anomalous supervision are optimised for the  
 176 source distribution and are not guaranteed to benefit the target.  
 177

### 3 METHODS

179 This section details the proposed gradient-reuse framework. It first outlines the task setting and the  
 180 motivation for exploiting source-domain information under severe anomaly scarcity. The subsequent  
 181 subsection, ‘‘Cross-Domain Consistent Component Selection (CCCS),’’ explains how components  
 182 that exhibit domain-invariant behavior are identified and preserved. Finally, ‘‘Adaptive Domain-  
 183 Specific Perturbation Removal (ADPR)’’ describes how these components are leveraged to learn  
 184 cross-domain gradients, enabling the estimation of target anomaly gradient. A rigorous convergence  
 185 proof of DAGR is provided in the Appendix, establishing the theoretical soundness of the method.  
 186

#### 3.1 MOTIVATION

188 Let the *source domain*  $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$  contain both *normal* ( $y = 0$ ) and *anomalous* ( $y = 1$ )  
 189 instances, while the *target domain*  $\mathcal{D}_t = \{\mathbf{x}_j^t\}_{j=1}^{N_t}$  is assumed to comprise normal data only. A  
 190  $K$ -layer deep network is considered, whose layerwise parameters are collected as  
 191

$$\Phi = \{\phi_1, \phi_2, \dots, \phi_K\} \quad \text{and} \quad \Psi = \{\psi_1, \psi_2, \dots, \psi_K\} \quad (1)$$

193 for the source and target models, respectively. At each training step we compute stochastic gradients  
 194  $\nabla\Phi$  on the source mini-batch and  $\nabla\Psi$  on the target mini-batch. Breaking the source gradient into  
 195 class-conditioned components gives

$$\nabla\Phi = \nabla\Phi^+ + \nabla\Phi^-, \quad (2)$$

198 where  $\nabla\Phi^+$  and  $\nabla\Phi^-$  are from normal and anomalous samples, respectively. Because  $\mathcal{D}_t$  lacks  
 199 anomalies, only

$$\nabla\Psi^+ = \frac{1}{|\mathcal{B}_t|} \sum_{\mathbf{x}_j^t \in \mathcal{B}_t} \nabla_\Psi \mathcal{L}(f_\Psi(\mathbf{x}_j^t), 0) \quad (3)$$

202 is observable in the target domain, with  $\mathcal{B}_t$  denoting the current target mini-batch.  
 203

204 **Empirical observation.** Figure 1 plots the distributions of  $\nabla\Phi^+$ ,  $\nabla\Phi^-$ ,  $\nabla\Psi^+$  and  $\nabla\Psi^-$  over  
 205 training epochs. The divergence  $\mathcal{D}(\nabla\Phi^+, \nabla\Phi^-) \approx \mathcal{D}(\nabla\Psi^+, \nabla\Psi^-)$  remains low, where  $\mathcal{D}$  is  
 206 instantiated as the 1-Wasserstein distance. This alignment suggests that both domains share a  
 207 *domain-invariant* gradient component despite being collected in different environments.  
 208

209 **Gradient decomposition hypothesis.** We therefore posit that every mini-batch gradient can be  
 210 decomposed into

$$\nabla\Phi = \Omega + \nabla\Phi_a, \quad \nabla\Psi = \Omega + \nabla\Psi_a, \quad (4)$$

212 where  $\Omega$  encodes *cross-domain* knowledge that is useful for both domains;  $\nabla\Phi_a$  and  $\nabla\Psi_a$  capture  
 213 *domain-specific perturbations*.

214 Because both domains provide abundant normal data, we learn a mapping  
 215

$$\mathcal{F} : \nabla\Phi^+ \longmapsto \Omega^+ \quad (5)$$

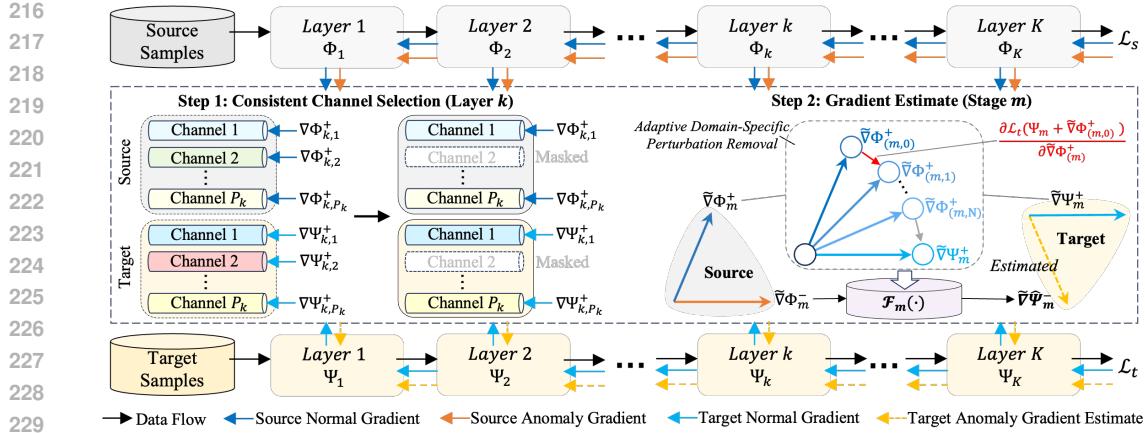


Figure 3: **Overall workflow of the proposed cross-domain gradient reuse framework.** The pipeline is executed on a *source network* ( $\Phi_{1:K}$ , top) and an architecturally identical *target network* ( $\Psi_{1:K}$ , bottom). **Step 1: Consistent Channel Selection (Layer  $k$ )**. For every layer  $k$ , the cosine similarity between the *normal-sample* gradients  $\nabla\Phi_{k,p}^+$  and  $\nabla\Psi_{k,p}^+$  of each channel  $c_{k,p}$  is evaluated; channels whose similarity falls below the threshold  $\gamma$  are masked (grey), leaving only domain-invariant components (solid colour). **Step 2: Adaptive Domain-Specific Perturbation Removal**. Given the masked source normal gradient  $\tilde{\nabla}\Phi_m^+$  at outer stage  $m$ , an inner optimisation loop (blue trajectory) adapts it to the target loss, yielding the de-domainised estimate  $\Omega_m^+ = \mathcal{F}_m(\tilde{\nabla}\Phi_m^+)$ . The same map  $\mathcal{F}_m(\cdot)$  is then reused to transform the masked *anomalous* gradient, producing  $\Omega_m^-$ . Both components are aggregated as  $\Omega_m = \Omega_m^+ + \Omega_m^-$  and injected into the target network update (yellow dashed arrows), enabling anomaly knowledge transfer without exposing target data to anomalies.

that removes domain-specific noise from normal gradients. Under the *label-independent shift* assumption, Ben-David et al. (Ben-David et al., 2010) bound the target risk by the source risk plus the distribution divergence between domains. Coupled with the *Gradient Distribution Alignment* principle, this implies that the same  $\mathcal{F}$  generalises to anomalous gradients:

$$\mathcal{F}(\nabla\Phi^-) \approx \Omega^- \quad (6)$$

Aggregating  $\Omega = \Omega^+ \cup \Omega^-$  yields a low-variance estimate of the domain-invariant descent direction.

**Transferring anomalous knowledge.** Finally, the target model is updated by

$$\psi_k \leftarrow \psi_k - \eta \cdot \Omega_k, \quad k = 1, \dots, K, \quad (7)$$

where  $\eta$  is the learning rate. Equation 7 enables *implicit reuse* of anomalous gradients without exposing the target model to any anomalous data. We prove in Appendix that, under mild smoothness conditions (Gao et al., 2021a), the update rule in equation 7 reduces the target risk upper bound monotonically, thereby accelerating convergence.

The key insight is that *gradients—rather than features or logits—exhibit strong cross-domain regularities*. By denoising source gradients (equation 5), estimating the shared component  $\Omega$  (equation 4), and injecting it into target updates (equation 7), our method transfers *anomalous knowledge* to a domain where no anomalies are observable. Formal justification is provided in Appendix A.1.

### 3.2 CROSS-DOMAIN CONSISTENT COMPONENT SELECTION

Although the source and target networks share an identical architecture, individual sub-modules (e.g., convolutional channels or Transformer heads) may specialise in *domain-specific* patterns. If such components participate in learning the de-domainisation map  $\mathcal{F}$ , the resulting estimate of the shared descent direction  $\Omega$  would be biased. Hence, before training  $\mathcal{F}$ , we automatically identify and retain only those components whose behaviour is consistent across domains.

270 **Channel-wise gradient similarity.** Consider a  $K$ -layer CNN. Let  $k \in \{1, \dots, K\}$  index the layers and  $\mathcal{C}_k = \{c_k^1, \dots, c_k^{P_k}\}$  denote the  $P_k$  output channels of layer  $k$ . For each channel  $c_k^p \in \mathcal{C}_k$  we measure the *normal-sample* gradients in the source and target domains:

$$274 \quad \nabla\Phi_{k,p}^+ = \frac{1}{|\mathcal{B}_s^+|} \sum_{\mathbf{x}_i^s \in \mathcal{B}_s^+} \nabla_{\phi_{k,p}} \mathcal{L}(f_\Phi(\mathbf{x}_i^s), 0), \quad (8)$$

$$277 \quad \nabla\Psi_{k,p}^+ = \frac{1}{|\mathcal{B}_t|} \sum_{\mathbf{x}_j^t \in \mathcal{B}_t} \nabla_{\psi_{k,p}} \mathcal{L}(f_\Psi(\mathbf{x}_j^t), 0), \quad (9)$$

280 where  $\phi_{k,p}$  (resp.  $\psi_{k,p}$ ) collects the weights associated with channel  $c_k^p$  in the source (resp. target) 281 model. Their cosine similarity

$$282 \quad \rho_{k,p} = \frac{\langle \nabla\Phi_{k,p}^+, \nabla\Psi_{k,p}^+ \rangle}{\|\nabla\Phi_{k,p}^+\|_2 \|\nabla\Psi_{k,p}^+\|_2} \in [-1, 1] \quad (10)$$

285 reflects the extent to which channel  $c_k^p$  reacts similarly to *normal* data from both domains.

287 **Domain-invariant channel mask.** Given a global *channel-masking ratio*  $\alpha \in (0, 1)$ , the 288 layer-wise threshold  $\gamma_k$  is chosen as the  $\alpha$ -percentile of the cosine similarities  $\{\rho_{k,p}\}_{p=1}^{C_k}$  in layer  $k$ :

$$290 \quad \gamma_k = \text{Percentile}_\alpha(\{\rho_{k,p}\}_{p=1}^{C_k}). \quad (11)$$

292 **Binary mask.** Using the data-driven threshold equation 11, a binary mask is defined as

$$293 \quad m_{k,p} = \mathbb{I}[\rho_{k,p} \geq \gamma_k], \quad c_k^p \in \mathcal{C}_k, \quad (12)$$

295 where  $\mathbb{I}[\cdot]$  is the indicator function. Channels with  $\rho_{k,p} < \gamma_k$  are treated as domain-specific and *de-*  
296 *activated* by nullifying their gradients:  $\tilde{\nabla}\Phi_{k,p}^+ = m_{k,p} \nabla\Phi_{k,p}^+$ ,  $\tilde{\nabla}\Psi_{k,p}^+ = m_{k,p} \nabla\Psi_{k,p}^+$ ,  $\tilde{\nabla}\Phi_{k,p}^- =$   
297  $m_{k,p} \nabla\Phi_{k,p}^-$ . Aggregating over all layers yields the final masked gradients  $\tilde{\nabla}\Phi^\pm$  and  $\tilde{\nabla}\Psi^\pm$ .

298 The filtered gradients are fed into Eq. equation 5:

$$300 \quad \Omega^+ = \mathcal{F}(\tilde{\nabla}\Phi^+), \quad \Omega^- \approx \mathcal{F}(\tilde{\nabla}\Phi^-). \quad (13)$$

301 By explicitly excising domain-specific channels, the variance of the shared estimate  $\Omega$  is further  
302 reduced, which empirically accelerates convergence and stabilises the target update rule equation 7.

### 304 3.3 ADAPTIVE DOMAIN-SPECIFIC PERTURBATION REMOVAL

306 The masked normal gradients  $\tilde{\nabla}\Phi^+ \in \mathbb{R}^D$  and  $\tilde{\nabla}\Psi^+ \in \mathbb{R}^D$  are extremely high-dimensional ( $D \approx$   
307  $10^{6-9}$ ) and exhibit complex, *non-linear* cross-domain discrepancies. Simple statistics (e.g., means  
308 or linear projections) are therefore insufficient for extracting the shared component  $\Omega^+$ . We instead  
309 implement the adaptive cross-Domain gradient distiller  $\mathcal{F}(\cdot)$  of Eq. equation 5 as a *gradient-based*,  
310 *end-to-end adaptive procedure* that removes domain-specific perturbations from source gradients by  
311 directly minimising the *target* loss.

312 **Outer-inner optimisation view.** At global training step  $m \in \mathbb{N}$  let  $\Phi_m$  and  $\Psi_m$  denote the current  
313 source and target network parameters, and define the (normal-sample) source gradient  $g_m^{(0)} =$   
314  $\tilde{\nabla}\Phi_m^+$ . We treat  $g_m^{(0)}$  as the *optimisable variable* and run an inner loop of  $N$  steps to obtain its  
315 domain-invariant component  $g_m^{(N)}$ .

317 **Target-aligned inner loop.** Starting from the *fast weight*

$$319 \quad \Psi_{m,0} = \Psi_m + \alpha g_m^{(0)}, \quad (14)$$

321 where  $\alpha > 0$  is a small, fixed step size, we freeze the backbone parameters  $\Psi$  and iteratively refine  
322  $g_m^{(n)}$  by descending the target loss  $\mathcal{L}_t \triangleq \mathcal{L}(\mathcal{D}_t; \cdot)$ :

$$323 \quad g_m^{(n+1)} = g_m^{(n)} - \beta \nabla_g \mathcal{L}_t(\Psi_m + \alpha g_m^{(n)}), \quad (15)$$

324 where  $n = 0, 1, \dots, N - 1$  and  $\beta > 0$  denotes the inner-loop learning rate. The inner gradient  
 325  $\nabla_g \mathcal{L}_t$  is computed over the current target mini-batch  $\mathcal{B}_t$  and back-propagated *through* the fast weight  
 326 construction in Eq. equation 14, thereby allowing  $g_m^{(n)}$  to adapt to target-domain feedback.  
 327

328 After  $N$  refinement steps we define

$$329 \quad \Omega_m^+ = \mathcal{F}(\tilde{\nabla}\Phi_m^+) \triangleq g_m^{(N)}, \quad (16)$$

331 and inject  $\Omega_m^+$  into the target update rule equation 7. Because the inner optimisation equation 15 is  
 332 conditioned *solely* on target normal data,  $\Omega_m^+$  is empirically free of domain-specific artefacts.  
 333

334 **Complexity.** A detailed time/space complexity discussion and implementation notes are provided  
 335 in Appendix.  
 336

337 **Algorithmic summary.** The overall training alternates between (i) sampling a source normal  
 338 mini-batch to obtain  $\tilde{\nabla}\Phi_m^+$ , (ii) executing the inner loop equation 14–equation 15 to produce  $\Omega_m^+$ ,  
 339 and (iii) updating the target parameters via Eq. equation 7.  
 340

341 **Gradient reuse.** Once the de-domainisation map  $\mathcal{F}$  has been obtained via the inner loop in  
 342 Eqs. equation 14–equation 16, it is *reused* to process the *anomalous* source gradients:  
 343

$$344 \quad \Omega_m^- = \mathcal{F}(\tilde{\nabla}\Phi_m^-). \quad (17)$$

346 We then aggregate the normal and anomalous components,  
 347

$$348 \quad \Omega_m = \Omega_m^+ + \Omega_m^-, \quad (18)$$

349 and apply the shared descent direction to the target network using the update rule of Eq. equation 7.  
 350 In this way, *anomalous knowledge* is injected into the target model purely through gradient transfer,  
 351 with no anomalous samples ever appearing in the target domain.  
 352

## 354 4 EXPERIMENTS

### 356 4.1 EXPERIMENTS SETUP

358 **Datasets and source–target protocol.** DAGR was evaluated on DCASE 2020 Task 2 (Koizumi  
 359 et al., 2020) and DAGM (Wieler et al., 2007). Source–target pairs were formed between *adjacent domains*  
 360 sharing sensing modality and generative mechanism—DCASE features motor-driven  
 361 machines recorded under a common acoustic pipeline, while DAGM comprises homogeneous man-  
 362 ufactured textures—ensuring a meaningful transfer basis. For DCASE, *Fan* is fixed as the source  
 363 for stable in-domain performance and gradients; *Pump*, *ToyCar*, and *ToyConv.* serve as typical  
 364 motor-driven targets with prominent motor signatures, whereas *Valve*, dominated by electromagnetic  
 365 actuation and airflow pulses, and *Slider*, driven by reciprocating stage motion with weak motor har-  
 366 monics, are included as *heterogeneity stress tests* to probe robustness under stronger cross-domain  
 367 differences. For DAGM, *Class 2* is chosen as the source owing to its strongest in-domain per-  
 368 formance; *Class 1*, *Class 3*, and *Class 6* are selected as more challenging targets, whereas *Class 4/5*  
 369 are omitted because unsupervised baselines already saturate. In all settings, all labelled source data  
 370 are available, each target exposes only 10% of its normal samples with no anomalies, and a single  
 371 source is applied per dataset without per-target tuning.

372 **Baseline.** To ensure a thorough comparison with current state-of-the-art approaches, four baseline  
 373 categories were evaluated. The unsupervised group comprised General-AD (Sträter et al., 2024) and  
 374 GLASS (Chen et al., 2024). The partial domain adaptation group consisted of PDA (Bai et al., 2024),  
 375 CMKD (Zhou & Zhou, 2024), UniNet (Wei et al., 2025b), ANC (Zhang et al., 2024), JWO (Chen,  
 376 2024), PWAN (Wang et al., 2025) and MLWE (Wen et al., 2024). The domain generalisation group  
 377 included BDC (Zhang et al., 2025b), DDDG (Zhang et al., 2025a), GGA (Ballas & Diou, 2025) and  
 378 DKGPL (Wei et al., 2025a).

378  
 379 Table 1: AUROC (%) comparison on 8 benchmark datasets. Best per column in **bold**, second best  
 380 is underlined.

381 Methods	382 Source Domain	DCASE (Fan)					DAGM (Class 2)			Ave.
		383 Target Domain	384 Pump	Slider	385 Valve	386 ToyCar	387 ToyConv.	388 Class 1	389 Class 3	390 Class 6
383 Unsupervised 384 Methods	385 General-AD	69.80	82.20	66.09	58.39	58.80	59.95	70.89	90.11	69.03
	386 GLASS	65.93	88.37	<b>67.42</b>	63.22	59.76	<u>90.98</u>	80.22	70.04	<u>73.45</u>
385 Partial 386 Domain 387 Adaptation	388 JWO	41.32	64.99	64.40	65.01	54.73	55.71	64.53	81.45	61.52
	389 PWAN	60.97	53.83	52.66	58.10	59.67	50.21	59.94	56.69	56.51
	390 MLWE	42.75	63.30	49.59	54.46	63.35	49.31	47.63	50.62	52.63
	391 CMKD	53.47	66.68	56.69	48.49	54.93	54.93	78.55	<u>94.51</u>	63.53
388 Domain 389 Adaptation	390 UniNet	45.62	49.32	52.72	44.92	51.72	48.64	56.08	59.08	51.01
	391 ANC	62.08	45.19	51.55	<u>68.10</u>	<u>76.86</u>	56.89	57.63	51.96	58.78
	392 FFTAT	57.72	71.06	55.92	64.42	57.15	58.38	<u>80.59</u>	86.04	66.41
391 Domain 392 Generalization	393 GGA	69.26	76.77	62.48	67.08	70.35	58.20	52.89	74.88	66.49
	394 BDC	60.22	52.75	53.00	56.38	54.78	51.53	56.81	62.30	55.97
	395 DDDG	58.73	59.95	52.74	53.12	52.28	57.65	56.73	63.53	56.84
	396 PMGDG	68.27	45.55	54.87	56.45	60.21	51.53	44.72	71.44	56.63
	397 DKGPL	50.62	63.87	52.57	64.62	60.79	50.15	60.63	58.72	57.75
398 Proposed	399 DAGR	<b>83.42</b>	<b>88.96</b>	55.62	<b>72.08</b>	<b>79.14</b>	<b>93.31</b>	<b>81.69</b>	<b>95.27</b>	<b>81.19</b>

396  
 397 **Evaluation Metrics.** Three metrics are considered to assess the performance of MDPE: AUROC,  
 398 AUPRC, and Rec@K. **AUROC** (area under the ROC) quantifies the ability to distinguish positive  
 399 from negative classes and is widely regarded as a stable, threshold-agnostic indicator of discrimination  
 400 performance. **AUPRC** (area under the PRC) summarizes the trade-off between precision  
 401 and recall across thresholds and is particularly informative under severe class imbalance. **Rec@K**  
 402 (recall at rank  $K$ ) measures the proportion of true anomalies retrieved among the top- $K$  ranked  
 403 instances, where  $K$  equals the number of anomalous samples in the test set. Owing to space con-  
 404 straints, AUROC is adopted as the primary metric and its results are reported in the main text.

405 The *source codes* and more results about *AUPRC, Rec@K* are given in the *supplementary materials*.

406  
 407 **Comparison with SOTA.** Table 1 reports AUROC on all targets. DAGR achieves the best score  
 408 on seven of eight domains—every DCASE target except *Valve*—and ranks first on all three MVTec  
 409 defects; *Valve* is non-stationary and cross-mechanism, hence outside our adjacent-domain scope.  
 410 Averaged across benchmarks, it attains 81.19% AUROC, exceeding GLASS (73.45%) by +7.7 pp  
 411 and GGA (66.49%) by +14.7 pp. The gains are consistent across acoustic targets (*Pump, Slider,*  
 412 *ToyCar, ToyConveyor*) and visual targets (*Cable, Capsule, Hazelnut*), supporting the effectiveness  
 413 of the proposed gradient-reuse strategy.

## 414 4.2 ABLATION STUDY

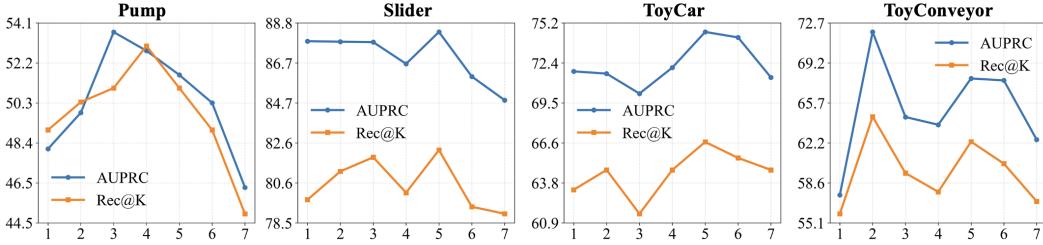
415 The ablation study investigates the contribution of each core component and the influence of  
 416 channel-masking ratio ( $\alpha$ ).

417  
 418 **Effectiveness of CCCS and ADPR** Table 2 compares the full DAGR model with four ablated  
 419 variants. Dropping the Cross-Domain Consistent Component Selection (w/o CCCS) reduces mean  
 420 AUROC by about 3 percentage points and mean AUPRC by about 6 points, showing that filtering  
 421 out gradient-inconsistent channels offers a clear yet secondary gain. In contrast, eliminating the  
 422 Adaptive Domain-Specific Perturbation Removal (w/o ADPR) causes a sharp decline of roughly  
 423 16 points in AUROC and 33 points in AUPRC, indicating that learning a cross-domain gradient  
 424 transformation is essential for successful reuse of source information. Simply substituting ADPR  
 425 with conventional feature alignment (w FA) or a linear gradient mapping (w LT) only partially  
 426 restores performance; both alternatives still trail the complete model by more than 15 points in  
 427 AUROC and more than 30 points in AUPRC on average. **Beyond accuracy, these variants exhibit**  
 428 **stable training dynamics: removing CCCS/ADPR reduces accuracy without inducing divergence,**  
 429 **indicating that both modules act as variance-reducing regularisers for gradient-space updates.** These  
 430 results confirm that CCCS helps but ADPR is the primary driver of DAGR’s effectiveness, and that  
 431 sophisticated gradient-space adaptation is required to fully exploit source-domain knowledge.

432

433  
434  
Table 2: The results (AUROC, %) of the ablation study on different modules. The highest score is  
highlighted in **bold**.

435 436 Model	437 Pump		438 Slider		439 Valve		440 ToyCar		441 ToyConveyor	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
w/o CCCS	82.65	51.86	87.20	86.79	48.95	15.93	70.46	69.72	74.14	57.58
w/o ADPR	<b>57.79</b>	28.14	69.86	39.92	54.91	16.08	55.03	25.52	62.07	36.23
w FA	61.53	25.61	70.86	42.23	42.39	22.98	66.52	34.66	59.45	33.80
w LT	59.97	20.98	55.93	31.05	42.07	24.08	60.65	29.19	64.35	34.89
DAGR	<b>83.42</b>	<b>53.71</b>	<b>88.96</b>	<b>87.79</b>	<b>55.62</b>	<b>26.52</b>	<b>72.08</b>	<b>71.76</b>	<b>79.14</b>	<b>71.91</b>

452  
453  
Figure 4: **Effect of channel-masking ratio ( $\alpha$ )**. Detection performance (AUPRC and Rec@K)  
454  
versus channel-masking ratio (%) on four representative benchmarks.455  
456  
**Effect of the channel-masking ratio ( $\alpha$ )**. Figure 4 reports AUPRC (green) and Recall@K (red)  
457  
as the masking ratio  $\alpha$  is varied from 1% to 7% on four MIMII tasks. All curves rise when a  
458  
small fraction of gradient-divergent channels is suppressed, peak in the 3% – 5% range, and decline  
459  
thereafter. Peak values appear at 59.8% / 57.0% on Fan, 87.8% / 81.9% on Slider, and 74.6% /  
460  
66.7% on ToyCar. ToyConveyor reaches its first maximum at 2% and a secondary, gentler high near  
461  
5%. These results indicate that masking roughly 5% of channels achieves the best trade-off between  
462  
noise removal and information retention; therefore  $\alpha = 5\%$  is used in the remaining experiments.463  
4.3 DISCUSSION AND FUTURE WORK464  
DAGR delivers the strongest overall performance, reaching an average AUROC of 81.19% and ranking  
465  
first on seven of eight targets, including all three DAGM classes, with consistent gains across  
466  
acoustic and visual domains. The underperformance on *Valve* delineates the method’s boundary  
467  
rather than contradicting it. The *Fan* source exhibits quasi-stationary harmonic spectra from ro-  
468  
tating parts, whereas *Valve* is dominated by non-stationary flow transients and a different physical  
469  
process; the gradient-consistency premise is therefore not satisfied and transferability is limited.  
470  
This clarifies the intended scope of DAGR: adjacent domains that share sensing modality and gen-  
471  
erative mechanism, such as motor-driven machinery and manufactured textures. For deployment,  
472  
a proximity screen on normal-gradient geometry should be used ; when proximity falls below a  
473  
threshold, the mapped anomalous component should be down-weighted or disabled, reverting to  
474  
a conservative target-only update to avoid negative transfer. Accordingly, future work will enable  
475  
proximity-aware gating by default and assess applicability under broader cross-domain shifts.476  
5 CONCLUSION477  
This paper presents Domain-Aware Gradient Reuse (DAGR), a transfer-learning framework that  
478  
reinterprets domain adaptation as the selective reuse of source-domain gradients. By integrating  
479  
gradient-consistency filtering with adaptive perturbation removal, DAGR provides both a formal  
480  
convergence guarantee and a practical pathway to cross-domain generalisation. Extensive experi-  
481  
ments on eight acoustic and visual anomaly detection benchmarks achieve state-of-the-art per-  
482  
formance, showing that gradients, rather than features, can serve as an effective conduit for knowl-  
483  
edge transfer. These results introduce a gradient-centric perspective to anomaly detection and open  
484  
promising avenues for future adaptation strategies grounded in gradient compatibility.

486 ETHICS STATEMENT  
487488 This work adheres to the ICLR Code of Ethics. Experiments rely on publicly available or appropriately  
489 licensed datasets; where data may contain personal or sensitive attributes, de-identification and  
490 license terms are respected, and no attempt is made to re-identify individuals. The study does not  
491 target or enable discriminatory or unsafe use; foreseeable dual-use risks are discussed and mitigation  
492 strategies (e.g., responsible release, robustness and bias checks) are described in the supplementary  
493 materials. No human-subject intervention, clinical decision-making, or deployment in safety-critical  
494 settings was conducted; any future deployment will follow applicable legal and institutional review  
495 requirements. Funding sources and potential conflicts of interest are disclosed. All procedures, data  
496 handling, and reporting were conducted with attention to privacy, fairness, and research integrity.  
497498 REPRODUCIBILITY STATEMENT  
499500 Reproducibility has been prioritized. The main text specifies datasets and splits, preprocessing  
501 pipelines, model architectures, training schedules, and evaluation protocols, while exact hyperpa-  
502 rameters, random seeds, and ablation configurations are provided in the appendix. An anonymized  
503 repository with source code, configuration files, and scripts to regenerate all tables and figures is  
504 supplied in the supplementary materials, including environment setup instructions and hardware re-  
505 quirements. For theoretical results, assumptions are stated in the paper and complete proofs are  
506 included in the appendix. External baselines are referenced with versions or commit hashes, and  
507 dataset licenses and checksums are reported to ensure fidelity. Together, these materials enable  
508 end-to-end reproduction of the reported results.  
509510 REFERENCES  
511512 Jinwon An and Sungsoon Cho. Variational autoencoder based anomaly detection using reconstruc-  
513 tion probability. *Special lecture on IE*, 2(1):1–18, 2015.  
514 Shuanghao Bai, Min Zhang, Wanqi Zhou, Siteng Huang, Zhirong Luan, Donglin Wang, and Badong  
515 Chen. Prompt-based distribution alignment for unsupervised domain adaptation. In *Proceedings  
516 of the AAAI conference on artificial intelligence*, volume 38, pp. 729–737, 2024.  
517 Aristotelis Ballas and Christos Diou. Gradient-guided annealing for domain generalization. In *Pro-  
518 ceedings of the Computer Vision and Pattern Recognition Conference*, pp. 20558–20568, 2025.  
519  
520 Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations  
521 for domain adaptation. *Advances in neural information processing systems*, 19, 2006.  
522  
523 Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wort-  
524 man Vaughan. A theory of learning from different domains. *Mach. Learn.*, 79(1–2):151–175,  
525 May 2010. ISSN 0885-6125. doi: 10.1007/s10994-009-5152-4. URL <https://doi.org/10.1007/s10994-009-5152-4>.  
526  
527 Dimitri P Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3):  
528 334–334, 1997.  
529  
530 Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and  
531 Stan Z Li. A survey on generative diffusion models. *IEEE transactions on knowledge and data  
532 engineering*, 36(7):2814–2830, 2024.  
533  
534 Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu,  
535 and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *Pro-  
536 ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 627–636,  
537 2019.  
538  
539 Qiyu Chen, Huiyuan Luo, Chengkan Lv, and Zhengtao Zhang. A unified anomaly synthesis strategy  
with gradient ascent for industrial anomaly detection and localization. In *European Conference  
on Computer Vision*, pp. 37–54. Springer, 2024.

540 Sentao Chen. Joint weight optimization for partial domain adaptation via kernel statistical distance  
 541 estimation. *Neural Networks*, 180:106739, 2024.

542

543 Zhaomin Chen, Chai Kiat Yeo, Bu Sung Lee, and Chiew Tong Lau. Autoencoder-based network  
 544 anomaly detection. In *2018 Wireless telecommunications symposium (WTS)*, pp. 1–5. IEEE, 2018.

545 Yizhou Dang, Enneng Yang, Guibing Guo, Linying Jiang, Xingwei Wang, Xiaoxiao Xu, Qinghui  
 546 Sun, and Hong Liu. Uniform sequence better: Time interval aware data augmentation for sequen-  
 547 tial recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37,  
 548 pp. 4225–4232, 2023.

549

550 Zhekai Du, Jingjing Li, Hongzu Su, Lei Zhu, and Ke Lu. Cross-domain gradient discrepancy min-  
 551 imization for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on*  
 552 *computer vision and pattern recognition*, pp. 3937–3946, 2021.

553 Zhiqiang Gao, Shufei Zhang, Kaizhu Huang, Qiufeng Wang, and Chaoliang Zhong. Gradient dis-  
 554 tribution alignment certificates better adversarial domain adaptation. In *2021 IEEE/CVF Interna-  
 555 tional Conference on Computer Vision (ICCV)*, pp. 8917–8926, 2021a. doi: 10.1109/ICCV48922.  
 556 2021.00881.

557 Zhiqiang Gao, Shufei Zhang, Kaizhu Huang, Qiufeng Wang, and Chaoliang Zhong. Gradient  
 558 distribution alignment certificates better adversarial domain adaptation. In *Proceedings of the*  
 559 *IEEE/CVF international conference on computer vision*, pp. 8937–8946, 2021b.

560

561 Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n)  
 562 equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34:4181–  
 563 4192, 2021.

564 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 565 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the*  
 566 *ACM*, 63(11):139–144, 2020.

567

568 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola.  
 569 A kernel two-sample test. *The journal of machine learning research*, 13(1):723–773, 2012.

570

571 Jubum Han, Mateusz Matuszewski, Olaf Sikorski, Hosang Sung, and Hoonyoung Cho. Randmask-  
 572 ing augment: A simple and randomized data augmentation for acoustic scene classification. In  
 573 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing*  
 (ICASSP), pp. 1–5. IEEE, 2023.

574

575 Shady Abu Hussein, Tom Tirer, and Raja Giryes. Image-adaptive gan based reconstruction. In  
 576 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3121–3129, 2020.

577

578 Jongheon Jeong and Jinwoo Shin. Consistency regularization for certified robustness of smoothed  
 579 classifiers. *Advances in Neural Information Processing Systems*, 33:10558–10570, 2020.

580

581 Arsham Gholamzadeh Khoei, Yinan Yu, and Robert Feldt. Domain generalization through meta-  
 582 learning: a survey. *Artificial Intelligence Review*, 57(10):285, 2024.

583

584 Yuma Koizumi, Yohei Kawaguchi, Keisuke Imoto, Toshiaki Nakamura, Yuki Nikaido, Ryo Tanabe,  
 585 Harsh Purohit, Kaori Suefusa, Takashi Endo, Masahiro Yasuda, et al. Description and discussion  
 on dcase2020 challenge task2: Unsupervised anomalous sound detection for machine condition  
 monitoring. *arXiv preprint arXiv:2006.05822*, 2020.

586

587 Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind the  
 588 gap: Understanding the modality gap in multi-modal contrastive representation learning. *Ad-  
 589 vances in Neural Information Processing Systems*, 35:17612–17625, 2022.

590

591 Chuang Liu, Yichao Cao, Xiu Su, and Haogang Zhu. Universal frequency domain perturbation for  
 592 single-source domain generalization. In *Proceedings of the 32nd ACM International Conference*  
 593 *on Multimedia*, pp. 6250–6259, 2024.

594

595 Alhassan Mumuni, Fuseini Mumuni, and Nana Kobina Gerrar. A survey of synthetic data augmen-  
 596 tation methods in machine vision. *Machine Intelligence Research*, 21(5):831–869, 2024.

594 Bach Hoai Nguyen, Bing Xue, Peter Andrae, and Mengjie Zhang. Evolutionary instance selection  
 595 with multiple partial adaptive classifiers for domain adaptation. *IEEE Transactions on Evolutionary  
 596 Computation*, 2023.

597 Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. Cold decoding: Energy-based con-  
 598 strained text generation with langevin dynamics. *Advances in Neural Information Processing  
 599 Systems*, 35:9538–9551, 2022.

600 Luc PJ Sträter, Mohammadreza Salehi, Efstratios Gavves, Cees GM Snoek, and Yuki M Asano.  
 601 Generalad: Anomaly detection across domains by attending to distorted features. In *European  
 602 Conference on Computer Vision*, pp. 448–465. Springer, 2024.

603 KSV Swarna, Arangarajan Vinayagam, M Belsam Jeba Ananth, P Venkatesh Kumar, Veerapandiyan  
 604 Veerasamy, and Padmavathi Radhakrishnan. A knn based random subspace ensemble classifier  
 605 for detection and discrimination of high impedance fault in pv integrated power network. *Mea-  
 606 surement*, 187:110333, 2022.

607 Shixiang Tang, Peng Su, Dapeng Chen, and Wanli Ouyang. Gradient regularized contrastive learning  
 608 for continual domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*,  
 609 volume 35, pp. 2665–2673, 2021.

610 Arun Venkatraman, Martial Hebert, and J Bagnell. Improving multi-step prediction of learned time  
 611 series models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

612 Zi-Ming Wang, Nan Xue, Ling Lei, Rebecka Jörnsten, and Gui-Song Xia. Partial distribution match-  
 613 ing via partial wasserstein adversarial networks. *IEEE Transactions on Pattern Analysis and Ma-  
 614 chine Intelligence*, 2025.

615 Jiacheng Wei, Qibin Wang, Guowei Zhang, Hongbo Ma, and Yi Wang. Domain knowledge guided  
 616 pseudo-label generation framework for semi-supervised domain generalization fault diagnosis.  
 617 *Advanced Engineering Informatics*, 67:103540, 2025a.

618 Shun Wei, Jielin Jiang, and Xiaolong Xu. Uninet: A contrastive learning-guided unified framework  
 619 with feature selection for anomaly detection. In *Proceedings of the Computer Vision and Pattern  
 620 Recognition Conference*, pp. 9994–10003, 2025b.

621 Lisheng Wen, Sentao Chen, Zijie Hong, and Lin Zheng. Maximum likelihood weight estimation for  
 622 partial domain adaptation. *Information Sciences*, 676:120800, 2024.

623 Matthias Wieler, Tobias Hahn, and Fred A. Hamprecht. Weakly supervised learning for indus-  
 624 trial optical inspection. In *Pattern Recognition: 29th DAGM Symposium, Heidelberg, Germany,  
 625 September 12–14, 2007, Proceedings*, volume 4713 of *Lecture Notes in Computer Science*, Berlin,  
 626 Heidelberg, 2007. Springer. Competition/dataset description; data available at HCI or Zenodo.

627 Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P  
 628 Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with  
 629 diffusion priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern  
 630 recognition*, pp. 21551–21561, 2024.

631 Deqiang Xie, Jianxin Yi, Xianrong Wan, and Hao Jiang. Masking effect mitigation for fm-based  
 632 passive radar via nonlinear sparse recovery. *IEEE Transactions on Aerospace and Electronic  
 633 Systems*, 59(6):8246–8262, 2023.

634 Chi Zhang, Cheng Jingpu, Yanyu Xu, and Qianxiao Li. Parameter-efficient fine-tuning with controls.  
 635 In *Forty-first International Conference on Machine Learning*, 2024.

636 Guowei Zhang, Xianguang Kong, Hongbo Ma, Qibin Wang, Jingli Du, and Jinrui Wang. Dual dis-  
 637 entanglement domain generalization method for rotating machinery fault diagnosis. *Mechanical  
 638 Systems and Signal Processing*, 228:112460, 2025a.

639 Guowei Zhang, Xianguang Kong, Qibin Wang, Jingli Du, and Hongbo Ma. Balancing discrepancy  
 640 and consistency: Adversarial single domain generalization in fault diagnosis. *IEEE Transactions  
 641 on Industrial Informatics*, 2025b.

648 Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets  
 649 for partial domain adaptation. In *Proceedings of the IEEE conference on computer vision and*  
 650 *pattern recognition*, pp. 8156–8164, 2018.

651

652 Jiayuan Zhou, Lijun Wang, Huchuan Lu, Kaining Huang, Xinchu Shi, and Bocong Liu. Mvsalnet:  
 653 Multi-view augmentation for rgb-d salient object detection. In *European Conference on Computer*  
 654 *Vision*, pp. 270–287. Springer, 2022.

655

656 Lihua Zhou, Mao Ye, Xinpeng Li, Ce Zhu, Yiguang Liu, and Xue Li. Disentanglement then recon-  
 657 struction: Unsupervised domain adaptation by twice distribution alignments. *Expert Systems with*  
 658 *Applications*, 237:121498, 2024.

659

660 Wenlve Zhou and Zhiheng Zhou. Unsupervised domain adaption harnessing vision-language pre-  
 661 training. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(9):8201–8214,  
 662 2024. doi: 10.1109/TCSVT.2024.3391304.

663

664 Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Graph contrastive learning  
 665 with adaptive augmentation. In *Proceedings of the web conference 2021*, pp. 2069–2080, 2021.

666

667

## A APPENDIX

668 The supplementary material provides additional information, including the theoretical foundations  
 669 of Domain-Aware Gradient Reuse (*Section A.1*), datasets descriptions (*Section A.2*), further com-  
 670 parative experiments (*Section A.3*) and disclosure of language model assistance (*Section A.4*). The  
 671 source code for the proposed method is provided in the `code` directory.

### A.1 THEORETICAL FOUNDATIONS

672 In this section, **Supplementary A.1.1** demonstrates a *cross-class generalisation* property for the  
 673 proposed de-domain mapping. After the mapping is trained *exclusively* on normal-class gradients, it  
 674 is able to project source-domain abnormal gradients so that their distribution matches that of the (un-  
 675 observed) target-domain abnormal gradients up to the same small tolerance level. **Supplementary**  
 676 **A.1.2** leverages this result to study the optimisation trajectory of the Domain-aware Gradient Reuse  
 677 (DAGR) algorithm. It is shown that, at every training step, the surrogate gradient used by DAGR  
 678 differs from the exact target-domain gradient by a *uniform* and *time-independent* margin whose size  
 679 is the sum of the tolerances proved in Supplementary A.1.1 and a bounded domain-specific pertur-  
 680 bation term. Under standard smoothness and Polyak–Łojasiewicz conditions on the loss function,  
 681 classical inexact-gradient descent theory (Bertsekas, 1999) then guarantees that the model parame-  
 682 ters converge to a neighbourhood of the optimum whose radius is proportional to the square of this  
 683 margin, and—crucially—this error does *not* accumulate over epochs.

684 Together, the two appendices provide a complete theoretical foundation for DAGR, simultaneously  
 685 validating the reuse of source abnormal gradients and establishing the global convergence of the  
 686 training procedure.

#### A.1.1 PROOF OF CROSS-CLASS GENERALISATION

687 **Purpose.** This section proves that, under a *label-independent shift*, a mapping  $\mathcal{F}$  learnt *solely*  
 688 from **normal-class** gradients aligns **abnormal-class** gradients to the target domain with the same  
 689 error upper-bound  $\varepsilon$ . This result substantiates the main-paper strategy of re-using source abnormal  
 690 information—encoded in  $\nabla\Phi^-$ —even though the target domain contains no abnormal samples. This  
 691 justifies the main-paper statement:

$$692 \mathcal{F}(\nabla\Phi^+) \approx \nabla\Psi^+ \implies \mathcal{F}(\nabla\Phi^-) \approx \nabla\Psi^-,$$

693 and explains why the abnormal-class information contained in  $\nabla\Phi^-$  can be safely reused in the  
 694 target domain even when no abnormal samples are available there.

**Symbols and Decomposition.** Let the  $K$ -layer source model have parameters  $\Phi = \{\phi_1, \dots, \phi_K\}$  and the target model  $\Psi = \{\psi_1, \dots, \psi_K\}$ . For any mini-batch we obtain *expected* gradients  $\nabla\Phi^+$ ,  $\nabla\Phi^-$  (source, normal / abnormal) and  $\nabla\Psi^+$  (target, normal). Following the main paper, *each* gradient splits into a domain-invariant component  $\Omega$  and a domain-specific perturbation:

$$\nabla\Phi = \Omega + \nabla\Phi_a, \quad \nabla\Psi = \Omega + \nabla\Psi_a.$$

We further write  $\Omega^+$  (normal) and  $\Omega^-$  (abnormal); note that  $\Omega^-$  exactly coincides with what the main text formerly denoted  $\nabla\Phi_g^- = \nabla\Psi_g^-$ , i.e. the abnormal but *domain-invariant* gradient component shared by both domains.

The random variables  $G_\Phi^y$  and  $G_\Psi^y$  ( $y \in \{+, -\}$ ) represent per-sample gradients whose distributions are  $\mathcal{P}_\Phi^y$  and  $\mathcal{P}_\Psi^y$ . A 1-Lipschitz distance  $\mathcal{D}(\cdot, \cdot)$ —concretely the kernel Maximum Mean Discrepancy (MMD; see (Gretton et al., 2012))—measures distribution gaps.

**Label-Independent Shift Assumption.** [Uniform Translation–Perturbation] There exists an *invertible* map  $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$  such that

$$G_\Psi^y = T(G_\Phi^y), \quad \forall y \in \{+, -\}. \quad (\text{A1})$$

Because the same  $T$  applies to both labels, it transports the whole pair  $(\Omega^y, \nabla\Phi_a^y)$  to  $(\Omega^y, \nabla\Psi_a^y)$  without altering  $\Omega^y$ . Empirical t-SNE plots in Fig. 1 verify this behaviour.

**Learning  $F$  from Normal Gradients Only.** With access to  $\mathcal{P}_\Phi^+$  and  $\mathcal{P}_\Psi^+$  we solve

$$\theta^* = \arg \min_{\theta} \mathcal{D}(F_\theta(\mathcal{P}_\Phi^+), \mathcal{P}_\Psi^+), \quad (19)$$

producing a *de-domain* mapping  $F_{\theta^*}$ . Its residual normal-class mismatch is

$$\varepsilon = \mathcal{D}(F_{\theta^*}(\mathcal{P}_\Phi^+), \mathcal{P}_\Psi^+).$$

MMD ensures  $\mathbb{E}[\varepsilon] = O(N^{-1/2})$  with  $N$  normal samples (Gretton et al., 2012).

**Cross-Class Generalisation Theorem.** Under Assumption A1 and with  $F_{\theta^*}$  from equation 19,

$$\mathcal{D}(F_{\theta^*}(\mathcal{P}_\Phi^-), \mathcal{P}_\Psi^-) \leq \varepsilon.$$

Triangle inequality yields

$$\mathcal{D}(F_{\theta^*}(\mathcal{P}_\Phi^-), \mathcal{P}_\Psi^-) \leq \mathcal{D}(F_{\theta^*}(\mathcal{P}_\Phi^-), T(\mathcal{P}_\Phi^-)) + 0,$$

where the zero comes from Assumption A1. Because  $\mathcal{D}$  is 1-Lipschitz,  $\mathcal{D}(F_{\theta^*}(z), T(z)) \leq \|F_{\theta^*}(z) - T(z)\|_2$ . Let  $\delta = \sup_z \|F_{\theta^*}(z) - T(z)\|_2$ ; the same reasoning on the *normal* class gives  $\varepsilon \leq \delta$ , hence the abnormal-class distance is bounded by  $\varepsilon$ .

**Relation to  $\Omega^-$ .** Because  $T$  preserves the invariant part,  $T(\Omega^-) = \Omega^-$ . Applying  $F_{\theta^*}$  to source abnormal gradients gives

$$F_{\theta^*}(\nabla\Phi^-) = F_{\theta^*}(\Omega^- + \nabla\Phi_a^-) \approx \Omega^-,$$

up to error  $\varepsilon$ . Thus the mapped gradient contains (almost) exclusively the domain-invariant abnormal component  $\Omega^-$ , meeting the requirement expressed in the main paper as  $F(\nabla\Phi^-) = \nabla\Phi_g^- \approx \nabla\Psi_g^- = \Omega^-$ .

**Implications.** By Theorem A.1.1,

$$\Omega^- \approx F_{\theta^*}(\nabla\Phi^-) =: \nabla\Phi_g^-,$$

which can be injected into the target update rule, despite the absence of abnormal target samples. Combining Ben-David’s risk bound (Ben-David et al., 2006) with the fact that equation 19 shrinks the domain distance for *both* classes guarantees safe transfer. Gradient-space alignment has empirically achieved lower domain discrepancies than feature-space alignment (Gao et al., 2021b), supporting our choice of operating in gradient space.

756 A.1.2 CONVERGENCE ANALYSIS OF DAGR  
757

758 **Purpose.** DAGR updates the **target-domain** parameters  $\Psi = \{\psi_1, \dots, \psi_K\}$  by a surrogate gradi-  
759 ent  $g_t = \mathcal{F}(\nabla\Phi^+) + F(\nabla\Phi^-)$ , because *abnormal* target samples are absent. To justify its reliability,  
760 we prove that (i)  $g_t$  deviates from the *true* target gradient  $\nabla L(\Psi_t)$  by a *uniform* bound, and (ii) the  
761 resulting inexact-descent iterates  $\{\Psi_t\}_{t \geq 0}$  converge to an  $\mathcal{O}(\varepsilon)$  neighbourhood of the optimum with-  
762 out error accumulation.

763 Notation (identical to Supplementary A.1.1)  
764

- 765 • Source parameters  $\Phi = \{\phi_1, \dots, \phi_K\}$ , target parameters  $\Psi$ .
- 766 • Per-label gradients decompose as  $\nabla\Phi^y = \Omega^y + \nabla\Phi_a^y$  and  $\nabla\Psi^y = \Omega^y + \nabla\Psi_a^y$ ,  $y \in \{+, -\}$ .
- 767 •  $F$  is the *de-domain* map from Supplementary A, trained with normal data; its precision  
768 satisfies  $\|F(\nabla\Phi^y) - \Omega^y\|_2 \leq \varepsilon = O(N^{-1/2})$ .
- 769 •  $L(\Psi)$  denotes the empirical target-domain loss, assumed  $L$ -smooth and  $\mu$ -PL.
- 770

772 **Assumptions A1 (label-independent shift).** There exists an invertible  $T$  transporting all  
773 class-conditional gradient distributions (see Supplementary A).

774 **A2 (bounded perturbations).**  $\|\nabla\Phi_a^y\|_2, \|\nabla\Psi_a^y\|_2 \leq B$  and  $\mathbb{E}[\nabla\Phi_a^y] = \mathbb{E}[\nabla\Psi_a^y] = \mathbf{0}$ .

775 **A3 (loss geometry).**  $L(\cdot)$  is  $L$ -smooth and satisfies the PL-inequality with constant  $\mu$ .

776 **One-Step Gradient Error** The true full gradient at step  $t$  is  $\nabla L(\Psi_t) = \Omega^+ + \Omega^- + \nabla\Psi_{a,t}^+ + \nabla\Psi_{a,t}^-$ ,  
777 whereas DAGR uses  $g_t$ . A direct triangle-inequality gives

$$778 \|g_t - \nabla L(\Psi_t)\|_2 \leq 2\varepsilon + 2B = \delta, \\ 779$$

780 which is *constant* in  $t$ . Hence no per-iteration error growth can occur.

781 **Inexact-Gradient Descent Dynamics** With step-size  $\eta \leq 1/(2L)$ , DAGR performs  $\Psi_{t+1} = \Psi_t -$   
782  $\eta g_t$ . Using standard inexact-descent analysis (Bertsekas, 1997) we obtain

$$783 L(\Psi_{t+1}) - L(\Psi^*) \leq (1 - \eta\mu)[L(\Psi_t) - L(\Psi^*)] + C\eta\delta^2, \\ 784$$

785 where  $C < 2$  is universal. Telescoping over  $T$  steps and letting  $T \rightarrow \infty$  yields the steady-state bound

$$786 \limsup_{t \rightarrow \infty} [L(\Psi_t) - L(\Psi^*)] \leq \frac{C\delta^2}{\mu} = \mathcal{O}((\varepsilon + B)^2). \\ 787$$

788 Consequences  
789

- 790 • **Convergence.** DAGR reaches an error ball whose radius contracts with  $\varepsilon$ ; increasing  
791 normal-sample pairs tightens the bound.
- 792 • **No accumulation.**  $\delta$  is independent of  $t$ , so the error term in each step is con-  
793 stant—boundedness is preserved over indefinite epochs.
- 794 • **Practical implication.** When  $B$  is empirically small (consistent dispersion of Fig. 1) and  
795  $N$  large, DAGR approaches the optimum as closely as exact gradient descent.

800 A.2 DATASETS DESCRIPTIONS  
801

802 **DCASE 2020 Task 2 Benchmark** The DCASE 2020 Challenge Task 2 dataset(Koizumi et al.,  
803 2020) is a standard benchmark for unsupervised detection of anomalous sounds for machine con-  
804 dition monitoring. It features six distinct machine types: *ToyCar*, *ToyConveyor*, *Valve*, *Pump*, *Fan* and  
805 *Slide rail*. Each recording is a single-channel, 10-second audio clip sampled at 16 kHz. To lever-  
806 age powerful feature extraction techniques from the vision domain, we first transform these one-  
807 dimensional audio signals into two-dimensional spectrograms, which reframes the original acoustic  
808 anomaly detection task into a visual anomaly detection problem. We designate the *Fan* subset as the  
809 source domain. The remaining five machine types—*ToyCar*, *ToyConveyor*, *Valve*, *Pump* and *Slide*  
810 *rail*—are treated as unseen target domains.

810  
 811 Table 3: AUPRC (%) comparison on 8 benchmark datasets. Best result per column is in **bold**,  
 812 second best is underlined.

813 814 Methods	815 816 Target Domain	817 818 DCASE (Fan)					819 820 DAGM (Class 2)			821 822 Ave.
		823 Pump	824 Slider	825 Valve	826 ToyCar	827 ToyConv.	828 Class 1	829 Class 3	830 Class 6	
Unsupervised Methods	General-AD GLASS	66.40 <b>68.32</b>	78.79 75.16	62.77 <b>66.00</b>	48.90 <u>65.18</u>	40.02 <u>53.76</u>	32.81 <u>79.13</u>	37.64 <u>67.01</u>	83.72 43.75	58.63 <b>72.04</b>
Partial Domain Adaptation	JWO PWAN MLWE CMKD	6.50 16.10 6.00 49.14	25.99 20.36 17.77 80.40	14.44 14.68 15.25 59.75	25.19 17.32 14.20 43.45	26.25 24.04 24.34 38.14	20.01 16.07 15.99 17.51	26.72 20.07 14.78 55.15	50.72 22.32 17.82 <u>85.93</u>	24.48 18.87 15.77 54.68
Domain Adaptation	UniNet ANC FFAT	12.00 12.44 16.53	14.84 12.09 37.38	20.17 15.92 26.26	11.61 21.63 32.19	19.31 35.37 34.99	16.86 16.49 22.98	18.54 22.12 41.42	23.01 16.67 70.45	17.04 19.09 35.27
Domain Generalization	GGA BDC DDDGG PMGDD DKGPL	18.43 12.15 11.56 11.94 7.32	16.10 15.58 22.13 12.34 25.68	18.95 15.69 16.68 14.99 14.24	17.13 19.43 13.94 14.24 22.25	31.70 18.02 19.80 21.75 19.98	19.74 16.66 20.81 16.75 17.14	17.43 21.08 20.75 13.91 20.69	42.78 24.19 24.01 35.94 22.54	22.78 17.85 18.71 17.73 18.73
Proposed	DAGR	53.71	<b>87.79</b>	26.52	<b>71.76</b>	<b>71.91</b>	<b>81.77</b>	<b>73.29</b>	<b>86.97</b>	69.22

828  
 829  
 830 Table 4: Rec@K (%) comparison on 8 benchmark datasets. Best result per column is in **bold**, second  
 831 best is underlined.

832 833 Methods	834 835 Target Domain	836 837 DCASE (Fan)					838 839 DAGM (Class 2)			840 841 Ave.
		842 Pump	843 Slider	844 Valve	845 ToyCar	846 ToyConv.	847 Class 1	848 Class 3	849 Class 6	
Unsupervised Methods	General-AD GLASS	58.41 <b>62.34</b>	72.93 71.69	63.52 <b>67.66</b>	49.00 <u>58.74</u>	41.29 <u>45.23</u>	32.00 <u>70.67</u>	41.33 <u>58.67</u>	73.33 41.33	56.46 <b>67.29</b>
Partial Domain Adaptation	JWO PWAN MLWE CMKD	5.80 23.19 2.90 45.14	27.20 26.40 17.60 76.07	10.77 20.77 19.23 58.66	27.33 18.00 12.00 40.89	29.11 27.22 31.01 35.05	14.84 15.79 17.11 16.25	31.08 23.68 10.53 46.25	51.32 23.68 21.05 <u>76.25</u>	24.68 22.34 16.43 50.70
Domain Adaptation	UniNet ANC FFTAT	10.77 17.65 18.20	16.00 6.40 40.60	23.47 17.69 28.59	10.00 23.33 34.53	22.78 35.44 36.99	19.74 11.84 20.51	19.74 26.32 43.59	23.68 15.79 64.10	18.27 19.31 35.89
Domain Generalization	GGA BDC DDDGG PMGDD DKGPL	20.29 19.12 12.70 13.04 4.41	18.40 13.60 28.80 5.60 28.00	21.54 16.92 17.69 16.15 15.38	22.00 25.33 18.00 14.00 25.33	32.91 18.99 24.05 25.32 19.62	21.05 15.79 18.92 15.79 15.79	17.57 28.95 18.92 10.53 18.42	44.74 27.03 24.86 36.84 23.68	24.81 20.72 20.49 17.16 18.83
Proposed	DAGR	53.02	<b>81.16</b>	26.96	<b>66.29</b>	<b>64.46</b>	<b>72.69</b>	<b>67.75</b>	<b>82.27</b>	64.33

847  
 848  
 849 **DAGM Dataset** The DAGM 2007 dataset(Wieler et al., 2007) is a widely-used benchmark de-  
 850 signed for unsupervised visual anomaly detection. The dataset features 10 different classes of  
 851 grayscale texture images, created to simulate various industrial surfaces. These images have a res-  
 852 olution of  $512 \times 512$  pixels and contain several types of artificially generated defects. For each class,  
 853 the dataset provides 1,000 defect-free images for training and 150 images for testing, which may  
 854 or may not contain defects. In our experiment, all defects are treated as anomalies. The *Class 2*  
 855 category served as the source domain, with the *Class 1*, *Class 3*, and *Class 6* categories acting as  
 856 target domains.

### 857 A.3 ADDITIONAL COMPARISON EXPERIMENTS

858 **Overall results.** Tables 3 and 4 report AUPRC and Rec@K on eight transfer tasks. The proposed  
 859 DAGR attains the **best per-column score on 6/8 tasks** under both metrics, and ranks second in  
 860 the overall average. Specifically, in AUPRC DAGR leads on *Slider*, *ToyCar*, *ToyConv.*, *DAGM*-  
 861 *Class 1/3/6*, and in Rec@K it again leads on the same six targets. The average, however, is depressed  
 862 by a single outlier—the *Fan*  $\rightarrow$  *Valve* transfer.

864 **Why the average is lower: the Valve outlier.** On *Valve*, DAGR is substantially weaker  
 865 (AUPRC 26.52, Rec@K 26.96) than on the other targets. This behaviour is expected and consistent  
 866 with the method’s scope: DAGR explicitly relies on *domain proximity* diagnosed from *normal* data  
 867 (Sec. §2/§3), i.e., the cross-domain normal-gradient geometry must be compatible enough to learn  
 868 a reliable transport map. Most DCASE targets (*Pump*, *Slider*, *ToyCar*, *ToyConv*) are *motor-driven*  
 869 *rotating machinery*. Their acoustics are dominated by quasi-stationary tonal components at the ro-  
 870 tation frequency and its sidebands, yielding stable spectra and slowly varying envelopes. *Valve* is  
 871 qualitatively different: its sound is driven by *turbulence*, *flow transients* and *opening/closing events*,  
 872 which are *non-stationary*, bursty and broadband. Consequently, the target normal-gradient subspace  
 873 for *Valve* diverges from that of *Fan*, leading Stage-1 gating to mask many channels and Stage-2 to  
 874 learn a weak transport; the mapped anomalous gradients are then down-weighted by the reliability  
 875 scheme, effectively reverting to normal-only updates. In contrast, unsupervised baselines such as  
 876 GLASS/General-AD train purely on target normals and are not affected by cross-domain incompat-  
 877 ibility, hence their stronger scores on *Valve*.  
 878

879 **Quantifying the outlier effect.** The overall average in Tables 3 and 4 is a simple mean over eight  
 880 targets. For AUPRC, DAGR’s mean is **69.22%**, but *removing the single Valve column for refer-  
 881 ence* raises it to **75.31%**, exceeding the best competing average (72.04%). For **Rec@K**, the mean  
 882 increases from 64.33% to **69.66%** when *Valve* is excluded, again surpassing the best competing aver-  
 883 age (67.29%). These reference numbers are not substitutes for the official average; they merely  
 884 illustrate that the gap is driven by one incompatible target rather than by systematic underperfor-  
 885 mance.  
 886

887 **Summary.** (i) On *domain-proximal* transfers—rotating machinery in DCASE and texture-to-  
 888 texture transfers in DAGM—DAGR consistently outperforms strong baselines. (ii) On *non-proximal*  
 889 transfers such as *Fan*→*Valve*, the cross-domain normal-gradient geometry is not compatible; DAGR  
 890 therefore (by design) attenuates the mapped anomalous component and does not confer an advan-  
 891 tage over target-only unsupervised methods. (iii) This behaviour delineates the *intended operating*  
 892 *regime* of DAGR and aligns with the diagnostics introduced in the main paper. We include *Valve* as  
 893 a negative-control case to make the boundary explicit rather than to optimise it away.  
 894

#### A.4 DISCLOSURE OF LANGUAGE MODEL ASSISTANCE

895 Large language models were used only for editorial polishing (grammar, style, and minor rephras-  
 896 ing). They were not used for research design, methods, analysis, coding, figures/tables, or refer-  
 897 ences. All scientific content was authored and verified by the authors, and all edits were manually  
 898 reviewed. This use does not meet contributorship thresholds and does not affect reproducibility.  
 899

#### A.5 COMPLEXITY AND IMPLEMENTATION

900 **Notation.** Let the  $K$ -layer source and target models have parameters  $\Phi = \{\phi_1, \dots, \phi_K\}$  and  $\Psi =$   
 901  $\{\psi_1, \dots, \psi_K\}$ . Mini-batch gradients are  $\nabla\Phi$  and  $\nabla\Psi$ . CCCS (Sec. 3.2) yields masked gradients  
 902  $\tilde{\nabla}\Phi_{\pm}$  and  $\tilde{\nabla}\Psi_{\pm}$ . ADPR (Sec. 3.3) refines a gradient variable  $g_m^{(n)}$  for  $n = 0, \dots, N$  and outer step  
 903  $m$ , producing  $\Omega_m^+ = g_m^{(N)}$  (Eq. (11)), and reuses the same map for  $\Omega_m^-$ , then aggregates  $\Omega_m =$   
 904  $\Omega_m^+ + \Omega_m^-$  for the target update (Eq. (3)). We follow the main text and denote the fast-weight step in  
 905 Eq. (9) by  $\alpha$ ; when disambiguation is helpful we write  $\alpha_{\text{fw}}$ . The CCCS masking percentile is also  
 906 denoted  $\alpha$  in Sec. 3.2; context makes the meaning clear.  
 907

908 **Per-iteration time complexity.** Let  $C_{\text{fwd}}$  and  $C_{\text{bwd}}$  be the cost of one forward/backward pass of  
 909 the *target* network per mini-batch, and let  $D$  be the number of trainable parameters.  
 910

911 **CCCS.** For each layer  $k$  with channel set  $C_k = \{c_k^1, \dots, c_k^{P_k}\}$ , CCCS computes cosine similarities  
 912  $\rho_{k,p}$  between  $\tilde{\nabla}\Phi_{k,p}^+$  and  $\tilde{\nabla}\Psi_{k,p}^+$  and applies a percentile threshold (Eqs. (6)–(8)). All operations  
 913 are on gradients already available from back-propagation; the added arithmetic is  $O(\sum_{k,p} d_{k,p}) =$   
 914  $O(D)$ , where  $d_{k,p}$  is the parameter count of channel  $c_k^p$ . This overhead is negligible compared with  
 915 a single forward/backward pass.  
 916

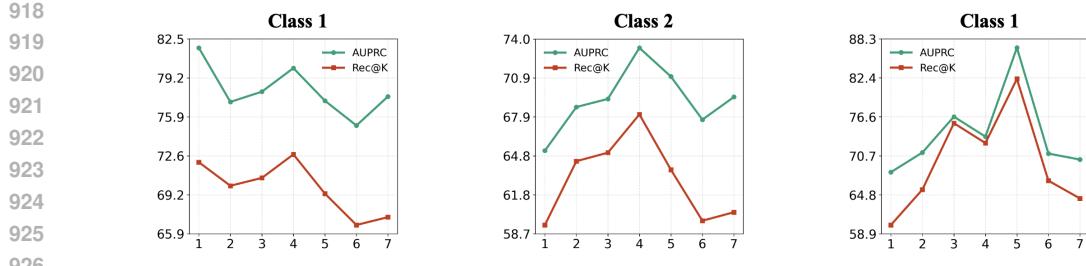


Figure 5: **Effect of channel-masking ratio  $\alpha$  on DAGM transfers.** Detection performance (AUPRC and Rec@K) as  $\alpha$  varies from 1% to 7% on Class 1, Class 3, and Class 6.

*ADPR inner loop.* At outer step  $m$ , set fast weights  $\Psi_{m,0} = \Psi_m + \alpha g_m^{(0)}$  (Eq. (9)). Each inner update  $g_m^{(n+1)} = g_m^{(n)} - \beta \nabla_g L_t(\Psi_m + \alpha g_m^{(n)})$  (Eq. (10)) requires one forward and one backward pass through the *target* network with fast weights. Forming fast weights is an  $O(D)$  axpy. With  $N$  inner steps, the inner-loop cost is

$$T_{\text{ADPR}} = N(C_{\text{fwd}} + C_{\text{bwd}}) + O(D).$$

*Outer update.* A standard target update costs  $(C_{\text{fwd}} + C_{\text{bwd}})$ . Aggregating  $\Omega_m$  (Eq. (12)) and updating  $\Psi$  (Eq. (3)) adds only  $O(D)$  work. Thus the per-outer-step time is

$$T_{\text{DAGR}} = (1 + N)(C_{\text{fwd}} + C_{\text{bwd}}) + O(D),$$

i.e., DAGR incurs a *constant-factor* overhead over standard training.

**Memory complexity.** Let  $M_{\text{base}}$  be the peak activation/optimizer memory of the base detector. DAGR adds: (i) channel-mask buffers  $\{m_{k,p}\}$  (Boolean;  $\sum_k P_k$  entries), (ii) the current gradient variable  $g_m^{(n)} \in \mathbb{R}^D$  and its masked initialization, and (iii) ephemeral activations for the  $N$  inner steps (not accumulated across steps). No per-sample gradients nor second-order tensors are stored. Peak memory therefore satisfies

$$M_{\text{DAGR}} \approx M_{\text{base}} + O(D) + O\left(\sum_k P_k\right),$$

and standard techniques (gradient checkpointing, mixed precision) remain fully applicable.

**Recommended settings (used in the main experiments).** The channel-masking percentile  $\alpha$  is swept from 1% to 7% with peaks typically in 3–5%;  $\alpha = 5\%$  is adopted thereafter (Sec. 4.2, Fig. 4). The inner-loop length  $N$  is kept small (a constant); the inner learning rate  $\beta$  and fast-weight step  $\alpha$  follow standard grids. These choices keep  $T_{\text{DAGR}}$  a modest constant multiple of the baseline cost, as observed empirically.

**Stability and theory link.** Appendix A.1.2 proves that DAGR’s surrogate gradient deviates from the exact target gradient by a time-independent bound, leading to inexact-descent convergence to an  $O((\varepsilon + B)^2)$  neighbourhood without error accumulation; this matches the smooth training curves observed across tasks.

**Practical deployment note (non-parametric).** When transfers involve potentially non-adjacent domains, practitioners may *screen* the proximity of *normal* gradient distributions (e.g., via  $W_1$  or MMD) before enabling gradient reuse. If the divergence is large, a conservative fallback is to disable reuse and proceed with target-only updates. This screening is advisory and does not alter the reported experiments; devising automated reliability weighting is left for future work.

## B EFFECT OF CHANNEL-MASKING RATIO ON DAGM BENCHMARKS

To examine whether the channel-masking ratio  $\alpha$  identified on the DCASE benchmarks is also reasonable for visual-texture domains, we conduct an additional ablation study on the DAGM transfers (Class 2  $\rightarrow$  Class 1/3/6, as shown in Figure 5). For each target class,  $\alpha$  is swept from 1% to 7%, and we report AUPRC and Rec@K averaged over three runs.

972     **Results.** The three DAGM targets exhibit heterogeneous behaviours. One class shows a mostly  
973     decreasing trend as  $\alpha$  grows, whereas the other two classes first benefit from masking a small fraction  
974     of channels and then degrade when too many channels are removed. Across all cases, however, the  
975      $\alpha \in [3\%, 5\%]$  range consistently yields competitive performance: the scores at  $\alpha = 5\%$  are either  
976     close to the per-class maximum (within a small margin) or lie on a relatively flat part of the curve  
977     without sharp deterioration. This indicates that the method is not overly sensitive to the exact choice  
978     of  $\alpha$  in this moderate region, even though the precise optimum is slightly domain dependent.

979  
980     **Conclusion.** Taken together with the DCASE study in the main paper, these experiments suggest  
981     that  $\alpha = 5\%$  is a *robust default* that performs reasonably well across both acoustic (DCASE) and  
982     visual-texture (DAGM) benchmarks, rather than a strictly optimal setting for every single transfer.  
983     In practice, mild per-domain tuning around this range could further improve performance if desired,  
984     but all reported results use the same fixed value  $\alpha = 5\%$  for simplicity and fairness.

985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025