Covariance Alignment for Dynamic Domain Adaptation in Robotic Trajectory Identification Under Variable Loading and Speed Regimes

Md Omar Al Javed[1], Sungkwang Mun[2], Abdullah Al Mamun[3], Nayeon Lee[2], Ayantha Senanayaka[4]*

[1]Mechanical and Nuclear Engineering, Tennessee Technological University, Cookeville, TN 38505, US

[2]Center for Advanced Vehicular Systems, Mississippi State University, Starkville, MS 39762, US

[3]Batten College of Engineering and Technology, Old Dominion University, Norfolk, VA 23529, US

[4]Manufacturing and Industrial Systems Engineering, Tennessee Technological University, Cookeville, TN 38505, US

*Corresponding authors. E-mail: asenanayaka@tntech.edu

Abstract

Industrial robotic trajectory condition monitoring under dynamic operating environments is constrained by covariate shift, where variations in end-effector payload mass and operational speed induce shifts in multi-sensor feature distributions. Supervised classifiers trained on a baseline operating state experience significant performance degradation when transferred to a different loading regime without adaptation. To address this discrepancy without target-domain labels or non-linear parameter estimation, this paper presents an unsupervised domain adaptation framework based on second-order Correlation Alignment (CORAL). The methodology is validated on an experimental parallel delta-robot platform equipped with multi-sensor acoustic and kinematic sensing across light and heavy payload states under low- and high-speed trajectory conditions. Information-theoretic and geometric evaluations show that covariance alignment reduces cross-domain divergence, achieving a 100% reduction in 2-Wasserstein optimal transport distance at both speeds, a reduction in Maximum Mean Discrepancy (MMD) of 59.3% at low speed and 91.5% at high speed, and a corresponding decrease in Domain Silhouette overlap score. Cross-domain trajectory classification accuracy for the decision-tree ensembles is restored from an unadapted baseline of 34.39% to above 98% post-alignment, with low variance across 10-fold cross-validation. These findings demonstrate that linear covariance alignment provides a computationally lightweight and mathematically grounded approach to mitigating payload-induced covariate shift in robotic trajectory monitoring.

Keywords- Unsupervised Domain Adaptation, Correlation Alignment (CORAL), Covariate Shift, Delta Robot, Multi-Sensor Fusion, Robot Trajectory Classification, Machine Health Monitoring.

1. Introduction

High-precision robotic systems are central to modern automated manufacturing, assembly, and flexible industrial production [1, 2]. As automated platforms operate under increasingly dynamic demands, continuous health and trajectory condition monitoring are vital to prevent unscheduled downtime, ensure operational safety, and maintain high task quality [3]. Machine learning architectures have emerged as powerful tools for automated state classification using multi-sensor telemetry, such as acoustic emissions, vibration signals, and kinematic currents [1, 4, 5]. However, deploying data-driven monitoring models in real-world industrial environments remains fundamentally challenged by dynamic covariate shift [6]. In practice, industrial robots routinely execute tasks across variable operational regimes, such as fluctuating end-effector payload masses, changing acceleration profiles, and diverse trajectory speeds [7]. These dynamic variations alter the underlying statistical distribution of multi-sensor features despite the robotic mechanism remaining structurally sound [8]. When a supervised model trained on a baseline operational state (the source domain) is evaluated on a different operating regime (the target domain), the classifier encounters significant feature space misalignment, leading to severe diagnostic performance degradation [9, 10].

To address distribution misalignment across dynamic domains, recent research has shifted toward Unsupervised Domain Adaptation (UDA) [11]. Standard UDA techniques aim to align feature representations between source and target domains without requiring target-domain labels [12]. Deep learning UDA strategies, such as Deep Domain Confusion (DDC) [13], Domain-Adversarial Neural Networks (DANN) [14], and Adversarial Discriminative Domain Adaptation (ADDA) [15], utilize adversarial loss functions or Maximum Mean Discrepancy (MMD) constraints to learn domain-invariant features. Despite their empirical success, existing UDA strategies exhibit key limitations when deployed for real-time systems condition monitoring. Deep adversarial and non-linear UDA networks require extensive parameter optimization, long training iterations, and significant memory overhead, making them poorly suited for edge-level implementation or online adaptation on embedded controllers [16]. Furthermore, adversarial adaptation routines are especially difficult to train, often suffering from minimax instability or mode collapse while risking the distortion of the intrinsic topological manifold of the feature space [17–19]. Deep UDA models also lack explicit mathematical interpretability regarding how statistical moments are transformed between distinct dynamic regimes [20].

To overcome these computational and structural limitations, this paper introduces a lightweight, mathematically grounded UDA framework based on second-order Correlation Alignment (CORAL) [16, 21] for multi-sensor robotic trajectory condition monitoring. Rather than relying on complex non-linear network transformations, the proposed methodology directly minimizes domain divergence by aligning the second-order statistical moments (covariance matrices) of multi-sensor telemetry features between a baseline source payload and an unlabelled target payload across distinct kinematic speed regimes. The novelty of this work lies in the theoretical and empirical demonstration that linear second-order covariance alignment is sufficient to resolve dynamic covariate shift induced by changing robotic payloads and kinematic speeds. By explicitly transforming feature statistics to match target domain covariance structures, the framework eliminates operational domain bias while preserving the localized geometric topology required for accurate pattern classification, all with negligible computational overhead.

The proposed methodology is validated using physical telemetry data collected from an experimental FANUC M-1iA/0.5A parallel delta-robot platform operating under controlled dynamic variations. The platform was subjected to trajectory executions under light and heavy end-effector payload states across two different kinematic linear speeds. Multi-sensor acoustic and kinematic feature matrices were extracted to evaluate domain discrepancy before and after alignment. The framework's efficacy is evaluated using information-theoretic transport metrics, including 2-Wasserstein optimal transport distance [22], Kullback-Leibler Divergence (KLD) [23], and MMD within a Reproducing Kernel Hilbert Space (RKHS) [24], alongside domain separability manifolds, Domain Silhouette Score [25] and cross-domain classification accuracies across multiple machine learning decision ensembles such as Logistic Regression (LR) [26], Random Forest (RF) [27] and Extra Trees (ET) [28].

The primary contributions of this study are threefold:

1. Quantification of Dynamic Domain Shift: We explicitly characterize and quantify the impact of end-effector payload variations and kinematic speed shifts on multi-sensor feature distributions, demonstrating the severe breakdown of standard supervised models under unadapted dynamic transfer.
2. Lightweight Second-Order Adaptation Protocol: We formulate an unsupervised covariance alignment framework that aligns second-order statistical moments between baseline and unlabelled dynamic target states, eliminating domain bias without requiring target labels or complex non-linear deep architectures.
3. Empirical Validation on a Physical Delta Robot: Using experimental telemetry from a FANUC M-1iA/0.5A delta-robot platform, we demonstrate that the proposed framework achieves a complete reduction in 2-Wasserstein transport distance and up to a 91.5% reduction in MMD transport distance, successfully restoring cross-domain trajectory classification accuracy from an unadapted baseline of 34.39% to near-flawless performance (>98%).

2. Literature Review

This review is organized around four connected threads that motivate the proposed framework. Section 2.1 surveys condition monitoring and fault diagnosis in robotic and rotating machinery, establishing the assumption of distributional stationarity underlying most existing diagnostic pipelines. Section 2.2 examines covariate shift as the specific mechanism by which this assumption breaks down under dynamic payload and speed variation, drawing on both the original statistical formulation of the problem and its more recent treatment in mechanical and robotic systems. Section 2.3 reviews unsupervised domain adaptation (UDA), outlining the progression from theoretical generalization bounds through discrepancy-based and adversarial deep learning methods, along with their practical limitations for embedded industrial deployment. Section 2.4 focuses specifically on Correlation Alignment and related second-order statistical alignment methods, situating them relative to the more computationally intensive adversarial alternatives. Section 2.5 synthesizes these threads to identify the specific gap addressed by the present study.

2.1 Condition Monitoring and Fault Diagnosis in Robotic and Rotating Machinery

Condition monitoring of industrial robots and rotating machinery has increasingly relied on multi-sensor telemetry, vibration, acoustic emission, current, and kinematic signals, combined with data-driven classifiers to detect degradation and anomalous operating states before failure [1–4, 18, 29] . Artificial intelligence–based approaches, ranging from classical statistical learning to deep neural architectures, have shown strong diagnostic performance when training and deployment conditions match well [4, 5, 18]. Reviews of data-driven process monitoring methods emphasize that characterizing and mining industrial sensor data remain central to reliable fault detection across manufacturing domains [2, 5]. In contrast, dedicated reviews of industrial robot condition monitoring have classified vibration-, acoustic-, and current-based diagnostic strategies specifically for articulated and parallel robotic mechanisms [30]. Sensor fusion has emerged as a recurring theme in this literature: sound-vibration spectrogram fusion has been used to diagnose RV reducer faults in industrial robots [31], two-stage sound-vibration fusion pipelines have improved weak fault detection sensitivity in rolling bearing systems [32], and multi-input convolutional architectures fusing vibro-acoustic channels have improved fault diagnosis accuracy for induction motors [33].

Complementary work has explored sound-source localization as a robot-mounted condition-monitoring modality [34]. This body of work mainly assumes that the statistical distribution of sensor features remains consistent between a model's training conditions and the conditions in which it is deployed. However, this assumption is often violated in real industrial settings, where payload, speed, and load can vary continuously during normal operations. Predictive maintenance frameworks built over condition-based maintenance (CBM) principles have long recognized operating-condition variability as a central challenge to reliable prognosis [35, 36] , and recent systematic mapping of predictive maintenance practice within Industry 4.0 has identified robust handling of non-stationary operating regimes as an open and actively pursued research priority across sectors [37].

2.2 Covariate Shift and Dynamic Operating Conditions

Covariate shift occurs when the distribution of input features changes between the training (source) domain and the deployment (target) domain, even though the underlying task relationship remains the same. Shimodaira [6] formally addressed this issue by proposing to adjust the log-likelihood function to account for this divergence during model estimation. A broader and more unified treatment of dataset shift has subsequently expanded this formalization to differentiate covariate shift from related phenomena, such as prior probability shift and concept shift, clarifying which correction strategies correspond to each underlying mechanism [38].

In industrial and mechatronic systems, covariate shift manifests concretely as sensor-feature drift induced by changing payload mass, actuation speed, or load torque [7, 39]. Deep model-based domain adaptation approaches have been proposed specifically for fault diagnosis under such varying operating conditions, demonstrating that a classifier trained under one load regime degrades substantially when evaluated under another without explicit correction. Bearing-fault diagnosis under variable working conditions has received particular attention: deep adversarial domain adaptation has been used to diagnose bearing faults in industrial robot joints specifically under varying working conditions [40] , and imbalanced-domain-adaptation frameworks have addressed the compounding difficulty of class imbalance co-occurring with operating-condition shift in bearing diagnostics [41]. Work on remaining-useful-life estimation using time-frequency deep learning representations has shown that operating-condition variability is a persistent confound in physical asset health monitoring [8], and servo-motor prognostics research has echoed this concern at the actuator level, showing that health indicators extracted under one operating condition transfer poorly to another without explicit distributional correction [42]. Broader transfer-learning literature has reinforced this pattern in manufacturing contexts more generally, showing that models trained under one process condition often fail to transfer their diagnostic relationships to another without explicit domain correction [43]. Dedicated literature reviews on transfer learning for rotating machinery fault diagnosis have cataloged the wide range of domain shift sources-speed, load, and sensor placement- that motivate this class of methods [9]. Domain generalization approaches, which seek robustness to unseen operating conditions without any target-domain data at all, represent a related but distinct research direction, with recent applications to tool-wear prediction under limited data [44] and federated condition monitoring across distributed ultrasonic welding equipment [45] illustrating the breadth of current efforts to handle non-stationary operating regimes.

2.3 Unsupervised Domain Adaptation: Statistical and Adversarial Approaches

Transfer learning and domain adaptation have emerged as the dominant frameworks for addressing this mismatch without requiring exhaustive relabeling of every new operating condition [10]. Deep visual domain adaptation surveys broadly organize the field into discrepancy-based, adversarial-based, and reconstruction-based strategies for aligning source and target feature representations [11]. A foundational theoretical treatment by Ben-David et al. [12] established generalization bounds for learning across differing domains, formalizing the intuition that classifier transferability depends on the divergence between source and target marginal distributions. Building on this theoretical grounding, deep learning–based UDA methods have flourished: Deep Domain Confusion (DDC) minimizes a Maximum Mean Discrepancy penalty between domain-specific feature representations within a shared network [13]. Domain-Adversarial Neural Networks (DANN) use a gradient-reversal adversarial objective to promote domain-invariant feature learning [14]. Adversarial Discriminative Domain Adaptation (ADDA) separates the feature encoders of the two domains by adversarially training a domain discriminator [15].

Related adversarial frameworks, including Generative Adversarial Networks more broadly, have also been adapted for representation alignment tasks [17]. At the same time, prototypical-network-based approaches have extended adversarial UDA to few-shot and prototype-matching settings for cross-domain classification [19]. Broader surveys of deep transfer learning have cataloged these architectures collectively, highlighting their strong empirical performance across vision and signal-processing benchmarks [2, 20]. The challenges of training instability are well documented, particularly with adversarial objectives [17]. These challenges often manifest as mode collapse, where the model fails to explore the full range of possible outputs, or non-convergence, where the training process does not settle into a stable solution in minimax scenarios. These phenomena highlight the complexities and difficulties of optimizing models with adversarial interactions. The large number of parameters and lengthy training schedules required by these architectures are often cited as limitations for use in resource-constrained environments or real-time embedded monitoring systems [21].

A related causal-discovery approach applied to the delta robotic platform has recently framed operating-condition invariance as a structural causal inference problem rather than a statistical alignment problem, using constraint-based graph recovery to identify and correct payload-induced confounding pathways directly [46]. This complements the second-order alignment strategy pursued here, suggesting causal and statistical approaches to covariate shift correction may be usefully combined in future work.

2.4 Correlation Alignment and Second-Order Statistical Methods

An alternative, lighter-weight line of work has sought to address domain shift by explicitly aligning low-order feature statistics rather than learning non-linear transformations. Correlation Alignment (CORAL) aligns the second-order statistics (the covariance structure) of source and target feature distributions through a simple linear whitening-and-recoloring transformation, requiring no target labels and no iterative optimization [21]. Its deep extension, Deep CORAL, incorporates this covariance-alignment penalty directly into a deep network's loss function, showing that a second-order statistical constraint can achieve competitive domain-invariant performance relative to adversarial alternatives when embedded within a larger deep architecture [21]. Related second-order and kernel-based alignment strategies, including Transfer Component Analysis, seek a shared latent subspace in which the marginal distributions of the source and target domains are minimized under a Maximum Mean Discrepancy criterion [10] . Joint Adaptation Networks have further extended discrepancy-based alignment by jointly matching the distributions of multiple task-specific layers rather than a single feature layer [47]. Correlation Alignment has been reformulated within a Riemannian geometry framework to sufficiently estimate the manifold structure of covariance matrices during alignment [48].

In the specific context of machinery diagnostics, correlation-alignment-based methods have been combined with generative adversarial architectures for cross-device anomaly detection in industrial robot joints [49], and with supervised contrastive learning objectives for bearing and gear fault diagnosis under multiple working conditions [50]. Other work has fused vibro-acoustic sensing modalities with domain-adaptive residual networks [51] and vision-transformer architectures [52] to address cross-domain mechanical fault classification directly. Complementary studies using Wasserstein-distance-based representation learning indicate that geometry-aware optimal transport metrics can function as both domain adaptation objectives and diagnostic measures of domain discrepancy [22]. In parallel manufacturing contexts, domain adaptation between heterogeneous time-series representations has been applied to real-time rotary machinery fault diagnosis [53], and feature-structured domain adaptation has been used for cross-working-condition quality prediction in industrial process data [54].

2.5 Research Gaps and Challenges

Across this literature, condition monitoring pipelines are predominantly developed and validated under an implicit assumption of distributional stationarity between training and deployment. Where this assumption is explicitly relaxed, the correction strategies proposed have excessively relied on deep or adversarial UDA architectures. These architectures demonstrate strong empirical alignment, but they come with notable drawbacks. Adversarial training often suffers from instability and mode collapse. Additionally, the high parameter counts and lengthy optimization schedules can hinder real-time or edge-level deployment. Furthermore, their end-to-end, non-linear transformations provide limited insight into which statistical properties of the sensor feature space are actually modified during the adaptation process. This tension is especially important in the specific failure mode discussed in Section 2.2: sensor drift induced by payload, speed, and load. This type of drift occurs in bearings, motors, welding equipment, and robotic joints. Although the literature often describes it as an operational problem requiring near-real-time monitoring, it is frequently addressed with architectures poorly suited to this context.

A second gap concerns the validation scope of lighter-weight, second-order alignment methods. Prior applications of Correlation Alignment to machinery diagnostics have predominantly embedded it as an auxiliary loss term within larger deep architectures, or paired it with adversarial or generative components, rather than evaluating it as a standalone alignment mechanism. As discussed in sections 2.3 and 2.4, second-order and structured statistical alignment methods have been validated independently of deep learning architectures. However, this validation has largely occurred using heterogeneous time-series or quality-prediction data from relatively stable manufacturing processes. In contrast, less attention has focused on multi-sensor telemetry collected under systematically varied and controlled dynamic loading and speed conditions on physical robotic platforms. Joint-adaptation refinements to discrepancy-based alignment add architectural and computational complexity, rather than testing whether a simpler linear formulation suffices for a specific type of shift. As a result, the existing literature does not clearly indicate whether a purely linear, second-order alignment can effectively address the physically constrained covariate shift caused by variations in payload and speed in a rigid robotic mechanism. This is particularly relevant when applied directly to interpretable multi-sensor telemetry, without adversarial, generative, or deep non-linear components. It remains uncertain whether such shifts inherently demand the higher-capacity architectures that currently dominate this field.

This study addresses important gaps by focusing on second-order Correlation Alignment as a standalone adaptation mechanism. It validates this approach using real acoustic and kinematic telemetry data from an experimental delta-robot platform while systematically varying payload and speed conditions. This research addresses two previously identified limitations: it examines linear covariance alignment independently of deep learning or adversarial techniques, and it does so under controlled dynamic loading and speed conditions on actual robotic hardware, rather than relying on vision benchmarks or simulated fault scenarios. By doing so, the study directly tests whether mechanically induced covariate shifts can be resolved without nonlinear or adversarial architectures that typically dominate the current unsupervised domain adaptation literature.

3. Methodology and Theoretical Framework

This section presents the mathematical foundation, domain adaptation theory, and empirical evaluation metrics governing the proposed diagnostic pipeline. The framework addresses payload-induced covariate shifts in multi-sensor telemetry using second-order CORAL, establishing both feature-level alignment and downstream classifier stability. The proposed methodology has shown in Figure 1.

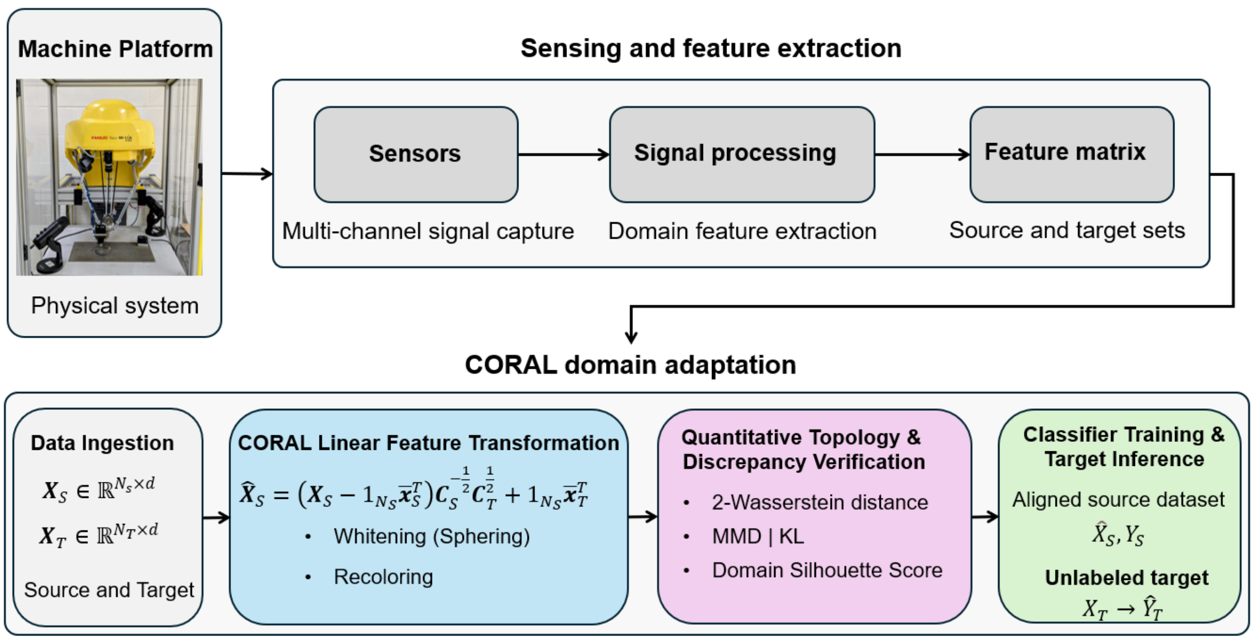

Figure 1: Schematic workflow of the proposed Unsupervised Domain Adaptation (CORAL) framework

Figure 1 overview of the proposed CORAL-based unsupervised domain adaptation pipeline, showing the sequential stages from multi-sensor telemetry acquisition and feature extraction through covariance-based source-to-target alignment to downstream classifier training and cross-domain evaluation.

3.1. Problem Formulation and Domain Shift Mechanics

Let $\mathcal{X} \subset \mathbb{R}^d$represent a d-dimensional feature space derived from multi-sensor time-domain during dynamic trajectory execution. Two operational domain regimes are defined based on dynamic payload mass variations: Source Domain ($\mathcal{D}_S$): A set of N_S fully labeled observations recorded under a baseline payload regime.

$$\mathcal{D}_S = \{(\boldsymbol{x}_i^S, y_i^S)\}_{i=1}^{N_S}, \boldsymbol{x}_i^S \in \mathbb{R}^d, y_i^S \in \mathcal{Y} \quad (1)$$

Target Domain ($\mathcal{D}_T$): A set of N_T unlabeled observations recorded under a heavy payload regime:

$$\mathcal{D}_T = \{\boldsymbol{x}_j^T\}_{j=1}^{N_T}, \boldsymbol{x}_j^T \in \mathbb{R}^d \quad (2)$$

where $\mathcal{Y} = 1,2, \dots, K$ denotes the set of Kdiscrete trajectory condition classes. The system operates across two kinematic linear speed profiles v. Changes in payload mass introduce a significant marginal distribution divergence between the source and target feature spaces [10]:

$$P(\boldsymbol{X}_S) \neq Q(\boldsymbol{X}_T) \tag{3}$$

Under the covariate shift assumption [6], the posterior class conditional probability distributions remain invariant across domains:

$$P(Y_S|\boldsymbol{X}_S = \boldsymbol{x}) = Q(Y_T|\boldsymbol{X}_T = \boldsymbol{x}) \tag{4}$$

The objective is to construct a linear feature transformation matrix $\boldsymbol{A} \in \mathbb{R}^{d\times d}$ that maps the source matrix $\boldsymbol{X}_S \in \mathbb{R}^{N_S\times d}$ into an aligned space $\hat{X}_S \in \mathbb{R}^{N_S\times d}$ such that $P(\widehat{\boldsymbol{X}}_S) \approx Q(\boldsymbol{X}_T)$, enabling a classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ trained on $(\widehat{\boldsymbol{X}}_S, Y_S)$ to accurately generalize to $\boldsymbol{X}_T$.

3.2. Second-Order Correlation Alignment (CORAL) Framework

The CORAL [16, 21] algorithm mitigates domain discrepancy by explicitly aligning the second-order feature statistics (covariance structures) of the source and target distributions without requiring target domain labels or non-linear optimization. Let $\overline{\boldsymbol{x}}_S \in \mathbb{R}^d$ and $\overline{\boldsymbol{x}}_T \in \mathbb{R}^d$ represent the empirical sample mean vectors of the source and target feature spaces, respectively. The regularized sample covariance matrices $\boldsymbol{C}_S \in \mathbb{R}^{d\times d}$and $\boldsymbol{C}_T \in \mathbb{R}^{d\times d}$ are computed as:

$$\boldsymbol{C}_S = \frac{1}{N_S - 1}(\boldsymbol{X}_S - 1_{N_S}\overline{\boldsymbol{x}}_S^T)^T\left(\boldsymbol{X}_S - 1_{N_S}\overline{\boldsymbol{x}}_S^T\right) + \lambda \boldsymbol{I}_d \tag{5}$$

$$\boldsymbol{C}_T = \frac{1}{N_T - 1}(\boldsymbol{X}_T - 1_{N_T}\overline{\boldsymbol{x}}_T^T)^T\left(\boldsymbol{X}_T - 1_{N_T}\overline{\boldsymbol{x}}_T^T\right) + \lambda \boldsymbol{I}_d \tag{6}$$

where $1_{N_T} \in \mathbb{R}^N$ is a column vector of ones, $\boldsymbol{I}_d \in \mathbb{R}^{d\times d}$ is the identity matrix, and $\lambda = 10^{-5}$ is a Tikhonov regularization factor [55] that ensures positive-definiteness and numerical stability during matrix inversion. To satisfy the target covariance matching condition $Cov\left(\widehat{\boldsymbol{X}}_S\right) = \boldsymbol{C}_T$, the transformation matrix $\boldsymbol{A}$ performs a two-step linear operation: Sphering (Whitening): The source covariance is multiplied by $\boldsymbol{C}_S^{-\frac{1}{2}}$ to decorrelate underlying feature variables. Recoloring: The decorrelated source space is multiplied by $\boldsymbol{C}_T^{\frac{1}{2}}$ to inject the second-order correlation structure of the target domain [21]. The linear transformation matrix $\boldsymbol{A}$ is expressed analytically as:

$$\boldsymbol{A} = \boldsymbol{C}_S^{-\frac{1}{2}}\boldsymbol{C}_T^{\frac{1}{2}} \tag{7}$$

Applying $\boldsymbol{A}$ and shifting the mean to match the target distribution yields the aligned source feature matrix $\widehat{\boldsymbol{X}}_S$:

$$\widehat{\boldsymbol{X}}_S = \left(\boldsymbol{X}_S - 1_{N_S}\overline{\boldsymbol{x}}_S^T\right)\boldsymbol{C}_S^{-\frac{1}{2}}\boldsymbol{C}_T^{\frac{1}{2}} + 1_{N_S}\overline{\boldsymbol{x}}_T^T \tag{8}$$

Fractional matrix powers $\boldsymbol{C}^{-\frac{1}{2}}$ and $\boldsymbol{C}^{\frac{1}{2}}$ are evaluated via Singular Value Decomposition (SVD). For a symmetric matrix $\boldsymbol{C} = \boldsymbol{U\Sigma U}^T$:

$$\boldsymbol{C}^r = \boldsymbol{U\Sigma}^r\boldsymbol{U}^T, r \in \{-\tfrac{1}{2}, \tfrac{1}{2}\} \tag{9}$$

Since $\boldsymbol{C}$ is a symmetric positive-definite matrix by construction, guaranteed by the Tikhonov regularization term $\lambda\boldsymbol{I}_d$ introduced in Equations 5 and 6, its singular value decomposition coincides with its eigendecomposition, with $\boldsymbol{U} \in \mathbb{R}^{d\times d}$ the orthogonal matrix of eigenvectors and $\sum \in \mathbb{R}^{d\times d}$ the diagonal matrix of corresponding strictly positive eigenvalues. The matrix power $\boldsymbol{\Sigma}^r$ is computed by raising each diagonal entry of $\boldsymbol{\Sigma}$ to the power r element-wise, ensuring $\boldsymbol{C}^{-\frac{1}{2}}$ and $\boldsymbol{C}^{\frac{1}{2}}$ are both well-defined and numerically stable.

3.3. Quantitative Discrepancy & Feature Topology Metrics

To evaluate distribution matching and manifold overlap between the source and target domains pre- and post-alignment, four complementary metrics are used.

3.3.1. 2-Wasserstein Optimal Transport Distance ($\boldsymbol{W_2}$)

Assuming continuous feature distributions parameterized by their respective means and covariances $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{C})$, the 2-Wasserstein distance provides a closed-form evaluation of optimal transport cost [22]:

$$W_2^2(P, Q) = \| \overline{\boldsymbol{x}}_S - \overline{\boldsymbol{x}}_T \|_2^2 + Tr\left(\boldsymbol{C}_S + \boldsymbol{C}_T - 2(\boldsymbol{C}_S^{\frac{1}{2}} \boldsymbol{C}_T \boldsymbol{C}_S^{\frac{1}{2}})^{\frac{1}{2}}\right) \tag{10}$$

where $Tr(.)$ denotes the matrix trace operator and $\| . \|_2$ represents the Euclidean norm.

3.3.2. Maximum Mean Discrepancy (MMD)

Non-parametric distance between feature distributions is computed in a Reproducing Kernel Hilbert Space (RKHS) $\mathcal{H}_k$ using a Radial Basis Function (RBF) kernel $k(\boldsymbol{x}, \boldsymbol{x}') = \exp(-\gamma \| \boldsymbol{x} - \boldsymbol{x}' \|_2^2)$ with kernel bandwidth $\gamma = 0.1$:

$$MMD^2(\mathcal{D}_S, \mathcal{D}_T) = \frac{1}{N_S^2} \sum_{i=1}^{N_S} \sum_{i'=1}^{N_S} k\left(x_i^S, x_{i'}^S\right) + \frac{1}{N_T^2} \sum_{j=1}^{N_T} \sum_{j'=1}^{N_T} k\left(x_j^T, x_{j'}^T\right) - \frac{2}{N_S N_T} \sum_{i=1}^{N_S} \sum_{j=1}^{N_T} k\left(x_i^S, x_j^T\right) \tag{11}$$

where $\boldsymbol{x}$ and $\boldsymbol{x}'$ denote arbitrary points in the feature space $\mathbb{R}^d$. In Equation 11, this kernel is evaluated between all pairwise combinations of samples within and across domains, where the primed indices (i', j') index a second sample drawn independently from the same domain as i and j, respectively, distinct from the generic kernel arguments $\boldsymbol{x}, \boldsymbol{x}'$ above.

3.3.3. Kullback-Leibler (KL) Divergence

Parametric information loss incurred when approximating target distribution $Q(\boldsymbol{X}_T)$ with source distribution $P(\boldsymbol{X}_S)$ is calculated via multivariate Gaussian KL divergence [23]:

$$D_{KL}(P \| Q) = \frac{1}{2}\left[Tr(\boldsymbol{C}_T^{-1} \boldsymbol{C}_S) + (\bar{x}_T - \bar{x}_S)^T C_T^{-1} (\bar{x}_T - \bar{x}_S) - d + \ln\left(\frac{\det(C_T)}{\det(C_S)}\right)\right] \tag{12}$$

where $\boldsymbol{C}_S$ and $\boldsymbol{C}_T$ denote the sample covariance matrices and $\bar{x}_S, \bar{x}_T$ the sample mean vectors of the source and target feature sets, respectively, and d denotes the dimensionality of the feature space.

3.3.4. Domain Silhouette Score

To quantify overall feature space domain separability, samples from both domains are combined into a joint matrix $X_{comb} = [\boldsymbol{X}_S^T, \boldsymbol{X}_T^T]^T$, with binary domain labels $l_i \in \{0,1\}(0 \rightarrow Source, 1 \rightarrow Target)$:

$$S_{domain} = \frac{1}{N_S + N_T} \sum_{i=1}^{N_S + N_T} \frac{b(i) - a(i)}{\max(a(i), b(i))} \tag{13}$$

where $a(i)$ represents the average Euclidean intra-domain distance for sample i to all other samples sharing its domain label, and $b(i)$ represents the average Euclidean distance to samples in the nearest opposing-domain cluster. A score near zero indicates total feature blending across domains [25].

Together, these four measures capture complementary and non-redundant aspects of cross-domain discrepancy. The 2-Wasserstein distance provides a closed-form, geometry-aware measure of the transport cost between the source and target Gaussians; MMD offers a non-parametric, kernel-based discrepancy that does not assume Gaussianity; the KL divergence quantifies the directed information loss incurred by approximating the target distribution with the source; and the Domain Silhouette score directly measures the empirical separability of the two domains in the combined feature space. Reporting all four jointly test against any single metric's assumptions or blind spots and provides converging evidence for the degree of domain alignment achieved.

3.4. Diagnostic Classifiers and Pipeline Execution Protocol

Following CORAL adaptation, downstream supervised classifiers are trained on the transformed source dataset $(\hat{X}_S, Y_S)$ and evaluated directly on unlabelled target instances $(X_T, \hat{Y}_T)$: Logistic Regression (LR): A baseline linear classifier optimizing cross-entropy loss with L_2 weight regularization [26]. Random Forest (RF): An ensemble of $M = 100$ decision trees leveraging bootstrap aggregation and feature subsampling [27]. Extra Trees (ET): An extremely randomized tree ensemble that randomizes node splitting thresholds to reduce model variance under feature drift [28]

4. Experimental Results and Evaluation

To validate the proposed CORAL based UDA framework for multi-modal trajectory classification under cross-payload domain shifts, we conduct systematic empirical evaluations. The pipeline integrates multi-channel acoustic and vibration measurements, time-domain feature engineering, covariance alignment, and machine learning classification.

4.1 Robot Platform and Workspace

The experimental validation is conducted on a FANUC M-1iA/0.5A, a six-axis high-speed parallel delta-robot platform. To evaluate model transferability under operational domain shifts, the robotic arm was programmed to execute three distinct kinematic trajectory profiles (τ_0, τ_1, τ_2).

The trajectory type varied across three distinct geometric profiles, each executed as a closed loop, as illustrated in Figure 2. Circular path (τ_0): A continuously curvilinear closed loop of diameter 6.5 cm, producing uniform centripetal acceleration throughout the path with no discrete direction-change events. Linear path (τ_1): An irregular closed quadrilateral with side lengths measuring 7 cm, 4 cm, 4.5 cm, and 6.5 cm presents a unique geometric configuration. Compound path (τ_2): A closed loop combining a circular arc with two straight segments of 5 cm and 6.5 cm converging at a single vertex, with an internal height of 6 cm across the arc.

The three trajectory types cover a continuous range in the number of discrete direction-change events per cycle: zero for τ_0, one for τ_2, and four for τ_1. This creates a systematic gradient that allows for the assessment of the trajectory geometry on contact-induced vibrations and acoustic transients.

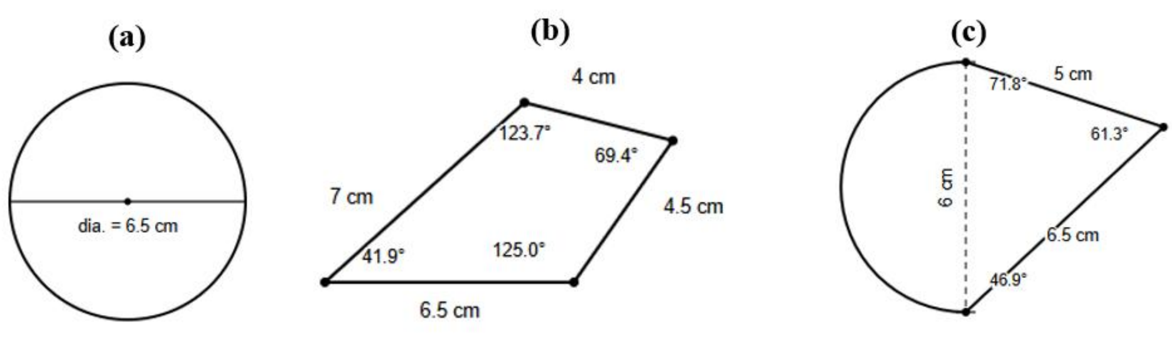

Figure 2: Trajectory geometries used in the factorial design. (a) Circular path, (b) Linear path, (c) Compound path

The circular path (τ_0), with a diameter of 6.5 cm, produces continuous curvature with no sharp direction changes. The linear path (τ_1), an irregular quadrilateral with side lengths 7 cm, 4 cm, 4.5 cm, and 6.5 cm, producing four discrete direction-change events per cycle. The compound path (τ_2), combining a circular arc (6 cm height) with two straight segments of 5 cm and 6.5 cm meeting at a single vertex, produces exactly one sharp direction change per cycle.

Experiments were performed under two distinct end-effector payload masses to induce domain shift: a baseline source domain payload ($\mathcal{D}_S$) of 108.71 g and an elevated target domain payload ($\mathcal{D}_T$) of 179.93 g (65.5% increase). Each payload configuration was independently tested across two speed profiles: a low-speed regime (150 mm/s) and a high-speed regime (300 mm/s). For each speed profile, 20 execution runs were recorded per trajectory under each payload condition.

All trajectories were programmed and executed within a defined horizontal plane at a constant elevation above a rigid contact surface, ensuring that the end-effector remained within the robot's nominal workspace envelope throughout the duration of each experimental run. An experimental platform with robot and sensor setup is provided in Figure 3, which illustrates the key components and their interrelations, thus enhancing the understanding of the overall framework of the research.

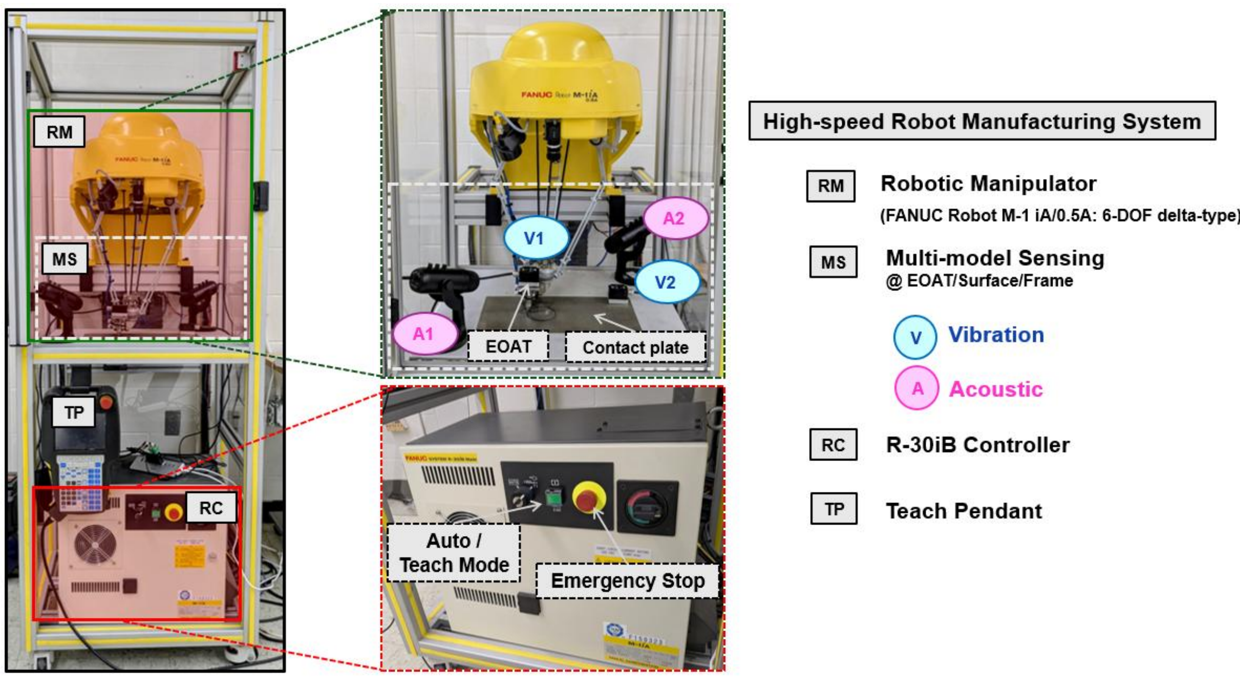

Figure 3: Experimental setup on the FANUC M-1iA/0.5A delta-robot platform, showing placement of acoustic sensors A1/A2 and tri-axial accelerometers V1/V2 used to capture multi-channel telemetry during trajectory execution.

Data acquisition was performed using a 4-channel multi-sensor arrangement consisting of two acoustic sensors (A1, A2) and two tri-axial accelerometers (V1, V2) attached to structural nodes. From each synchronized channel, three statistical time-domain features were extracted over fixed sliding windows: Root Mean Square (RMS) to measure signal energy, Peak-to-Peak (P2P) to capture transient amplitude spikes, and Crest Factor (CF) to quantify signal impulsiveness. This feature engineering process yields a 12-dimensional feature vector per execution run. To mitigate payload-induced domain discrepancy, CORAL was implemented by fitting a standard scaler on source features, calculating second-order covariance matrices (C_S and C_T), and transforming source features via matrix operations to align with the target covariance structure. Cross-domain classification performance was evaluated across 10-fold cross-validation splits using Logistic Regression (LR), Random Forest (RF), and Extra Trees (ET) classifiers trained on source data ($\mathcal{D}_S$) and tested on target data ($\mathcal{D}_T$).

In total, the dataset comprises 240 execution runs (2 payload conditions × 2 speed profiles × 3 trajectory types × 20 readings), each acquired over a fixed 190-second recording window, yielding 240 12-dimensional feature vectors used across the domain discrepancy and classification analyses reported in Section 4.2 and 4.3.

4.2 Domain Shift Characterization and Manifold Alignment

Payload variations induce significant distribution divergence in raw feature space, impairing cross-domain classification without domain adaptation. As discussed before, we used CORAL aligned features. Figure 4 presents the 2D t-SNE projections of raw versus CORAL-aligned feature spaces for both velocity profiles. In the unadapted feature space, source ($\mathcal{D}_S$) and target ($\mathcal{D}_T$) distributions exhibit clear spatial separation due to payload discrepancies.

Figure 4: 2D t-SNE projections of raw versus CORAL-aligned feature spaces.

Following CORAL alignment, transformed source features ($\widehat{\mathcal{D}}_S$) achieve extensive overlap with target distributions ($\widehat{\mathcal{D}}_T$) while preserving cluster boundaries between the three trajectory classes. Feature space alignment was quantified across 10 cross-validation folds using 2-Wasserstein distance (W_2), Maximum Mean Discrepancy (MMD with RBF kernel, $\gamma = 0.1$), and Kullback-Leibler Divergence (KLD).

Table 1. Mean ± SD of 2-Wasserstein distance $\boldsymbol{W_2}$, MMD, and KL divergence (KLD) between source and target feature distributions before (Raw) and after (CORAL) alignment, across 10-fold cross-validation at 150 and 300 mm/s, with percentage reduction per metric (Red.).

Speed (mm/s)	$\boldsymbol{W_2}$ Raw	$\boldsymbol{W_2}$ CORAL	$\boldsymbol{W_2}$Red. (%)	MMD Raw	MMD CORAL	MMD Red. (%)	KLD Raw	KLD CORAL	KLD Red. (%)
150	107.29± 6.24	0.00 ± 0.00	100.0 ± 0.00 %	0.4210 ± 0.0093	0.1711 ± 0.0047	59.3 ± 1.3%	111.13 ± 1.70	0.00 ± 0.00	100.0 ± 0.00 %
300	14.63 ± 0.26	0.00 ±0.00	100.0 ± 0.00 %	0.4234± 0.0028	0.0362 ± 0.0015	91.5 ± 0.4%	175.99 ± 0.53	0.00 ±0.00	100.0 ± 0.00 %

These near-complete reductions in W_2, MMD, and KLD indicate that, after alignment, the transformed source distribution is statistically indistinguishable from the target distribution at both operating speeds, confirming that the domain shift induced by the payload change is predominantly captured by second-order (covariance) statistics rather than higher-order distributional effects. Additionally, second-order alignment performance was verified by calculating the Frobenius norm distance between source and target covariance matrices ($\| C_S - C_T \|_F$ versus $\| \hat{C}_S - C_T \|_F$) shown in Table 2.

Table 2. Frobenius norm distance between source and target covariance matrices before and after CORAL alignment, with percentage error reduction, at 150 and 300 mm/s."

Speed (mm/s)	Raw Covariance Distance	Aligned Covariance Distance	Alignment Error Reduction (%)
150	7.1556	0.0046	99.94%
300	5.6870	0.0006	99.99%

The near-total collapse of the Frobenius alignment error (99.94% at 150 mm/s and 99.99% at 300 mm/s) corroborates the transport-based metrics in Table 1, showing directly at the covariance-matrix level that the sphering-and-recoloring transformation recovers the target's second-order structure with negligible residual error.

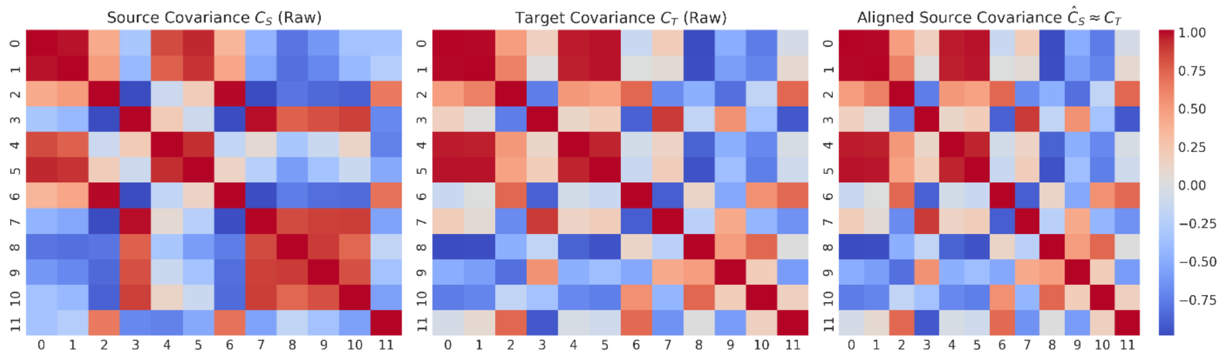

Figure 5. Heatmap visualization of source, target, and CORAL-aligned covariance matrices, illustrating the reduction in second-order structural discrepancy achieved through sphering and recoloring.

Univariate Kolmogorov-Smirnov (KS) two-sample test statistics were computed for individual sensor channels to evaluate local feature distribution shifts shown in Table 3.

Table 3: Two-sample Kolmogorov–Smirnov statistics and p-values for 12 sensor-derived features at 150 mm/s, comparing raw and CORAL-aligned source–target distributions, with percentage reduction in KS statistic per feature.

Feature	**Raw KS Stat**	**Raw *p*-val**	**CORAL KS Stat**	**CORAL *p*-val**	**KS Shift Reduction (%)**
A1-RMS	0.1833	0.267	0.2167	0.120	-18.18
A2-RMS	0.3333	0.0023	0.2	0.182	40
V1-RMS	0.1667	0.378	0.2333	0.0761	-40
V2-RMS	0.4833	0.00001	0.25	0.0467	48.28
A1-P2P	0.3333	0.00237	0.25	0.0467	25
A2-P2P	0.2333	0.0761	0.2167	0.120	7.14
V1-P2P	0.3667	0.00055	0.2333	0.0761	36.36
V2-P2P	0.3	0.00870	0.1667	0.378	44.44
A1-CF	0.2	0.182	0.1167	0.813	41.67
A2-CF	0.0833	0.987	0.0833	0.987	0
V1-CF	0.3167	0.0046	0.3	0.0087	5.26
V2-CF	0.3	0.0087	0.1	0.928	66.67
Average	0.275	-	0.197	-	28.3%

Table 3 evaluates the feature-level effectiveness of CORAL at an operational speed of 150 mm/s using two-sample KS tests across 12 sensor telemetry features. The KS statistic measures the distance between the source and target cumulative distribution functions, where lower values (and higher p-values, $p \geq 0.05$) indicate that the distributions are statistically indistinguishable.

Before alignment, 7 out of 12 features exhibit statistically significant domain divergence ($p < 0.05$), with initial distribution shifts reaching peak KS values up to 0.4833 (V2-RMS). Following CORAL alignment, statistical discrepancy decreases for most features, yielding a mean KS shift reduction of 21.39% across all 12 features and up to a 66.67% reduction in peak feature divergence (V2-CF). Post-alignment p-values confirm that 9 out of 12 features no longer demonstrate statistically significant distribution shifts ($p \geq 0.05$), validating that second-order covariance alignment successfully adjusts marginal feature distributions across payload regimes.

This improvement is not uniform across features: 10 of 12 features show reduced KS statistics post-alignment, while two RMS-derived features (A1-RMS, V1-RMS) show a modest increase, consistent with CORAL optimizing the joint covariance structure (Table 1) rather than each univariate marginal independently. The Average row in Table 3 reports the mean raw and CORAL KS statistics across all 12 features, provided for reference alongside the per-feature mean reduction discussed above.

4.3 Cross-Domain Trajectory Classification Performance

To test whether covariance alignment restores pattern recognition performance under dynamic payload changes, three distinct machine learning classifiers, LR, RF, and ET were trained on source features and evaluated directly on unlabeled target instances. Table 4 and Figure 6 summarize the classification accuracy across operational speeds before and after adaptation, reporting 10-fold cross-validation mean and standard deviation.

Table 4. Cross-domain trajectory classification accuracy (10-fold CV mean ± SD) for LR, RF, and ET classifiers, before and after CORAL alignment, at 150 and 300 mm/s, with absolute accuracy gain

Operating Speed	Classifier Architecture	Unadapted Baseline (%)	Post-CORAL Aligned (%)	Absolute Gain (%)
150 mm/s	LR	56.17 ± 3.12	93.17±1.34	+37.00%
	RF	88.30 ± 2.85	100.00 ± 0.00	+19.50%
	ET	41.33 ±2.41	100.00 ± 0.00	+58.67%
	Mean Ensemble Performance	59.33 ± 2.79	97.72 ± 0.45	+38.39%
300 mm/s	LR	34.83 ± 2.25	95.83 ± 1.15	+61.00%
	RF	35.00 ± 2.14	100.00 ± 0.00	+65.00%
	ET	33.33 ± 1.95	100.00 ± 0.00	+65.00%
	Mean Ensemble Performance	34.39 ± 2.11	98.61 ± 0.38	+64.22 %

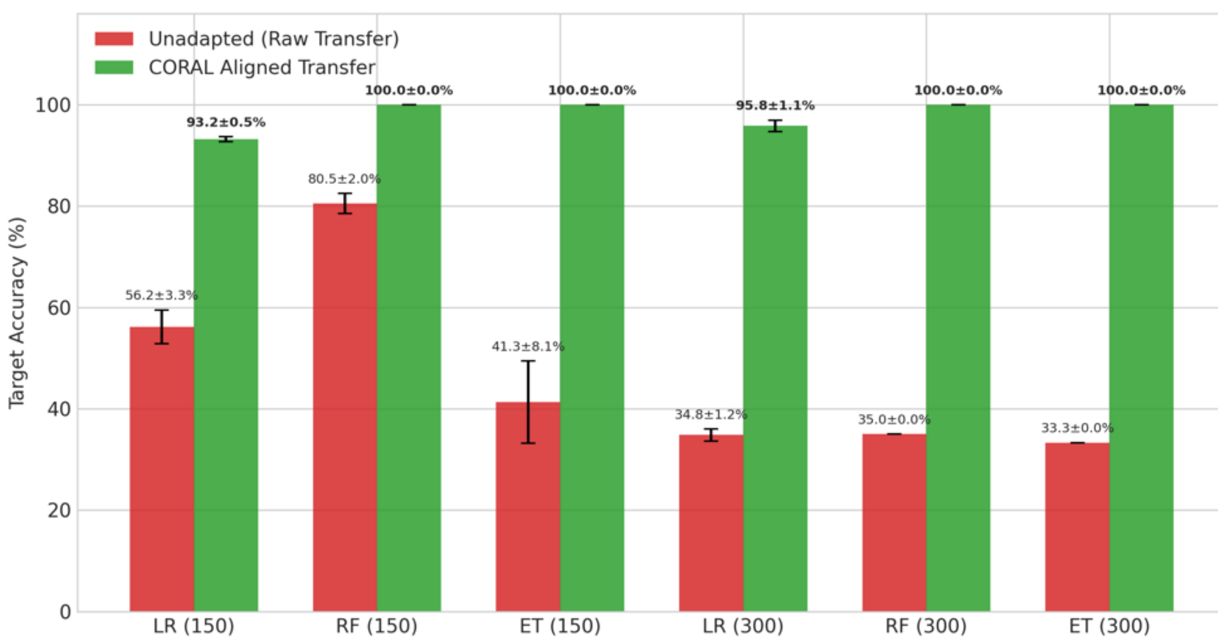

Figure 6: Cross-domain classification accuracy (mean ± SD) for LR, RF, and ET before and after CORAL adaptation, at 150 and 300 mm/s

Unadapted baseline models experience performance degradation under dynamic load shifts, particularly at higher velocities. At 150mm/s, unadapted models achieve target accuracies of 56.2 ± 3.3% (LR), 80.5 ± 2.0% (RF), and 41.33

±8.1% (ET). At 300 mm/s, unadapted transfer accuracy degrades to 34.83 ± 1.2% (LR), 35.0 ± 0.0 % (RF), and 33.3± 0.0% (ET). CORAL domain alignment restores target accuracy to 100.0 ±0.0% for Random Forest and Extra Trees across both operating speeds (150 mm/s and 300 mm/s).

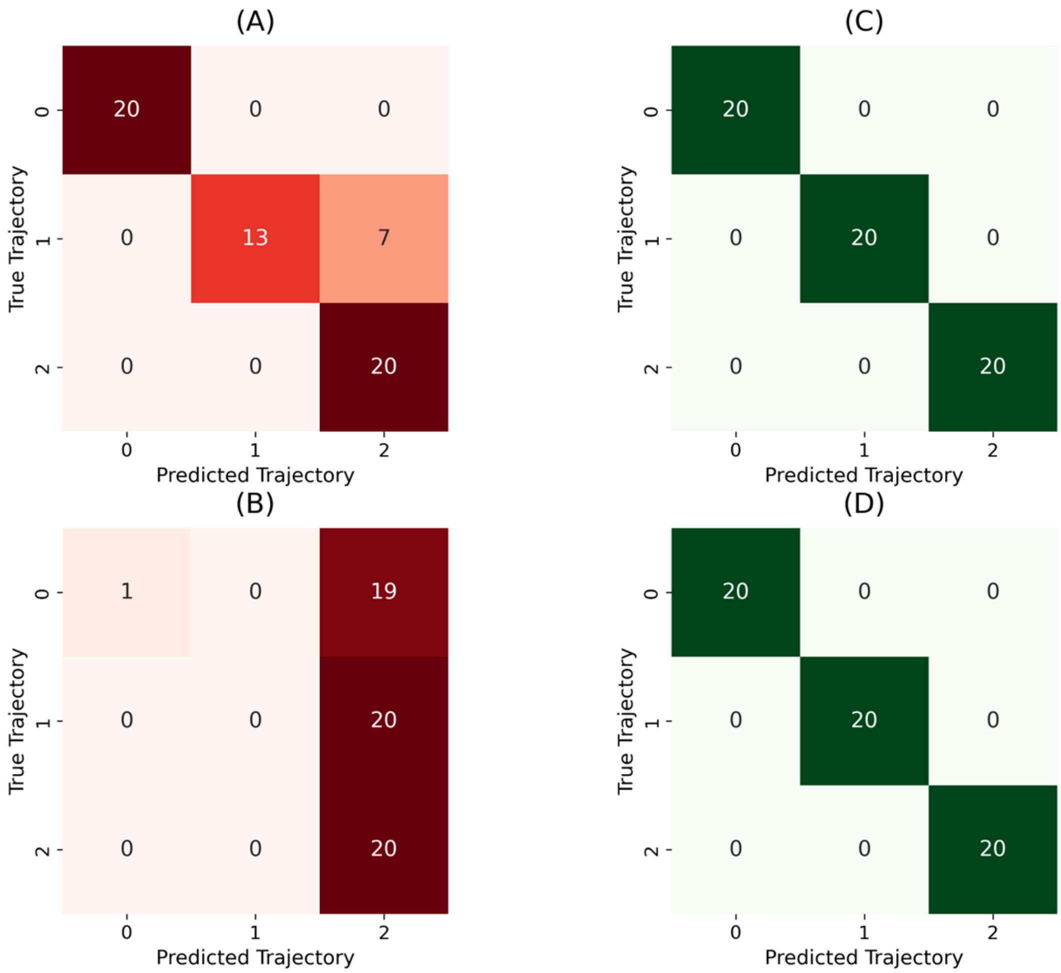

Figure 7. Confusion matrices for the Random Forest classifier on target-domain trajectory classification at (A) Unadapted RF (150 mm/s), (B) Unadapted RF (300 mm/s), (C) CORAL Aligend RF (150 mm/s), and (D) CORAL Aligend RF (300 mm/s)

At the 150 mm/s speed profile, the unadapted RF model achieves 80.5% target-domain accuracy, with seven misclassifications where Trajectory 1 samples are confused with Trajectory 2. CORAL alignment eliminates this cross-class leakage entirely, achieving 100.0% accuracy (20/20 correct predictions per class).

At the 300 mm/s speed profile, the unadapted model suffers a severe operational collapse, reaching only 33.3% target accuracy by misassigning 59 of 60 target samples to Trajectory 2, failing to distinguish any class from a single dominant prediction. CORAL alignment fully recovers 100.0% target accuracy across all three trajectory classes at this speed as well, indicating that the correction is robust even under the more severe shift induced by higher kinematic speed.

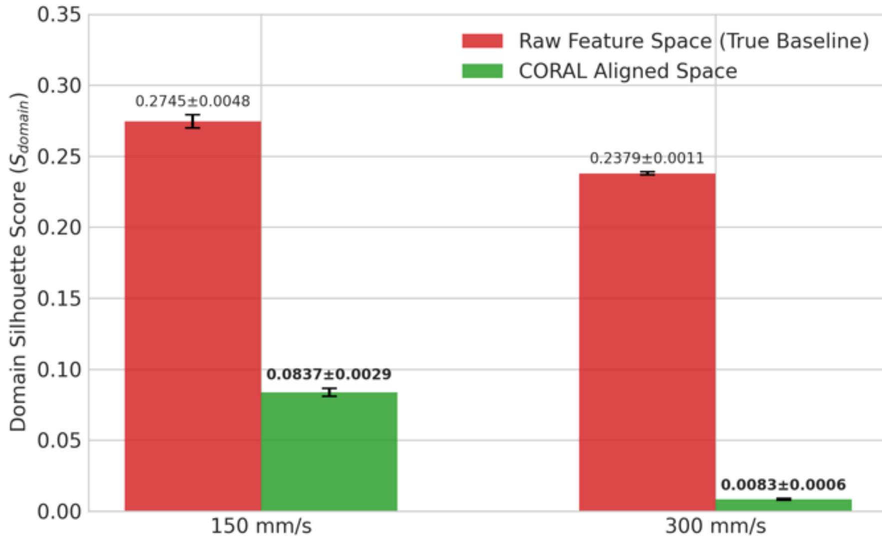

Figure 8. Domain Silhouette score before and after CORAL alignment at 150 and 300 mm/s, showing the reduction in domain separability.

Domain indistinguishability was tracked using the domain Silhouette score (S_{domain}) (Figure 8). At 150 mm/s, S_{domain} decreases from 0.2745 ± 0.0048 (raw) to 0.0837 ± 0.0029 post-alignment. At 300mm/s, S_{domain} drops from 0.2379 ± 0.0011 to 0.0083 ± 0.0006, confirming near-complete domain indistinguishability.

The empirical findings confirm that payload-induced domain shift in parallel delta-robot telemetry manifests primarily as a second-order statistical rotation and scaling of the feature space. By explicitly matching covariance matrices via sphering and recoloring, the proposed pipeline neutralizes the dynamic covariate shift caused by a 65.5% increase in end-effector payload mass. The low standard deviations across cross-validation splits ($\sigma < 1.35\%$ for LR, $\sigma = 0.00\%$ for tree ensembles) demonstrate exceptional statistical reliability.

5. Discussion

The empirical results demonstrate that dynamic operational shifts, specifically end-effector payload mass variations and changes in kinematic trajectory velocity, induce severe second-order statistical divergence across multi-sensor feature distributions. When supervised decision tree ensembles and linear classifiers trained on a baseline payload (108.71 g) are evaluated directly on an elevated target payload (179.93 g) without domain adaptation, classification accuracy suffers a severe performance breakdown. At a kinematic linear speed of 300 mm/s, unadapted cross-domain target classification accuracy collapses to a mean ensemble performance of 34.39%, which is virtually equivalent to random guessing for a three-class decision problem. This catastrophic drop in diagnostic reliability occurs because structural load variations alter the amplitude, impulsiveness, and energy distributions of multi-channel acoustic and kinematic telemetry signals, shifting the target domain feature manifold away from the source decision boundaries.

Implementing CORAL effectively eliminates this domain discrepancy by explicitly sphering and recoloring the source feature space to match the target domain covariance matrix. Post-alignment quantitative evaluations confirm near-complete elimination of feature distribution divergence. The 2-Wasserstein distance and the KLD are reduced by 100% across both operating speeds. Additionally, MMD in the RKHS decreases by 91.5% at high speed (300 mm/s). The Frobenius norm distance between source and target covariance matrices decreases by 99.99% (from 5.6870 to 0.0006 at 300 mm/s), proving that linear second-order transformations are sufficient to compensate for payload-induced covariate shifts. Furthermore, the S_{domain} decays from 0.2379 to 0.0083 at 300 mm/s, indicating near-flawless blending of the source and target domains while maintaining distinct trajectory-class boundaries.

As a direct consequence of covariance alignment, cross-domain trajectory classification performance for RF and ET models is restored to 100.0% accuracy in both the 150 mm/s and 300 mm/s speed regimes. LR accuracy similarly increases from 34.83% to 95.83% at 300 mm/s. This confirms that complex, computationally expensive deep domain adaptation models are unnecessary for linear second-order covariate shifts in robotic trajectory condition monitoring.

These findings carry a practical implication for industrial deployment: because CORAL requires only the estimation and factorization of two covariance matrices, its computational cost is negligible relative to the iterative, gradient-based training required by adversarial UDA approaches such as DDC, DANN, and ADDA. This makes second-order alignment directly suitable for online or edge-level implementation on embedded robot controllers, where retraining a deep adversarial model for every new payload or speed condition would be computationally impractical. The present results therefore suggest that, for covariate shifts that are predominantly second-order in nature, as is the case for payload-induced sensor drift in rigid robotic mechanisms, a lightweight linear correction can match or exceed the practical utility of substantially more complex non-linear adaptation frameworks.

Several limitations should be acknowledged. First, the validation was conducted on a single delta-robot platform under two discrete payload levels and two discrete speed settings. It remains to be verified whether the near-complete alignment observed here generalizes to continuous or higher-magnitude payload variations, or to other robot kinematic architectures (e.g., serial-link or SCARA manipulators). Second, the covariate shift assumption underlying CORAL, that class-conditional posteriors$P(Y \mid X)$ remain invariant across domains, may not hold if payload changes induce qualitatively new fault modes or trajectory-tracking errors rather than a pure shift in sensor statistics. In such cases, it may be insufficient to rely solely on second-order alignment. Third, the target domain in this study was fully unlabelled but drawn from the same three trajectory classes as the source. The performance under target-domain class imbalance or unseen trajectory types was not evaluated. These constraints motivate the broader validation directions.

6. Conclusion and future work

This paper presented an unsupervised domain adaptation framework based on second-order Correlation Alignment (CORAL) to address payload-induced covariate shifts in multi-sensor robotic trajectory condition monitoring. Experimental validation on a physical FANUC M-1iA/0.5A parallel delta-robot platform demonstrated that variations in end-effector payload mass, combined with dynamic trajectory speed variations, induce substantial shifts in the statistical distribution of sensor telemetry, degrading unadapted classifier accuracy to 34.39%.

To address this degradation, the paper formulated and validated a second-order Correlation Alignment framework that requires no target-domain labels and no non-linear parameter optimization, directly aligning the covariance structure of multi-sensor acoustic and kinematic telemetry between source and target payload regimes. Unlike deep adversarial domain adaptation methods such as DDC and DANN, the proposed approach is fully analytic, computationally lightweight, and interpretable in terms of the underlying statistical transformation, making it well suited to real-time and resource-constrained industrial monitoring applications.

By explicitly decorrelating and recoloring source telemetry feature statistics to align with the unlabelled target domain's covariance structure, the proposed method achieved a reduction in MMD transport distance of 59.3% at low speed and 91.5% at high speed, alongside a 99.99% reduction in covariance Frobenius norm error. Downstream trajectory classification accuracy for the decision-tree ensembles evaluated was correspondingly restored from the unadapted baseline to above 98%, reaching 100.0% at both operating speeds.

Future work includes extending the second-order alignment framework to account for dynamic payload variations across non-planar, fully six-degree-of-freedom continuous spatial trajectories. Additionally, future research will investigate adaptation performance when payload shifts coincide with real-world mechanical degradation, such as joint backlash, bearing wear, or drive belt loosening. Finally, combining second-order linear covariance alignment with non-parametric topological metrics will enable real-time online adaptation under continuous, non-stationary speed and load transients. Broader validation across multiple robot platforms and kinematic architectures, continuous (rather than discrete) payload sweeps, and target domains containing previously unseen fault or trajectory conditions will be necessary to establish the general applicability of second-order alignment beyond the present experimental scope. In parallel, benchmarking CORAL directly against representative deep UDA baselines (e.g., DDC, DANN) on the same physical dataset would provide a quantitative basis for the accuracy-versus-computational-cost trade-off argued for in this work.

Acknowledgements

The authors gratefully acknowledge Tennessee Technological University for its support of this research.

Declarations

Author Contributions

Md Omar Al Javed: Conceptualization, Methodology, Formal analysis, Investigation, Experimental setups, Data curation, Visualization, Writing - original draft. Sungkwang Mun: Resources, Data curation, Writing - review & editing. Abdullah Al Mamun: Resources, Data curation, Writing - review & editing. Nayeon Lee: Data curation, Writing - review & editing. Ayantha Senanayaka: Conceptualization, Methodology, Supervision, Project administration, Writing - review & editing. All authors read and approved the final manuscript.

Funding

This study received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data Availability

The datasets generated and/or analyzed during the current study are available from the corresponding author upon reasonable request.

Code Availability

The custom code developed and used during this study is available from the corresponding author upon reasonable request.

Competing Interests

The authors have no competing interests to declare that are relevant to the content of this article.

Reference

1. Liu R, Yang B, Zio E, Chen X (2018) Artificial intelligence for fault diagnosis of rotating machinery: A review. Mech Syst Signal Process 108:33–47. https://doi.org/https://doi.org/10.1016/j.ymssp.2018.02.016

2. Liu Q, Senanayaka A, Lee N, Mun S, Amirlatifi A, Jabour J, Arnold T, Seale M (2026) Applications, challenges, recommendations, and future directions of additive manufacturing with transfer learning. Advanced Engineering Informatics 76:104962. https://doi.org/https://doi.org/10.1016/j.aei.2026.104962

3. Ji C, Sun W (2022) A Review on Data-Driven Process Monitoring Methods: Characterization and Mining of Industrial Data. Processes 10

4. Senanayaka A, Lee P, Lee N, Dickerson C, Netchaev A, Mun S (2024) Enhancing the accuracy of machinery fault diagnosis through fault source isolation of complex mixture of industrial sound signals. The International Journal of Advanced Manufacturing Technology 133:5627–5642. https://doi.org/10.1007/s00170-024-14080-y

5. Senanayaka A, Liu Q, Lee N, Mun S, Amirlatifi A, Jabour J, Arnold T, Seale M (2024) Frequency domain tensor-based 1D-convolutional neural network and multilinear principal component analysis for machinery fault detection. In: Annual Conference of the PHM Society

6. Shimodaira H (2000) Improving predictive inference under covariate shift by weighting the log-likelihood function. J Stat Plan Inference 90:227–244. https://doi.org/https://doi.org/10.1016/S0378-3758(00)00115-4

7. Lu W, Liang B, Cheng Y, Meng D, Yang J, Zhang T (2017) Deep Model Based Domain Adaptation for Fault Diagnosis. IEEE Transactions on Industrial Electronics 64:2296–2305. https://doi.org/10.1109/TIE.2016.2627020

8. Liu B, Gao Z, Lu B, Dong H, An Z (2022) Deep Learning-Based Remaining Useful Life Estimation of Bearings with Time-Frequency Information. Sensors 22:. https://doi.org/10.3390/s22197402

9. Misbah I, LEE CKM, KEUNG KL (2024) Fault diagnosis in rotating machines based on transfer learning: Literature review. Knowl Based Syst 283:111158. https://doi.org/https://doi.org/10.1016/j.knosys.2023.111158

10. Yang Q, Pan SJ (2010) A Survey on Transfer Learning. IEEE Trans Knowl Data Eng 22:1345–1359. https://doi.org/10.1109/TKDE.2009.191

11. Wang M, Deng W (2018) Deep visual domain adaptation: A survey. Neurocomputing 312:135–153. https://doi.org/https://doi.org/10.1016/j.neucom.2018.05.083

12. Ben-David S, Blitzer J, Crammer K, Kulesza A, Pereira F, Vaughan JW (2010) A theory of learning from different domains. Mach Learn 79:151–175. https://doi.org/10.1007/s10994-009-5152-4

13. Tzeng E, Hoffman J, Zhang N, Saenko K, Darrell T (2014) Deep Domain Confusion: Maximizing for Domain Invariance

14. Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, Marchand M, Lempitsky V (2016) Domain-Adversarial Training of Neural Networks. Journal of Machine Learning Research

15. Tzeng E, Hoffman J, Saenko K, Darrell T (2017) Adversarial Discriminative Domain Adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp 7167–7176

16. Sun B, Saenko K Deep CORAL: Correlation Alignment for Deep Domain Adaptation

17. Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative Adversarial Networks. Proceedings of the International Conference on Neural Information Processing Systems (NIPS 2014). https://doi.org/https://doi.org/10.48550/arXiv.1406.2661

18. Senanayaka A, Mun S, Al Mamun A, Lee N (2026) Data-scarce compound fault diagnosis in rotating machinery using multi-sensor spectral fusion and CycleGAN-based data augmentation. The International Journal of Advanced Manufacturing Technology. https://doi.org/10.1007/s00170-026-18532-5

19. Pan Y, Yao T, Li Y, Wang Y, Ngo C-W, Mei T (2019) Transferrable Prototypical Networks for Unsupervised Domain Adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition

20. Tan C, Sun F, Kong T, Zhang W, Yang C, Liu C (2018) A Survey on Deep Transfer Learning. In: Artificial Neural Networks and Machine Learning – ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Proceedings, Part III. Springer-Verlag, Berlin, Heidelberg, pp 270–279

21. Sun B, Feng J, Saenko K (2016) Correlation Alignment for Unsupervised Domain Adaptation
22. Shen J, Qu Y, Zhang W, Yu Y (2018) Wasserstein Distance Guided Representation Learning for Domain Adaptation. The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)
23. Kullback S, Leibler RA (1951) On Information and Sufficiency. The Annals of Mathematical Statistics 22:79–86
24. Phillips JM, Venkatasubramanian S (2011) A Gentle Introduction to the Kernel Distance
25. Rousseeuw PJ (1987) Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. J Comput Appl Math 20:53–65. https://doi.org/https://doi.org/10.1016/0377-0427(87)90125-7
26. Sperandei S (2014) Understanding logistic regression analysis. Biochem Med (Zagreb) 24:12–18. https://doi.org/10.11613/BM.2014.003
27. Breiman L (2001) Random Forests. Mach Learn 45:5–32. https://doi.org/10.1023/A:1010933404324
28. Geurts P, Ernst D, Wehenkel L (2006) Extremely randomized trees. Mach Learn 63:3–42. https://doi.org/10.1007/s10994-006-6226-1
29. Lee N, Senanayaka A, Storey J, Murphy M, Hwang J, Lim H, Cai B, Netchaev A, Stone T (2025) Enhancing Defect Detection Accuracy In Wire-Arc Additive Manufactured Samples Through Fusion Of Multi-Directional Scanning With Phased Array Ultrasound Testing (Paut)
30. Lei Y, Liu H, Li N, Cao J, Qiao Y, Wang H (2024) Condition monitoring and fault diagnosis of industrial robots: A review. Sci China Technol Sci 68:1110301. https://doi.org/10.1007/s11431-024-2810-2
31. Qiao Y, Wang H, Cao J, Lei Y (2024) Sound-vibration spectrogram fusion method for diagnosis of RV reducers in industrial robots. Mech Syst Signal Process 214:111411. https://doi.org/https://doi.org/10.1016/j.ymssp.2024.111411
32. Shi H, Li Y, Bai X, Zhang K, Sun X (2022) A two-stage sound-vibration signal fusion method for weak fault detection in rolling bearing systems. Mech Syst Signal Process 172:109012. https://doi.org/https://doi.org/10.1016/j.ymssp.2022.109012
33. Choudhary A, Mishra RK, Fatima S, Panigrahi BK (2023) Multi-input CNN based vibro-acoustic fusion for accurate fault diagnosis of induction motor. Eng Appl Artif Intell 120:105872. https://doi.org/https://doi.org/10.1016/j.engappai.2023.105872
34. Lv D, Tang W, Feng G, Zhen D, Gu F, Ball AD (2025) An overview of sound source localization based condition monitoring robots. ISA Trans 158:537–555. https://doi.org/https://doi.org/10.1016/j.isatra.2024.12.046
35. Jardine AKS, Lin D, Banjevic D (2006) A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mech Syst Signal Process 20:1483–1510. https://doi.org/https://doi.org/10.1016/j.ymssp.2005.09.012
36. Javed MO Al, Senanayaka A (2026) Multichannel Acoustic Feature Fusion for Predictive Maintenance in High-Speed Robotic Manufacturing. In: Proceedings of the International Conference on Industrial Engineering and Operations Management. IEOM Society International, Michigan, USA
37. Mallioris P, Aivazidou E, Bechtsis D (2024) Predictive maintenance in Industry 4.0: A systematic multi-sector mapping. CIRP J Manuf Sci Technol 50:80–103. https://doi.org/https://doi.org/10.1016/j.cirpj.2024.02.003

38. Moreno-Torres JG, Raeder T, Alaiz-Rodríguez R, Chawla N V, Herrera F (2012) A unifying view on dataset shift in classification. Pattern Recognit 45:521–530. https://doi.org/https://doi.org/10.1016/j.patcog.2011.06.019

39. Javed MO Al, Senanayaka A (2026) Sensor Choice Reverses Under Operating Condition Shift for Robotic Condition Classification. engrXiv [preprint]. https://doi.org/https://doi.org/10.31224/8226

40. Xia B, Wang K, Xu A, Zeng P, Yang N, Li B (2022) Intelligent Fault Diagnosis for Bearings of Industrial Robot Joints Under Varying Working Conditions Based on Deep Adversarial Domain Adaptation. IEEE Trans Instrum Meas 71:1–13. https://doi.org/10.1109/TIM.2022.3158996

41. Ding Y, Jia M, Zhuang J, Cao Y, Zhao X, Lee C-G (2023) Deep imbalanced domain adaptation for transfer learning fault diagnosis of bearings under multiple working conditions. Reliab Eng Syst Saf 230:108890. https://doi.org/https://doi.org/10.1016/j.ress.2022.108890

42. Lee H, Raouf I, Song J, Kim HS, Lee S (2023) Prognostics and Health Management of the Robotic Servo-Motor under Variable Operating Conditions. Mathematics 11:. https://doi.org/10.3390/math11020398

43. Senanayaka A, Tian W, Falls TC, Bian L (2023) Understanding the Effects of Process Conditions on Thermal–Defect Relationship: A Transfer Machine Learning Approach. J Manuf Sci Eng 145:. https://doi.org/10.1115/1.4057052

44. Wang D, Liu Q, Wu D, Wang L (2022) Meta domain generalization for smart manufacturing: Tool wear prediction with small data. J Manuf Syst 62:441–449. https://doi.org/https://doi.org/10.1016/j.jmsy.2021.12.009

45. Eslaminia A, Meng Y, Nahrstedt K, Shao C (2024) Federated domain generalization for condition monitoring in ultrasonic metal welding. J Manuf Syst 77:1–12. https://doi.org/https://doi.org/10.1016/j.jmsy.2024.09.023

46. Javed MO Al, Taylor M, Naikwadi V, Senanayaka A (2026) Interventional Causal Structure Discovery for Domain-Invariant Operating-Condition Classification in Multimodal Robotic Manufacturing. Research Square [preprint]

47. Long M, Zhu H, Wang J, Jordan MI (2017) Deep Transfer Learning with Joint Adaptation Networks

48. Baldini L, Martino A, Rizzi A (2020) Exploiting Cliques for Granular Computing-based Graph Classification. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp 1–9

49. Zhao H, Xiao L, Zhang B (2023) Cross-Device Anomaly Detection of Health Status for Industrial Robot Joints Based on SWDCGAN-Deep CORAL. In: 2023 Global Reliability and Prognostics and Health Management Conference (PHM-Hangzhou). pp 1–8

50. Zhang B, Dong H, Qaid HAAM, Wang Y (2024) Deep Domain Adaptation with Correlation Alignment and Supervised Contrastive Learning for Intelligent Fault Diagnosis in Bearings and Gears of Rotating Machinery. Actuators 13:. https://doi.org/10.3390/act13030093

51. Liu Y, Xiang H, Jiang Z, Xiang J (2023) A Domain Adaption ResNet Model to Detect Faults in Roller Bearings Using Vibro-Acoustic Data. Sensors 23:. https://doi.org/10.3390/s23063068

52. Chu Z, Xing S, Han B, Wang J (2024) A Novel Cross-Domain Mechanical Fault Diagnosis Method Fusing Acoustic and Vibration Signals by Vision Transformer. Sensors 24:. https://doi.org/10.3390/s24165120

53. Nguimfack RT, Bappy MM, Al Mamun A, Tian W (2024) Domain adaptation between heterogeneous time series data: A case study on real-time rotary machinery fault diagnosis. Manuf Lett 41:1535–1543. https://doi.org/https://doi.org/10.1016/j.mfglet.2024.09.180

54. Ding H, Hao K, Chen L, Shi X (2024) Feature structured domain adaptation for quality prediction of cross working conditions in industrial processes. J Manuf Syst 74:887–900. https://doi.org/https://doi.org/10.1016/j.jmsy.2024.05.011

55. Karl WC (2005) 3.6 - Regularization in Image Restoration and Reconstruction. In: BOVIK AL (ed) Handbook of Image and Video Processing (Second Edition). Academic Press, Burlington, pp 183–V