DiFG-Net++: Physics-Informed Functional Learning for Cross-Speed Fault Diagnosis

David He[1], Miao He[2], and Muhammad Areeb [1]

[1,2]*Department of Mechanical and Industrial Engineering, University of Illinois-Chicago, Chicago, IL, 60607, USA*
davidhe@uic.edu

[2]*Siemens Corporation, Princeton, NJ, 08540, USA*
miao.he@siemens.com

ABSTRACT

Cross-speed fault diagnosis remains a significant challenge in data-driven prognostics and health management (PHM) because vibration signatures and fault characteristic frequencies vary systematically with rotational speed. The problem becomes particularly difficult in open-set scenarios where the target operating speed lies outside the range observed during training and fault data at the target speed are unavailable. Although recent advances in self-supervised learning (SSL), particularly Prediction of Functionals from Masked Latents (PFML), have demonstrated promising capabilities for learning fault-sensitive representations from unlabeled data, existing PFML-based frameworks do not explicitly address speed-induced distribution shifts. In particular, the recently developed Domain-Informed Functional Guidance Network (DiFG-Net) relies primarily on handcrafted functionals and lacks mechanisms for enforcing speed-invariant representation learning.

To address these limitations, this paper proposes DiFG-Net++, a physics-informed functional learning framework specifically designed for open-set cross-speed fault diagnosis. Building upon the DiFG-Net architecture, DiFG-Net++ integrates three complementary speed-adaptation strategies: (1) speed-extrapolative augmentation, which generates physically meaningful intermediate and extrapolated speed conditions through speed-guided time-scale transformations; (2) speed-invariant latent regularization, which encourages latent representations associated with the same machine health condition to remain consistent across operating speeds; and (3) healthy-state target-speed anchoring, which leverages readily available healthy vibration data at the target operating speed to facilitate adaptation without requiring target-speed fault samples. In addition, order-domain functionals are incorporated as physics-informed self-supervised targets to promote learning of speed-robust fault characteristics.

The effectiveness of the proposed framework is evaluated using a plastic-bearing fault dataset collected under four rotational speeds and multiple bearing health conditions. Experimental results demonstrate that the proposed speed-adaptation mechanisms substantially improve open-set cross-speed fault diagnosis performance. Compared with the baseline DiFG-Net, which achieved an overall validation accuracy of 46.68%, the proposed DiFG-Net++ achieved an overall accuracy of 95.85%, representing a relative improvement of more than 105%. These results demonstrate that integrating physics-informed functional learning with speed-aware representation learning provides an effective and practical solution for robust cross-speed fault diagnosis under previously unseen operating conditions.

1. INTRODUCTION

Over the past decade, data-driven fault diagnosis methods based on machine learning and deep learning have demonstrated remarkable success in automatically extracting discriminative features from monitoring signals for equipment fault diagnosis and prognosis. However, the performance of these methods often deteriorates significantly when the operating conditions during deployment differ from those observed during training. Transfer learning and domain adaptation techniques have therefore received considerable attention in intelligent fault diagnosis to address distribution shifts caused by variations in operating conditions, sensor locations, and machine configurations (Pan &Yang, 2010; Zhang, Li, Ma, Luo, & Li, 2019). Among various domain shifts encountered in industrial applications, variations in rotational speed remain one of the most challenging factors because monitoring signal signatures and fault characteristic frequencies are inherently dependent on shaft speed. Consequently, models trained at a single speed often fail to generalize to unseen operating speeds, particularly in open-set scenarios where the target speed falls outside the training range.

The challenge of cross-speed fault diagnosis stems from the difficulty of learning speed-invariant yet fault-sensitive representations from vibration data. Classical signal processing approaches, such as order tracking and speed normalization techniques, attempt to mitigate speed-induced variations by transforming vibration signals into the order domain, where characteristic frequencies are represented as ratios relative to shaft rotational speed. Although effective under certain conditions, these approaches often require accurate speed measurements and extensive domain expertise. More recently, domain adaptation and transfer learning methods have been developed to align feature distributions across different operating conditions using techniques such as Maximum Mean Discrepancy (MMD), adversarial domain adaptation, and discrepancy minimization strategies (Pan &Yang, 2010; Zhang, et al, 2019). Despite their success, many existing approaches assume that source and target domains exhibit overlapping operating conditions or rely on labeled target-domain data, assumptions that are often unrealistic in industrial settings. As a result, cross-speed fault diagnosis remains an open challenge, especially when fault data at the target speed are unavailable, and the model must extrapolate beyond the speed range observed during training.

Recent advances in self-supervised learning (SSL) have created new opportunities for addressing this challenge. SSL methods learn informative representations directly from large amounts of unlabeled data by solving auxiliary pretext tasks, thereby reducing dependence on labeled fault samples. Recent surveys have highlighted SSL as one of the most promising research directions for intelligent fault diagnosis because of its ability to exploit abundant unlabeled operational data while improving transferability across operating conditions (Wang, Chu, Han, Kong, & Huang, 2024). Among these methods, Prediction of Functionals from Masked Latents (PFML) has emerged as a promising framework for industrial condition monitoring (Vaaras, Airaksinen, & Räsänen, 2025). PFML learns latent representations by predicting physically meaningful signal functionals from masked observations, encouraging the model to capture intrinsic characteristics of the underlying signals. Building upon this concept, the Domain-Informed Functional Guidance Network (DiFG-Net) was recently proposed to incorporate domain knowledge into SSL-based functional learning for domain-adaptive fault diagnosis (He & He, 2026). By leveraging handcrafted signal functionals as learning targets, DiFG-Net demonstrated encouraging performance in transferring diagnostic knowledge across different fault domains.

Despite its effectiveness, DiFG-Net exhibits several limitations when applied to cross-speed fault diagnosis. First, the framework relies heavily on handcrafted functionals that primarily characterize signal statistics and spectral properties without explicitly accounting for rotational-speed effects. Second, the learned latent representations are not constrained to remain invariant across different speed conditions. Consequently, the model may inadvertently encode speed-dependent information, reducing its ability to generalize to unseen operating speeds. Finally, DiFG-Net lacks a mechanism to extrapolate beyond the speed range covered in the training data, limiting its applicability to open-set cross-speed diagnosis scenarios frequently encountered in real-world industrial systems.

To address these limitations, this paper proposes DiFG-Net++, a physics-informed functional learning framework specifically designed for cross-speed fault diagnosis. The proposed framework extends DiFG-Net by integrating three novel speed-adaptation strategies. First, a healthy-state target-speed anchoring mechanism leverages unlabeled healthy vibration data collected at the target operating speed to provide an operational reference for speed adaptation. This design is motivated by the practical observation that healthy-condition data are often readily available in industrial environments, whereas fault data at new operating speeds are scarce or unavailable. Second, a speed-invariant latent regularization strategy is introduced to encourage the learned representations to remain consistent across different speed conditions while preserving fault-discriminative characteristics. Third, a speed-extrapolative augmentation mechanism is developed to generate physically meaningful intermediate and extrapolated speed conditions through time-scale transformations guided by rotational-speed relationships. Together, these strategies enable the model to learn speed-invariant and fault-sensitive representations that generalize more effectively to unseen operating speeds.

The effectiveness of the proposed DiFG-Net++ framework is evaluated using a plastic bearing dataset collected under multiple rotational speeds and bearing health conditions. Extensive experiments are conducted under challenging cross-speed diagnosis settings, where the target operating speed is excluded from the training set. Comparative studies against baseline methods demonstrate that the proposed speed-adaptation strategies substantially improve diagnostic performance, particularly under open-set speed extrapolation scenarios. The results suggest that integrating physics-informed functional learning with speed-aware representation learning provides a promising approach to robust, practical cross-speed fault diagnosis for industrial prognostics and health management (PHM) applications.

2. BACKGROUND AND RELATED WORK

SSL and PFML related background information and a review of related work are provided in the following sections.

2.1. SSL and PFML

The remarkable success of deep learning in fault diagnosis has largely depended on the availability of large volumes of labeled condition-monitoring data. However, obtaining labeled fault data in industrial environments is often expensive, time-consuming, and impractical, as machinery typically operates under healthy conditions for most of its service life. To address this challenge, SSL has emerged as an attractive paradigm for learning informative representations from large volumes of unlabeled data. Rather than relying on manually annotated labels, SSL constructs auxiliary pretext tasks that enable neural networks to learn meaningful latent representations directly from raw observations. Recent advances in SSL, including contrastive learning, masked reconstruction, and predictive learning approaches, have demonstrated remarkable success across the domains of computer vision, natural language processing, and time-series analysis (Devlin, Chang, Lee, & Toutanova, 2019; Chen, Kornblith, Norouzi, & Hinton, 2020; Grill, Strub, Altché, Tallec, Richemond, Buchatskaya, Doersch, Pires, Guo, Azar, Piot, Kavukcuoglu, Munos, & Valko, 2020; He, Chen, Xie, Li, Dollár, & Girshick, 2022).

In the field of PHM, SSL has attracted growing attention for leveraging abundant unlabeled sensor data collected during machine operation. Existing SSL-based fault diagnosis approaches have employed contrastive objectives, masked signal reconstruction, and predictive coding techniques to learn representations that can subsequently be transferred to downstream diagnostic tasks (Mohsenvand, Izadi, & Maes, 2020; Eldele, Ragab, Chen, Wu, Kwoh, Li, & Guan, 2021; Zhang, Li, Ma, Luo, & Li, 2022; Wang, Chu, Han, Kong, & Huang, 2024). Although these methods reduce dependence on labeled data, many primarily focus on reconstructing signals or maximizing representation consistency without explicitly incorporating domain knowledge regarding machine behavior and fault characteristics.

To bridge this gap, Prediction of Functionals from Masked Latents (PFML) was recently proposed. Inspired by masked representation learning, PFML first partitions a time-series signal into frames and randomly masks a subset of latent representations generated by an encoder network. Instead of reconstructing the original signal, PFML predicts physically meaningful signal functionals derived from the input measurements. The implication of PFML for PHM applications is that, by predicting functionals rather than raw signals, PFML encourages the encoder to learn latent representations directly related to machine condition and fault-sensitive characteristics.

The PFML framework offers several advantages for fault diagnosis. First, the functional prediction task provides richer supervisory signals than conventional reconstruction-based SSL methods because the network must learn relationships between latent representations and physically meaningful signal characteristics. Second, the framework can effectively utilize large quantities of unlabeled operational data, making it particularly suitable for industrial applications where fault labels are scarce. Third, the learned representations can be transferred to downstream classification tasks through fine-tuning, enabling improved diagnostic performance under limited labeled data conditions. These properties make PFML an attractive foundation for developing physics-informed and domain-adaptive fault diagnosis frameworks. Nevertheless, the effectiveness of PFML depends heavily on the design of the predicted functionals and does not explicitly address domain shifts caused by varying operating conditions, motivating subsequent research on domain-informed functional learning approaches such as DiFG-Net.

2.2. DiFG-Net

Building upon the PFML framework, DiFG-Net was developed to improve representation learning under domain-shift conditions commonly encountered in industrial fault diagnosis. In practical PHM applications, machinery often operates under varying loads, environmental conditions, sensor configurations, and machine-to-machine variations, leading to significant discrepancies in distribution between the source and target domains. Such domain shifts frequently degrade the performance of conventional data-driven diagnostic models because the learned representations are sensitive to changes unrelated to machine health. Domain adaptation techniques have been widely investigated to address this challenge through distribution alignment, adversarial learning, and discrepancy minimization strategies (Pan & Yang, 2010; Ganin, Ustinova, Ajakan, Germain, Larochelle, Laviolette, Marchand, & Lempitsky, 2016; Wang & Deng, 2018; Zhang et al., 2019). However, many existing approaches focus primarily on feature alignment without explicitly incorporating engineering knowledge regarding fault-sensitive signal characteristics.

DiFG-Net addresses this limitation by integrating domain knowledge directly into the SSL representation learning process. The key idea is to guide latent representation learning using a set of domain-informed functionals that capture fault-relevant characteristics of vibration signals. Like PFML, DiFG-Net employs a masked latent prediction framework. However, instead of relying solely on generic signal descriptors, the framework uses carefully selected functionals that capture key diagnostic information about machinery faults. These functionals serve as intermediate learning targets that encourage the encoder to capture domain-invariant and fault-sensitive features while suppressing irrelevant variations introduced by operating conditions. The DiFG-Net framework is provided in Figure 1.

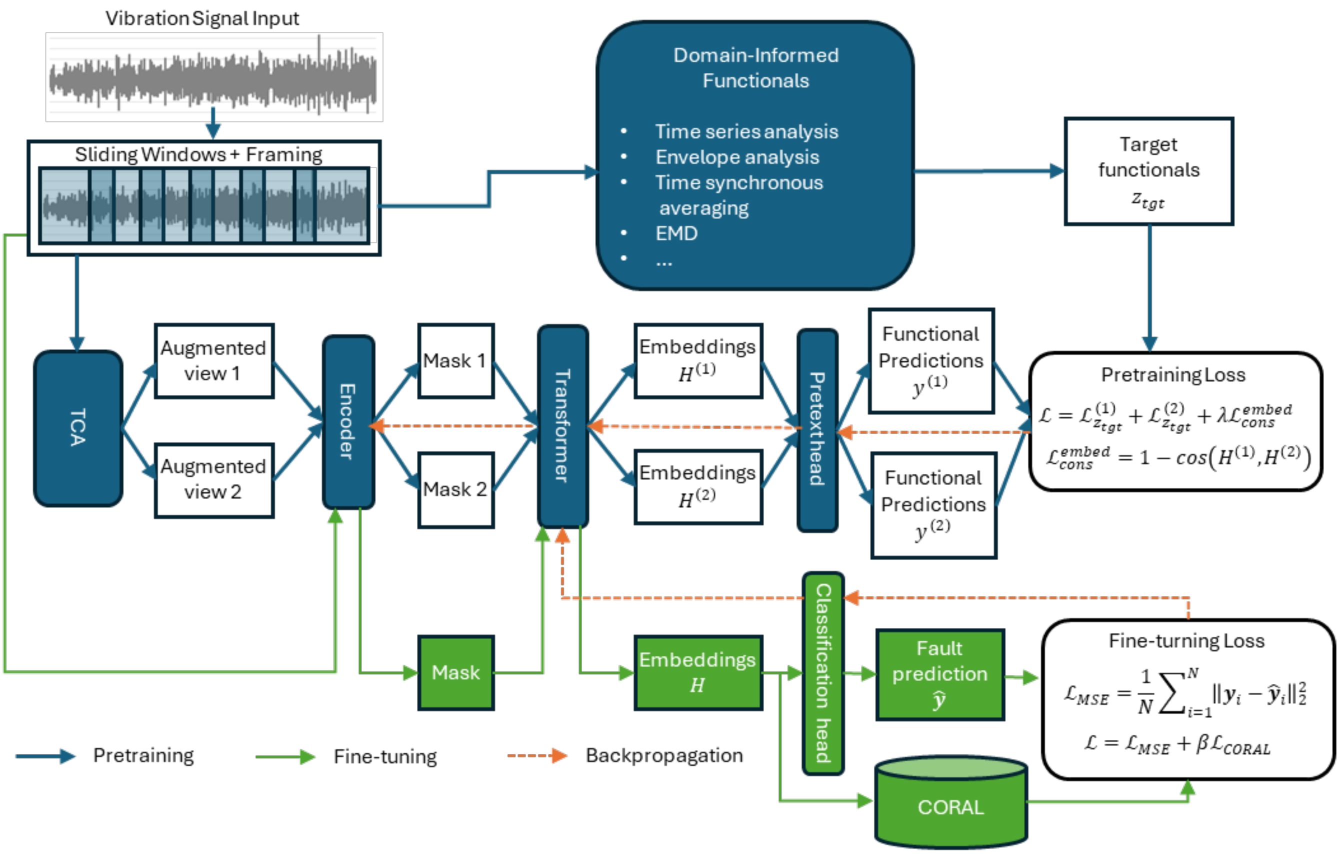

Figure 1. The DiFG-Net framework (He & He, 2026)

A distinguishing characteristic of DiFG-Net is its ability to leverage unlabeled data from both source and target domains during pretraining. Through functional prediction and latent consistency objectives, the model learns representations that preserve machine health information while reducing sensitivity to domain-specific variations. The learned encoder can then be transferred to downstream fault classification tasks through supervised fine-tuning. Experimental studies have demonstrated that DiFG-Net significantly improves cross-domain diagnostic performance compared with conventional SSL and transfer learning approaches under various domain-shift scenarios.

Despite these advantages, DiFG-Net exhibits several limitations when applied to cross-speed fault diagnosis. First, the functional targets primarily consist of handcrafted statistical and spectral descriptors that do not explicitly encode rotational-speed information. Second, the framework lacks mechanisms for enforcing speed-invariant latent representations across different operating speeds. Consequently, latent features may still capture speed-dependent characteristics rather than fault-related information. Third, DiFG-Net does not incorporate strategies for extrapolating to operating speeds outside the range observed during training, which is a critical requirement in open-set cross-speed fault diagnosis. These limitations motivate the development of DiFG-Net++, which extends the original framework through healthy-state target-speed anchoring, speed-invariant latent regularization, and speed-extrapolative augmentation to improve diagnostic robustness under unseen operating speeds.

Note that, since each vibration signal input file is sliced into multiple windows, classification is measured at either the window or file level.

The file-level accuracy is computed from window-level predictions, but not by the majority vote of window labels. Instead, it uses average logits across all windows belonging to the same file. For every window (chunk) i, the classifier produces: $logits[i]$ and these logits are stored according to the file path. For example, if a file is sliced into 12 windows:

File A

├── Window 1 → logits1

├── Window 2 → logits2

├── Window 3 → logits3

...

└── Window12 → logits12

Let $\mathbf{z}_{f,i}$ be the logits vector for window i of file f and N_f be the number of windows belonging to file f. Then the mean logits vector for file f can be computed as:

$$\bar{\mathbf{z}}_f = \frac{1}{N_f}\sum_{i=1}^{N_f} \mathbf{z}_{f,i}$$

Let $\bar{z}_{f,c}$ be the mean logit of file f for class c. Then the prediction for file f is:

$$prediction_f = \operatorname*{argmax}_c \bar{z}_{f,c}$$

Averaging logits is better than majority voting. For example, suppose we have 6 windows predicting ball fault and 6 predicting healthy. If we use majority voting, there will be a tie. If we use averaging logits, we will still produce a confidence-weighted decision because highly confident windows contribute more. This is why most modern MIL (Multiple Instance Learning) systems use average logits or average probabilities instead of hard voting. For PHM applications, file-level accuracy is usually the more meaningful metric because a maintenance engineer does not classify window #37; they classify record A.

3. DIFG-NET++

Although DiFG-Net has demonstrated promising performance for cross-domain fault diagnosis through domain-informed functional learning, its ability to generalize across unseen rotational speeds remains limited. In particular, cross-speed fault diagnosis presents a fundamentally different challenge from conventional domain adaptation because vibration signatures and fault characteristic frequencies vary systematically with shaft speed. The problem becomes even more difficult in open-set scenarios where the target operating speed lies outside the range observed during training, and fault data at that speed are unavailable. Under such conditions, simply aligning feature distributions across domains is insufficient because the learned representations must simultaneously preserve fault-sensitive information while remaining invariant to speed-induced variations.

To address these challenges, this paper proposes DiFG-Net++, an enhanced physics-informed functional learning framework specifically designed for cross-speed fault diagnosis. Building upon the PFML-based functional prediction framework of DiFG-Net, DiFG-Net++ introduces three complementary speed-adaptation strategies that operate at different stages of the learning process. First, a speed-extrapolative augmentation mechanism generates physically meaningful intermediate and extrapolated operating conditions through speed-guided time-scale transformations, enabling the model to observe a broader range of speed variations during self-supervised pretraining. Second, a speed-invariant latent regularization strategy encourages latent representations associated with the same machine health condition to remain consistent across different rotational speeds, thereby improving cross-speed generalization. Third, a healthy-state target-speed anchoring mechanism leverages unlabeled healthy vibration data collected at the target operating speed to provide an operational reference for adapting the learned representations to previously unseen speed conditions. Collectively, these three mechanisms enable DiFG-Net++ to learn representations that are both fault-sensitive and speed-invariant, thereby improving diagnostic robustness under challenging cross-speed operating conditions. The overall framework is illustrated in Figure 2, and its three key components are described in the following subsections.

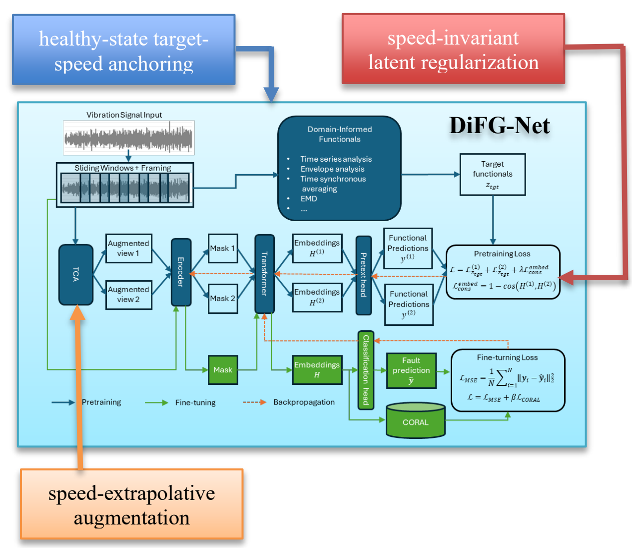

Figure 2. The framework of DiFG-Net++.

3.1. Speed-Extrapolative Augmentation

In SSL applications, if the target speed falls within the training speed range, e.g., target speed = 40 Hz and the training speed range is in the range of [10 Hz, 60 Hz], it is considered an interpolation case. On the other hand, if the target speed falls outside the training speed range, e.g., target speed = 60 Hz, and the training speed range is in the range of [10 Hz, 40 Hz], it is considered an extrapolation case. It is obvious that learning for the extrapolation cases is more difficult than for the interpolation cases. To facilitate the learning of the DiFG-Net++ for the extrapolation cases, we want to add an augmentation in the pre-training stage to teach the model:

10 Hz → looks like 20, 40, 60 Hz

20 Hz → looks like 10, 40, 60 Hz

40 Hz → looks like 10, 20, 60 Hz

We call this augmentation mechanism speed-extrapolative augmentation. In fact, many vibration-based functionals built in the order domain can scale with rotational speed. An order is defined as the ratio between a characteristic vibration frequency and the shaft rotational frequency. Unlike frequency-domain features, order-domain functionals remain invariant under speed changes because both the characteristic frequency and shaft frequency scale proportionally with rotational speed. Consequently, fault orders provide a physically meaningful representation for cross-speed fault diagnosis and serve as the foundation for the speed-adaptive mechanisms introduced in DiFG-Net++.

Here, we define fault order o as a dimensionless frequency ratio, i.e., a frequency normalized by the shaft rotational speed:

$$o = \frac{f}{f_r} \tag{1}$$

where:

f = characteristic fault frequency (Hz)

f_r= shaft rotational frequency (Hz)

The fault order is speed-invariant and is typically physics-based. Take bearing characteristic frequency as an example. For a given bearing geometry:

$$o_{BPFO} = \frac{n}{2}\left(1 - \frac{d}{D}\cos\theta\right) \tag{2}$$

$$o_{BPFI} = \frac{n}{2}\left(1 + \frac{d}{D}\cos\theta\right) \tag{3}$$

$$o_{BSF} = \frac{D}{2d}\left[1 - \left(1 - \frac{d}{D}\cos\theta\right)^2\right] \tag{4}$$

$$o_{FTF} = \frac{1}{2}\left(1 - \frac{d}{D}\cos\theta\right) \tag{5}$$

where:

n = number of rolling elements

d = ball diameter

D = pitch diameter

θ = contact angle

These bearing fault orders are determined by the bearing geometry and are approximately constant across speeds. For example, at 10 Hz, f_{BPFO} = 35.85 Hz; at 20 Hz, f_{BPFO} = 71.7 Hz; at 40 Hz, f_{BPFO} = 143.4 Hz. The bearing fault frequency changes with rotational speed, but the bearing fault order remains: $o_{BPFO} = \frac{f_{BPFO}}{f_r}$ =3.585.

Since fault orders are primarily determined by geometry and remain approximately invariant across operating speeds, they provide a physically meaningful representation for learning speed-invariant fault characteristics and a solid foundation for using speed-extrapolative augmentation during pretraining.

Although order-domain functionals are inherently invariant to rotational speed, the raw vibration signals used to train the encoder remain strongly speed-dependent because characteristic fault frequencies scale proportionally with shaft speed. Consequently, the encoder may fail to generalize to operating speeds outside the training range even when order-domain functional targets are employed. To expose the model to a broader spectrum of speed conditions and improve open-set speed generalization, a speed-extrapolative augmentation strategy is introduced to generate physically meaningful intermediate and extrapolated operating speeds based on the order-frequency relationship defined in Eq. (1).

In DiFG-Net++, the speed-extrapolative augmentation is implemented as follows. Let $f_{r,aug}$ denote the augmented rotational speed generated from the original speed $f_{r,s}$. The augmentation scale factor is defined as: $\alpha = \frac{f_{r,aug}}{f_{r,s}}$. $f_{r,aug}$ is randomly sampled based on the original speed and target speed as: `random.uniform(speed_range)`. Let $f_{r,t}$ be the target speed, then, *speed_range* = [$min(f_{r,s}, f_{r,t})$, $max(f_{r,s}, f_{r,t})$]. *speed_range* will be [10, 60] if the original speed is 10 Hz and the target speed is 60 Hz. *speed_range* will be [10, 40] if the original speed is 40 Hz and the target speed is 10 Hz. The input signal $x(t)$ is then time-scaled according to the augmented signal as: $x_{aug}(t) = x\left(\frac{t}{\alpha}\right)$. This transformation preserves the underlying fault orders while shifting the characteristic frequencies according to the speed-order relationship $f = of_r$ defined by Eq. (1). For example, converting 20 Hz to an augmented 60 Hz condition:

$$\alpha = \frac{60}{20} = 3$$

so the signal is compressed to one-third of its original length before padding/cropping back to the original window size. This shifts frequency components upward by a factor of 3, consistent with $f = of_r$ defined by Eq. (1).

3.2. Speed-invariant Latent Regularization

Although the proposed speed-extrapolative augmentation exposes the encoder to a broader range of operating speeds, it does not explicitly constrain the learned latent representations to be invariant to speed variations. In practice, vibration signals acquired under different rotational speeds often exhibit significant changes in their temporal and spectral characteristics because fault characteristic frequencies scale proportionally with shaft speed. Consequently, latent representations learned solely through functional prediction may inadvertently encode operating-speed information alongside fault-related characteristics. This issue becomes particularly problematic in open-set cross-speed fault diagnosis, where the target operating speed lies outside the speed range observed during training. In such cases, samples associated with the same fault type but collected at different rotational speeds may occupy separate regions of the latent space, leading to degraded generalization performance.

To address this challenge, a speed-invariant latent regularization strategy is introduced to encourage fault-consistent representations across different operating speeds. The key idea is that although vibration signatures shift in the frequency domain as rotational speed changes, the underlying fault mechanism remains unchanged. Therefore, latent representations associated with the same machine health condition should remain similar across operating speeds. By explicitly minimizing discrepancies among latent representations originating from different speed conditions, the proposed regularization guides the encoder to suppress speed-dependent variations while preserving fault-sensitive information. As a result, the learned latent space is organized by machine health conditions rather than operating speeds, thereby improving the model's ability to generalize to previously unseen rotational speeds.

Let $\boldsymbol{z}_i$ be the latent representation of the input $\boldsymbol{x}_i$ in the latent space as:

$$\boldsymbol{z}_i = g(\boldsymbol{x}_i)$$

In DiFG-Net, since rotational speed information may still be coded in the latent space, as a result, the latent representation of the same fault state may vary across different speeds:

$$\boldsymbol{z}_{10\text{ Hz}}^{ball\,fault} \neq \boldsymbol{z}_{20\text{ Hz}}^{ball\,fault} \neq \boldsymbol{z}_{40\text{ Hz}}^{ball\,fault}$$

This makes open-set cross-speed fault diagnosis difficult. By using speed-extrapolative augmentation and speed-invariant latent regularization in DiFG-Net++, we can create augmented versions of the same signal at different speeds:

$$\boldsymbol{x}_{10\text{ Hz}}^{ball\,fault}, \boldsymbol{x}_{15\text{ Hz}}^{ball\,fault}, \boldsymbol{x}_{20\text{ Hz}}^{ball\,fault}, \dots, \boldsymbol{x}_{65\text{ Hz}}^{ball\,fault}$$

Let us define a speed latent loss function for speed-invariant latent regularization as:

$$\mathcal{L}_{speed\ latent} = \frac{1}{N}\sum_i \left\|\boldsymbol{z}_i^k - \boldsymbol{z}_j^k\right\|_2^2 \qquad (6)$$

where: $\boldsymbol{z}_i^k$ and $\boldsymbol{z}_j^k$ represent the latent functional representation of the same state k at two different speeds i and j.

By minimizing $\mathcal{L}_{speed\ latent}$, we can enforce:

$$\boldsymbol{z}_{10\text{ Hz}}^{ball\,fault} \approx \boldsymbol{z}_{20\text{ Hz}}^{ball\,fault} \approx \boldsymbol{z}_{40\text{ Hz}}^{ball\,fault} \approx \boldsymbol{z}_{60\text{ Hz}}^{ball\,fault}$$

In this way, we can guide DiFG-Net++ to learn functions that are fault-sensitive but speed-invariant for effective open-set cross-speed fault diagnosis.

3.3. Healthy-state Target-speed Anchoring

Although speed-extrapolative augmentation and speed-invariant latent regularization could improve the ability of DiFG-Net++ to learn speed-robust representations, they do not explicitly leverage information from the target operating speed. In practical industrial environments, machines typically operate under healthy conditions for extended periods, making healthy vibration data readily available across a wide range of operating speeds. In contrast, fault data at newly introduced or previously unseen operating speeds are often unavailable due to the rarity of fault occurrences and the high cost of collecting them. Consequently, cross-speed fault diagnosis is frequently characterized by an asymmetric data availability scenario in which abundant healthy data exists at the target speed, while fault data are available only at source speeds. Conventional transfer learning and domain adaptation methods generally overlook this valuable source of information and treat the target domain as entirely unlabeled.

To exploit this practical characteristic, a healthy-state target-speed anchoring strategy is introduced in DiFG-Net++. The key idea is to utilize healthy vibration signals collected at the target operating speed as an operational reference for speed adaptation. Because healthy signals reflect the underlying dynamic characteristics of the target operating condition without contamination from fault-induced behaviors, they provide valuable information regarding the distribution shift introduced by the new speed. By anchoring the learned representations to the healthy-state characteristics of the target domain, the proposed approach enables the encoder to adapt to previously unseen operating speeds without requiring target-speed fault samples. As a result, the latent representations become more aligned with the target operating condition while preserving fault-sensitive information learned from the source-speed domains. This mechanism is particularly advantageous for open-set cross-speed fault diagnosis, where the target speed lies outside the training-speed range and fault data at that speed are unavailable.

From a representation-learning perspective, healthy-state target-speed anchoring serves as a bridge between the source-speed domain and the unseen target-speed domain. Since healthy machine behavior is governed primarily by operating conditions rather than fault-induced dynamics, healthy vibration signals collected at different rotational speeds provide a reliable characterization of speed-dependent distribution shifts. By exposing the encoder to healthy signals at the target speed during training, DiFG-Net++ learns the normal operating characteristics associated with the previously unseen speed while maintaining fault-discriminative knowledge acquired from source-speed fault data. Consequently, the model can distinguish variations due to changes in operating speed from those due to fault progression. This capability is particularly important in cross-speed fault diagnosis because the target-speed healthy data effectively establishes a reference operating manifold upon which fault-sensitive deviations can be detected and classified.

Healthy-state target-speed anchoring in DiFG-Net++ is implemented by incorporating healthy vibration signals collected at the target operating speed into the representation learning process. Unlike fault samples, which may be unavailable at previously unseen operating speeds, healthy operating data are typically abundant in industrial environments and provide valuable information regarding the operating characteristics of the target speed. During training, these healthy target-speed samples serve as an operational reference, enabling the encoder to learn the distribution shift associated with the new speed while preserving fault-sensitive knowledge learned from source-speed fault data. Consequently, the learned latent representations become aligned with the target operating condition, facilitating improved cross-speed fault diagnosis without requiring target-speed fault labels.

It is important to note that healthy-state target-speed anchoring complements rather than replaces the other speed-adaptation mechanisms introduced in DiFG-Net++. Speed-extrapolative augmentation expands the range of speed conditions observed during self-supervised learning, while speed-invariant latent regularization encourages fault-consistent representations across different operating speeds. Healthy-state target-speed anchoring further adapts these representations to the actual target operating condition by leveraging readily available healthy vibration data. Together, these three mechanisms operate at the data, representation, and domain levels, respectively, enabling DiFG-Net++ to achieve robust open-set cross-speed fault diagnosis without requiring fault samples from the target speed.

4. THE CASE STUDY

In this section, a case study on open-set cross-speed fault diagnosis using a plastic bearing dataset is presented to demonstrate the effectiveness of DiFG-Net++.

4.1. The Experimental Setup and the Dataset

The dataset used in this case study was collected from plastic-bearing seeded-fault tests conducted in the laboratory. Figure 3 shows the bearing test rig used for the tests.

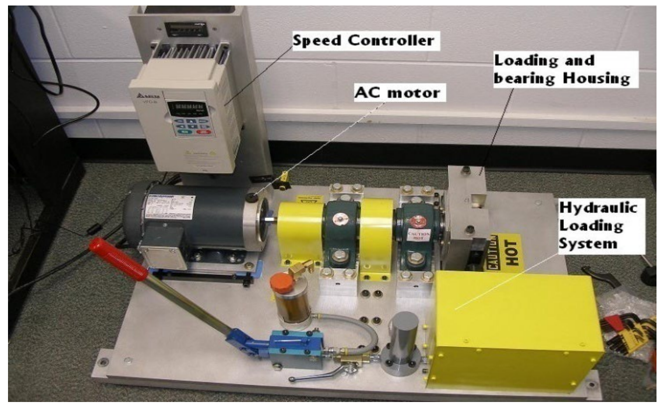

Figure 3. The bearing test rig for data collection in the laboratory.

Four types of bearing faults were seeded in 6205-2RS type plastic bearings: inner race fault, outer race fault, ball fault, and cage fault (see Fig. 4). The damage to the contact surfaces of the bearing inner and outer races was caused by scratching the race surfaces with an electric solder iron heated to a high temperature. The diameter of the damaged surface area was about one-third of the ball diameter. The ball damage was created by scratching one of the bearing balls with a grinding wheel. Roughly 40% of the ball volume was removed by grinding. The broken cage was created by cutting the Teflon retainer using a pair of sharp scissors.

During the tests, vibration data were collected at a sampling rate of 102.4 kHz. A total of 4 shaft rotational speeds, 10 Hz, 20 Hz, 40 Hz, and 60 Hz, were used during the testing. The accelerometers were mounted on the surface and on the side of the bearing housing. The locations of the accelerometers are shown in Figure 5. More details on the test settings can be found in He, Li, & Zhu (2013). In this case study, only the vibration data from the sensor mounted on top of the bearing housing was used.

Figure 4. The 4 seeded bearing faults used in the tests: inner race fault, outer race fault, ball fault, and cage fault.

(a)

Figure 5. The layout of the two accelerometers, screw-mounted on the bearing housing in the bearing test rig.

4.2. The Computational Setup

4.2.1. The Order-Domain Functionals

The original DiFG-Net already uses DiF-1, DiF-2, and DiF-3, including envelope-spectrum BFCF amplitudes and TSR-based functionals at operating speeds of 10, 20, 40, and 60 Hz. In this paper, we treat speed not just as an experimental condition but as a self-supervised physical variable that the model must learn consistently.

We added a new functional group, DiF-4: Normalized Order-Domain Functionals. The DiF-4 consists of the following functionals: A_1X_{norm}, A_2X_{norm}, A_3X_{norm}, A_BPFO_{norm}, A_BPFI_{norm}, A_BSF_{norm}, A_TFT_{norm}. A_1X_{norm}, A_2X_{norm}, and A_3X_{norm} are order-domain amplitudes extracted from the TSA signal and then normalized by the total energy of the selected order components. A_BPFO_{norm}, A_BPFI_{norm}, A_BSF_{norm}, and A_TFT_{norm} are normalized bearing fault amplitudes. Let A_1X, A_2X, and A_3X be the 1st, 2nd, and 3rd orders of the FFT amplitude of the TSA signal. Let $X_{TSA}(f)$denote the FFT of the TSA signal.

$$A\_BPFO = |X_{TSA}(o_{BPFO} f_r)| \tag{7}$$

$$A\_BPFI = |X_{TSA}(o_{BPFi} f_r)| \tag{8}$$

$$A\_BSF = |X_{TSA}(o_{BSF} f_r)| \tag{9}$$

$$A\_FTF = |X_{TSA}(o_{FTF} f_r)| \tag{10}$$

$$S = A\_1X + A\_2X + A\_3X + A\_BPFO + A\_BPFI + A\_BSF + A\_FTF \tag{11}$$

Then the normalized order-domain functionals can be computed as:

$$A\_1X_{norm} = \frac{A\_1X}{S} \tag{12}$$

$$A\_2X_{norm} = \frac{A\_2X}{S} \tag{13}$$

$$A\_3X_{norm} = \frac{A\_3X}{S} \tag{14}$$

$$A\_BPFO_{norm} = \frac{A\_BPFO}{S} \tag{15}$$

$$A\_BPFI_{norm} = \frac{A\_BPFI}{S} \tag{16}$$

$$A\_BSF_{norm} = \frac{A\_BSF}{S} \tag{17}$$

$$A\_TFT_{norm} = \frac{A\_FTF}{S} \tag{18}$$

Using these order-domain functionals as targets in the pre-training stage, the model will learn how energy is redistributed among physically meaningful orders rather than absolute vibration levels. This normalization helps targets work well for cross-speed diagnosis, as they encode the relative-order-energy signature of a fault, which is much more speed-invariant than raw spectral amplitudes.

4.2.2. Healthy-state Target-speed Anchoring Protocol

To evaluate the effectiveness of the proposed healthy-state target-speed anchoring strategy, a modified open-set cross-speed fault diagnosis protocol was adopted. For each holdout-speed experiment, three datasets were constructed: a self-supervised pre-training dataset, a supervised fine-tuning dataset, and a validation dataset. The pre-training dataset was used to train the DiFG-Net++ backbone using the proposed PFML objective, while the fine-tuning dataset was used to jointly optimize the backbone and the classification head. In both the pre-training and fine-tuning datasets, all healthy-state vibration signals collected at the holdout speed were retained, whereas all fault-condition signals collected at the holdout speed were excluded. For example, when 60 Hz was selected as the holdout speed, all 60 Hz healthy-state data were included in both training datasets, while the 60 Hz ball fault, cage fault, inner-race fault, and outer-race fault data were removed. The validation dataset consisted exclusively of fault-condition data collected at the holdout speed. Consequently, for the 60 Hz holdout experiment, only the 60 Hz ball fault, cage fault, inner-race fault, and outer-race fault data were included in the validation dataset, while the 60 Hz healthy-state data were excluded.

This protocol reflects practical industrial operating conditions in which healthy vibration data are routinely collected across a wide range of operating speeds, whereas fault data at newly introduced or previously unseen speeds are scarce or unavailable. As a result, the model has access to healthy operational information at the target speed but is not exposed to any target-speed fault samples during either the self-supervised pre-training or supervised fine-tuning stages. The resulting experimental setting, therefore, represents the problem for open-set cross-speed fault diagnosis, in which the model must transfer fault knowledge learned from source-speed domains to an unseen target-speed domain using only healthy target-speed information.

4.3. The Results

The results of the computational experiment are provided in Table 1.

Table 1. Summary of the computational experiment.

Method	Speed-extrapolative augmentation	Speed-invariant Latent Regularization	Healthy-state Target-speed Anchoring	Validation Accuracy (%)				
				Holdout speed 10 Hz	Holdout speed 20 Hz	Holdout speed 40 Hz	Holdout speed 60 Hz	Overall
1. DiFG-Net	No	No	No	60.00±0.1054	40.00±0.0982	46.70±0.1027	40.00±0.1003	46.68
2. DiFG-Net	No	No	Yes	75.00±0.0874	50.00±0.1012	91.70±0.0914	66.70±0.1011	70.85
3. DiFG-Net++	Yes	Yes	No	77.8±0.1101	86.7±0.1021	86.7±0.1051	86.70±0.1032	82.8
4. DiFG-Net++	Yes	Yes	Yes	93.85±0.088	93.85±0.0954	98.50±0.1022	97.20±0.0931	95.85

The results in Table 1 demonstrate that each of the proposed speed-adaptation strategies improves cross-speed fault diagnosis performance. The baseline DiFG-Net, which relies solely on PFML-based functional learning without explicit speed adaptation, achieved an overall validation accuracy of only 46.68%, indicating limited generalization to unseen operating speeds. The lowest accuracies were observed for the 20 Hz and 60 Hz holdout experiments, suggesting that the learned representations remained highly sensitive to speed-induced variations.

When healthy-state target-speed anchoring was introduced while retaining the original DiFG-Net architecture, the overall accuracy increased substantially from 46.68% to 70.85%. This result confirms that healthy vibration data collected at the target operating speed provides valuable information regarding the operating-condition shift associated with previously unseen speeds. By leveraging healthy target-speed data as an operational reference, the model became more robust to speed variations even without explicit speed-invariant learning mechanisms.

The introduction of the proposed DiFG-Net++ framework, which combines speed-extrapolative augmentation and speed-invariant latent regularization, further improved overall accuracy to 82.8%. Compared with the baseline DiFG-Net, the improvement exceeded 36 percentage points, demonstrating that exposing the encoder to a broader range of speed conditions while explicitly enforcing speed-invariant representations substantially enhances cross-speed generalization. Notably, accuracies became consistently high across all holdout-speed experiments, indicating that the learned latent space became less dependent on operating speed.

The best performance was achieved when all three proposed mechanisms were combined. DiFG-Net++ with healthy-state target-speed anchoring achieved an overall accuracy of 95.85%, representing an absolute improvement of approximately 49 percentage points over the original DiFG-Net. These results demonstrate the complementary nature of the three speed-adaptation strategies and validate the effectiveness of the proposed multi-level speed adaptation framework.

Figure 6 presents the t-SNE visualizations of the latent representations learned by the backbone and the classifier-input features after fine-tuning for the four evaluated methods under the 60 Hz holdout-speed condition. The progression of the feature distributions provides important insight into the performance improvements summarized in Table 1.

For Method 1 (DiFG-Net), the latent representations exhibit substantial overlap among several fault categories, and the classifier-input features remain fragmented after fine-tuning. Although some fault classes form local clusters, significant inter-class mixing persists, indicating that the learned representations remain sensitive to operating-speed variations. This behavior is consistent with the relatively poor classification performance reported in Table 1, where the overall validation accuracy is only 46.68%.

For Method 2 (DiFG-Net with healthy-state target-speed anchoring), the latent space becomes noticeably more structured. The incorporation of healthy target-speed data helps align the feature distributions toward the target operating condition, thereby improving cluster compactness after fine-tuning. Nevertheless, several classes remain partially intertwined, suggesting that speed-induced variations persist in the learned representations. Consequently, the overall accuracy improves to 70.85%, but there is substantial room for improvement.

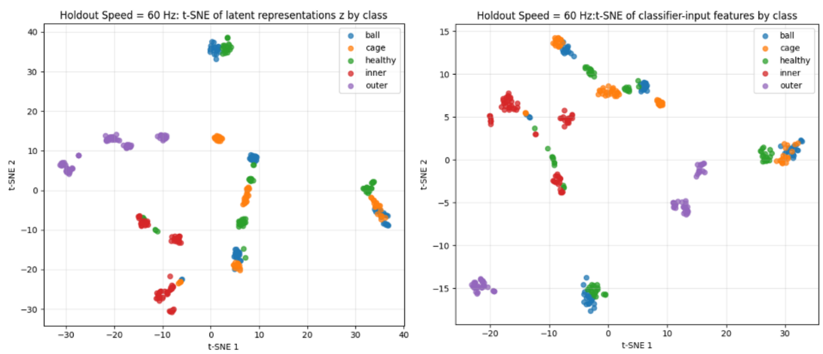

(a) Method 1: t-SNE plots of latent representations and classifier features at 60 Hz

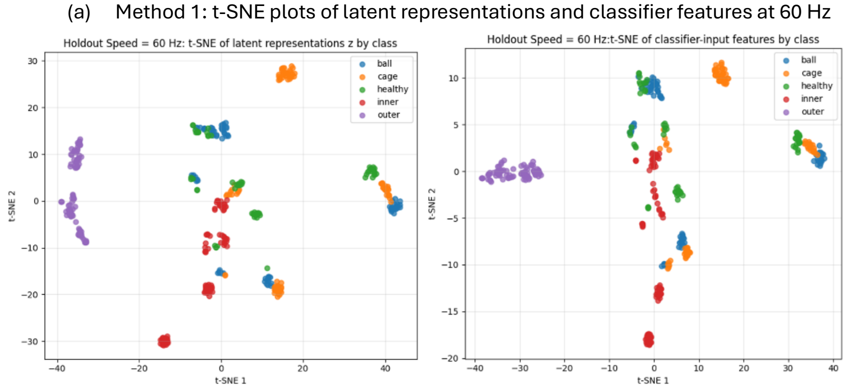

(b) Method 2: t-SNE plots of latent representations and classifier features at 60 Hz

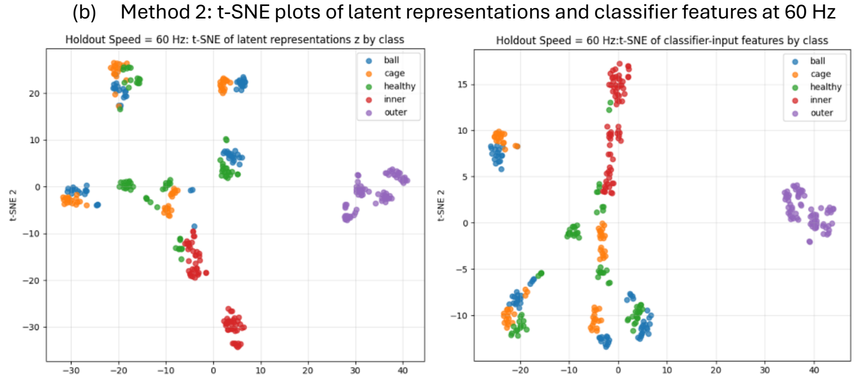

(c) Method 3: t-SNE plots of latent representations and classifier features at 60 Hz

(d) Method 4: t-SNE plots of latent representations and classifier features at 60 Hz

Figure 6. The t-SNE plots of latent representations and classifier features at holdout speed 60 Hz for the 4 methods in Table 1.

For Method 3 (DiFG-Net++ with speed-extrapolative augmentation and speed-invariant latent regularization), the separation among fault classes becomes significantly more pronounced. Compared with Methods 1 and 2, the latent representations exhibit improved class compactness and reduced overlap, indicating that the proposed speed-adaptation mechanisms successfully suppress a large portion of the speed-dependent variability. The classifier-input features further enhance this separation, producing more distinct fault clusters. These observations are reflected in the considerable increase in overall validation accuracy to 82.8%.

The most striking result is observed in Method 4 (DiFG-Net++ with all three proposed speed-adaptation strategies). Both the latent representations and the classifier-input features exhibit highly compact intra-class distributions and clear inter-class separation. In particular, the outer-race, inner-race, cage, and ball fault clusters become well separated with minimal overlap, while the healthy-state samples occupy distinct regions of the feature space. Compared with the preceding methods, the classifier-input features demonstrate substantially improved cluster compactness, indicating that the supervised fine-tuning stage further refines the PFML-learned representations into fault-discriminative and speed-robust features. The resulting feature geometry closely corresponds to the superior classification performance reported in Table 1, where Method 4 achieves the highest overall validation accuracy of 95.85%.

Overall, Figure 6 demonstrates a progressive evolution of the learned feature space from speed-sensitive and partially overlapping representations in DiFG-Net to highly discriminative and speed-invariant representations in DiFG-Net++. The results suggest that healthy-state target-speed anchoring, speed-extrapolative augmentation, and speed-invariant latent regularization contribute complementary benefits. Healthy-state target-speed anchoring improves target-domain alignment, while speed-extrapolative augmentation and speed-invariant latent regularization reduce speed-dependent variability and enhance fault discriminability. When combined, these mechanisms produce the most compact and well-separated feature distributions, which directly translate into the superior cross-speed fault diagnosis performance reported in Table 1.

In summary, the results in Table 1 indicate that speed-extrapolative augmentation, speed-invariant latent regularization, and healthy-state target-speed anchoring address distinct aspects of the cross-speed adaptation problem and collectively enable robust open-set fault diagnosis across unseen operating speeds, resulting in a 105% relative improvement in overall accuracy compared with the original DiFG-Net (95.85% vs. 46.68%).

Although the proposed DiFG-Net++ framework was evaluated using a plastic bearing fault dataset, the underlying principles are not limited to bearing applications and can be readily extended to other rotating machinery systems, including gearboxes, turbines, compressors, pumps, and electric motors. The key components of DiFG-Net++, physics-informed functional learning, speed-extrapolative augmentation, speed-invariant latent regularization, and healthy-state target-speed anchoring, are based on general physical relationships between machine dynamics, rotational speed, and fault characteristic frequencies. For example, gearbox faults are characterized by gear-mesh frequencies and sideband structures that scale with shaft speed, whereas turbine and compressor faults exhibit speed-dependent vibration signatures associated with blade-passing frequencies, imbalance, misalignment, and structural resonances. Like bearing fault characteristic orders, these diagnostic signatures can be represented using speed-normalized physical functionals that remain relatively invariant across operating speeds. Consequently, the PFML framework can be adapted to predict machinery-specific order-domain or physics-informed functionals, enabling the encoder to learn speed-robust representations for a wide range of rotating equipment. Furthermore, the healthy-state target-speed anchoring strategy is particularly attractive for industrial deployment because healthy operational data are routinely collected during commissioning and normal operation of most critical assets, whereas fault data are often scarce or unavailable. Therefore, the proposed framework provides a promising foundation for developing generalized open-set cross-speed fault diagnosis systems that can adapt to previously unseen operating conditions across diverse industrial applications.

Future work includes investigating the integration of equipment-specific physical knowledge, such as gear-mesh dynamics, blade-passing phenomena, and rotor-dynamic characteristics, into the PFML framework to further enhance the transferability of DiFG-Net++ across different machine types and operating environments.

5. CONCLUSION

This paper presented DiFG-Net++, a physics-informed self-supervised learning framework for open-set cross-speed fault diagnosis. The proposed approach extends the PFML-based DiFG-Net architecture by introducing three complementary speed-adaptation mechanisms: speed-extrapolative augmentation, speed-invariant latent regularization, and healthy-state target-speed anchoring. Unlike conventional transfer learning and domain adaptation methods that primarily focus on feature-distribution alignment, the proposed framework explicitly incorporates the physical relationship between rotational speed and fault characteristic frequencies into the representation-learning process. Furthermore, normalized order-domain functionals were introduced as physics-informed self-supervised targets to encourage the learning of speed-invariant and fault-sensitive latent representations.

The experimental results obtained from the plastic-bearing fault dataset demonstrated the effectiveness of the proposed framework under challenging open-set cross-speed diagnosis conditions. The baseline DiFG-Net achieved an overall validation accuracy of only 46.68%, indicating limited generalization to unseen operating speeds. By incorporating healthy-state target-speed anchoring, the overall accuracy increased to 70.85%, highlighting the value of utilizing healthy operational data collected at the target speed. The integration of speed-extrapolative augmentation and speed-invariant latent regularization further improved the overall accuracy to 82.8%. When all three speed-adaptation mechanisms were employed simultaneously, DiFG-Net++ achieved an overall validation accuracy of 95.85%, representing an absolute improvement of approximately 49 percentage points and a relative improvement exceeding 105% compared with the original DiFG-Net, demonstrating strong generalization capability to previously unseen operating speeds.

The results indicate that the three proposed mechanisms address different aspects of the cross-speed adaptation problem. Speed-extrapolative augmentation expands the range of operating conditions observed during pretraining, speed-invariant latent regularization promotes fault-consistent representations across speeds, and healthy-state target-speed anchoring provides a practical means of adapting the learned representations to previously unseen operating conditions using only healthy target-speed data. Their combined effect enables robust open-set cross-speed fault diagnosis without requiring fault samples from the target speed.

Although the case study in the paper focused on plastic-bearing fault diagnosis, the proposed framework is broadly applicable to other rotating machinery systems, including gearboxes, turbines, compressors, pumps, and electric motors. Future work will investigate the integration of additional physics-informed functionals derived from gear-mesh dynamics, blade-passing phenomena, and rotor-dynamic characteristics, as well as the application of DiFG-Net++ to more complex industrial systems involving multiple operating variables such as load, speed, and environmental conditions. The promising results of this study suggest that physics-informed functional learning is a powerful direction for developing generalized, practical PHM systems that can adapt to previously unseen operating environments.

REFERENCES

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A Simple Framework for Contrastive Learning of Visual Representations. Proceedings of the 37th International Conference on Machine Learning (ICML 2020). Virtual Conference, pp. 1597–1607.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT). Minneapolis, MN, USA, pp. 4171–4186. doi:10.18653/v1/N19-1423 doi: 10.1109/ACCESS.2025.3556957. doi:10.1016/j.ress.2024.109957

Eldele, E., Ragab, M., Chen, Z., Wu, M., Kwoh, C. K., Li, X., & Guan, C. (2021). Time-Series Representation Learning via Temporal and Contextual Contrasting. Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI 2021). Montreal, Canada, pp. 2352–2359. doi:10.24963/ijcai.2021/324.

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*, vol. 17, no. 59, pp. 1–35.

Grill, J. B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Pires, B., Guo, Z., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., & Valko, M. (2020). Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Vancouver, Canada, pp. 21271–21284.

He, D., & He, M. (2026). "Development of Autonomous PHM for Bearing Fault Diagnosis," *2026 IEEE Aerospace Conference*, Big Sky, MT, USA, 2026, pp. 1-13, doi: 10.1109/AERO66936.2026.11519902.

He, D., Li, R., & Zhu, J. (2013). "Plastic Bearing Fault Diagnosis Based on a Two-Step Data Mining Approach," *IEEE Transactions on Industrial Electronics*, vol. 60, no. 8, pp. 3429-3440, Aug. 2013, doi: 10.1109/TIE.2012.2192894.

He, K., Chen, X., Xie, S., Li, Y., Dollár, P., & Girshick, R. (2022). Masked Autoencoders Are Scalable Vision Learners. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). New Orleans, LA, USA, pp. 16000–16009. doi:10.1109/CVPR52688.2022.01553.

Mohsenvand, M. N., Izadi, M. R., & Maes, P. (2020). Contrastive Representation Learning for Electroencephalogram Classification. Proceedings of Machine Learning Research, vol. 136, pp. 238–253.

Pan, S. J., & Yang, Q. (2010). A Survey on Transfer Learning. IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359. doi:10.1109/TKDE.2009.191

Vaaras, E., Airaksinen, M., & Räsänen, O. (2025). "PFML: self-supervised learning of time-series data without representation collapse," *IEEE Access*, vol. 13, pp. 60233–60244, 2025.

Wang, M., & Deng, W. (2018). Deep Visual Domain Adaptation: A Survey. Neurocomputing, vol. 312, pp. 135–153. doi:10.1016/j.neucom.2018.05.083.

Wang, T., Chu, F., Han, Q., Kong, F., & Huang, X. (2024). Self-Supervised Learning for Intelligent Fault Diagnosis: Recent Advances, Challenges, and Future Directions. Reliability Engineering & System Safety, vol. 245, Article 109957.

Wang, T., Chu, F., Han, Q., Kong, F., & Huang, X. (2024). Self-Supervised Learning for Intelligent Fault Diagnosis: Recent Advances, Challenges, and Future Directions. Reliability Engineering & System Safety, vol. 245, Article 109957. doi:10.1016/j.ress.2024.109957.

Zhang, W., Li, X., Ma, H., Luo, Z., & Li, X. (2019). Deep Learning-Based Intelligent Fault Diagnosis Methods Toward Generalized and Transfer Diagnostics: A Review. Neurocomputing, vol. 360, pp. 215–234. doi:10.1016/j.neucom.2019.05.056.

Zhang, W., Li, X., Ma, H., Luo, Z., & Li, X. (2022). Self-Supervised Learning for Machinery Fault Diagnosis: A Review and New Perspectives. Mechanical Systems and Signal Processing, vol. 179, Article 109358. doi:10.1016/j.ymssp.2022.109358.

BIOGRAPHIES

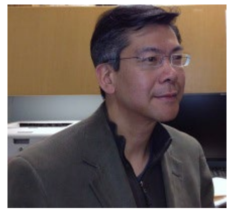

David He received a Ph.D. in Industrial Engineering from The University of Iowa. Dr. He is a Professor and Director of the Intelligent Systems Modeling & Development Laboratory in the Department of Mechanical and Industrial Engineering at The University of Illinois-Chicago. Dr. He is also a Fellow of the Prognostics and Health Management (PHM) Society. Dr. He's research areas include PHM, Industrial AI, smart manufacturing systems modeling and analysis, and quality and reliability engineering.

Miao He received a Ph.D. in Industrial Engineering and Operations Research from University of Illinois at Chicago, a B.E. degree in Vehicle Engineering from University of Science and Technology Beijing, Beijing, China and a M.S. degree in Mechanical Engineering from Purdue University, Hammond, IN. She is currently a senior research scientist at Siemens Corporation. Her research interests include digital signal processing, machinery health monitoring and fault diagnosis, prognostics, artificial neural network applications, and edge computing on factory automation.

Muhammad Areeb received the Bachelor of Engineering degree in Mechanical Engineering from NED University of Engineering and Technology, Karachi, Pakistan, in 2020, and the M.S. degree in Mechanical Engineering from the University of Illinois Chicago (UIC) in 2025. He is currently pursuing a Ph.D. in Industrial Engineering and Operations Research at UIC. He is a Graduate Research Assistant and Siemens Fellow in the Industrial AI and PHM Lab at UIC, where his research focuses on generative modeling for industrial audio and multimodal fault detection.