

# BearLLM: A Prior Knowledge-Enhanced Bearing Health Management Framework with Unified Vibration Signal Representation

Haotian Peng<sup>1, 2, 3, 4\*</sup>, Jiawei Liu<sup>1, 2, 3\*</sup>, Jinsong Du<sup>1, 2, 3</sup>, Jie Gao<sup>1, 2, 3†</sup>, Wei Wang<sup>1, 2, 3†</sup>,

<sup>1</sup>Shenyang Institute of Automation, Chinese Academy of Sciences, Shenyang, China

<sup>2</sup>Liaoning Liaohe Laboratory, Shenyang, China

<sup>3</sup>Key Laboratory on Intelligent Detection and Equipment Technology of Liaoning Province, Shenyang, China

<sup>4</sup>University of Chinese Academy of Sciences, Beijing, China

{penghaotian, liujiawei, jsdu, gaojie, wangwei2}@sia.cn

## Abstract

We propose a bearing health management framework leveraging large language models (BearLLM), a novel multimodal model that unifies multiple bearing-related tasks by processing user prompts and vibration signals. Specifically, we introduce a prior knowledge-enhanced unified vibration signal representation to handle various working conditions across multiple datasets. This involves adaptively sampling the vibration signals based on the sampling rate of the sensor, incorporating the frequency domain to unify input dimensions, and using a fault-free reference signal as an auxiliary input. To extract features from vibration signals, we first train a fault classification network, then convert and align the extracted features into word embedding, and finally concatenate these with text embedding as input to an LLM. To evaluate the performance of the proposed method, we constructed the first large-scale multimodal bearing health management (MBHM) dataset, including paired vibration signals and textual descriptions. With our unified vibration signal representation, BearLLM using one set of pre-trained weights achieves state-of-the-art performance on nine publicly available fault diagnosis benchmarks, outperforming specific methods designed for individual datasets. We provide a dataset, our model, and code to inspire future research on building more capable industrial multimodal models.

**Code** — <https://github.com/SIA-IDE/BearLLM>

## 1 Introduction

Bearings are the core components of mechanical rotating equipment but have high failure rates due to complex operational and environmental conditions (Wang et al. 2020b). Bearing health management (e.g., anomaly detection, fault diagnosis, and maintenance recommendations) is of great practical significance in industrial safety production to reduce economic losses and maintenance costs (Peng et al. 2022; Xiao et al. 2022; Ruan et al. 2023).

Current bearing health management frameworks rely on designing specialized methods for different working conditions and tasks, as shown in Fig. 1 (a). To apply specific methods to complex real-world industrial scenarios, domain

\*These authors contributed equally.

†Corresponding authors.

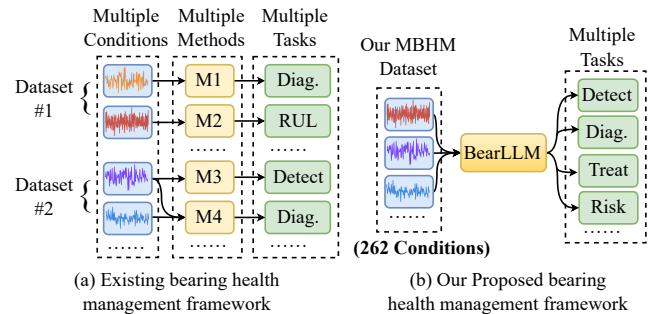


Figure 1: Comparison of existing bearing health management frameworks (Chaleshtori and Aghaie 2024; Ni et al. 2024) with our proposed approach. Our BearLLM replaces the complex operations of designing methods tailored to different conditions and tasks.

adaptation, and generalization have attracted widespread attention. Domain adaptation enables a model trained on one source domain to perform well on different but related target domains by reducing the domain shift or discrepancy (Wu et al. 2022; Zhang et al. 2022), but it suffers from low accuracy when the source and target domains are category-inconsistent (e.g., transitioning from working condition  $C_1$  with four fault types to  $C_2$  with five types). Domain generalization aims to extract domain-invariant features to improve performance on unseen domain (Li et al. 2020; Zheng et al. 2021; Chen et al. 2022a), but it is often constrained to a limited number of working conditions with small differences, e.g., fewer than ten working conditions in (Chen et al. 2022b; Lin et al. 2023). These purely data-driven methods often fail to strike an optimal balance between high accuracy and strong generalization for fault diagnosis.

In this paper, we propose a prior knowledge-enhanced bearing large language model (BearLLM), which can unify multiple bearing health management tasks over hundreds of different working conditions from multiple datasets, as shown in Fig. 1 (b). To handle various working conditions, we introduce a prior knowledge-enhanced unified vibration signal representation. Unlike most fault diagnosis methods that use fixed-length input segments, we sample vibration signals as variable-length but fixed-duration segments. These duration-consistent segments are then converted to the

frequency domain and are aligned. We further utilize a fault-free reference signal as a prior input, eliminating the need for complex mechanism analysis for various bearing designs (Zheng et al. 2021).

Specifically, we first design a fault classification network (FCN) to extract fault features based on the differences in frequency components between the query signal segment and the fault-free reference signal segment. This new frequency-based feature extraction paradigm for bearing fault diagnosis is more efficient (i.e., faster convergence and higher accuracy) and achieves stronger generalization, compared to previous methods that extract fault features directly from vibration signals. The extracted features are then transformed and aligned into word embedding, which is subsequently connected to user text embedding as inputs to the LLM. To evaluate the performance of the proposed method, we construct the first large-scale multimodal bearing health management (MBHM) dataset, including paired vibration signals and textual descriptions. Although the vibration signals from the nine public datasets differ significantly in distribution, BearLLM with a set of pre-trained weights achieves state-of-the-art performance using a unified vibration signal representation, outperforming specialized methods designed for individual datasets. The contributions of this paper are summarized as follows:

- We propose a novel bearing multimodal large language model, unifying multiple bearing health management tasks by aligning vibration signals and textual prompts.
- We propose a prior knowledge-enhanced unified vibration signal representation to handle various working conditions from multiple datasets.
- We construct the first large-scale multimodal dataset for bearing health management (MBHM), involving vibration signals with associated textual descriptions.
- Experimental results show that our BearLLM outperforms state-of-the-art fault diagnosis methods on nine publicly available benchmarks.

## 2 Related Works

**Multiple Working Condition:** Fault diagnosis under various working conditions from multiple datasets presents a challenge due to the heterogeneity of collected signals arising from variations in the test rigs, sensors, and environment, making it difficult to obtain unified features (Wen, Guo, and Li 2023). Existing domain adaptation methods (Chen, Gryllias, and Li 2020; Wan et al. 2022; Ma et al. 2023; Huo et al. 2023) typically involves training a model under known working conditions (source domain) and subsequently transferring knowledge to an unknown working condition (target domain). However, these approaches still necessitate individual transfer fine-tuning for each working condition in practice, hindering its ability to generalize across multiple scenarios. Domain generalization methods leverage training on multiple working conditions and aim to align the feature distributions of different domains through the design of network architectures and loss functions (Jia et al. 2023; Zhu, Chen, and Tang 2023; Hou et al. 2023b). However, these

approaches often rely on complex data preprocessing and augmentation techniques to help models learn fault features from vibration signals.

**Multiple Tasks:** Data-driven machinery health management have gained significant traction (Trabelsi et al. 2020). The concept of health management usually involves multiple tasks (Omri et al. 2021; Zio 2022), including anomaly detection, fault diagnosis, degradation prediction, maintenance decision-making, etc. LLMs such as ChatGPT-4 (OpenAI et al. 2024) have demonstrated exceptional capabilities across a wide range of tasks. The emergence of open-source foundational models like LLaMA 3 (Meta 2024) and Qwen 2 (Bai et al. 2023) have further empowered researchers in various disciplines to integrate these models into their own applications. In the aviation domain, Liu et al. (2024) applied generalized linear models to achieve multiple tasks, including assembly guidance and assembly error identification for aircraft engines. In the petroleum industry, Eckroth et al. (2023) designed a question-answering system based on LLM and knowledge graph, enabling retrieval of functionalities such as stratigraphy data and geological age determination. However, research integrating multiple tasks using LLMs for bearing health management remains limited (Li, Wang, and Sun 2024).

## 3 A Multimodal Bearing Health Management Dataset

Although several bearing-related datasets in Tab. 1 are available, they generally collect vibration signals on a single test rig, have a limited number of working conditions, and have no corresponding textual descriptions for training LLM. We have constructed a large-scale publicly multimodal dataset for bearing health management (MBHM).

Dataset	Sample Rate (kHz)	Conditions*	Fault Types	Time (s)	Text
CWRU	12 / 48	12	<b>10</b>	3932	×
DIRG	51.2	102	7	7140	×
HIT	20	40	3	9648	×
IMS	20	16	7	46480	×
JNU	100	45	4	3600	×
JUST	50	36	4	43986	×
MFPT	48.8 / 97.6	1	3	78	×
PU	64	4	5	7316	×
XJTU	25.6	6	<b>10</b>	13336	×
<b>MBHM</b>	<b>12 ~100</b>	<b>262</b>	<b>10</b>	<b>135516</b>	✓

\* The same working conditions represent the same load, speed, and sensor.

Table 1: Comparison of different datasets. Our MBHM dataset has the largest number of working conditions, the most complete fault types, and the longest time, paired textual prompts/responses.

The MBHM contains 135,516 pairs of vibration signal segments and fault types, and 542,064 pairs of text cues and responses, of which each sample is shown in Fig. 3, contains a vibration signal, a fault label, an operating condition id, a user prompt, and a text response, ie,  $(X_v, L_v, C, X_t, L_t) \in \text{MBHM}$ . Our dataset contains 262 working conditions collected from nine publicly accessible datasets, i.e., CWRU

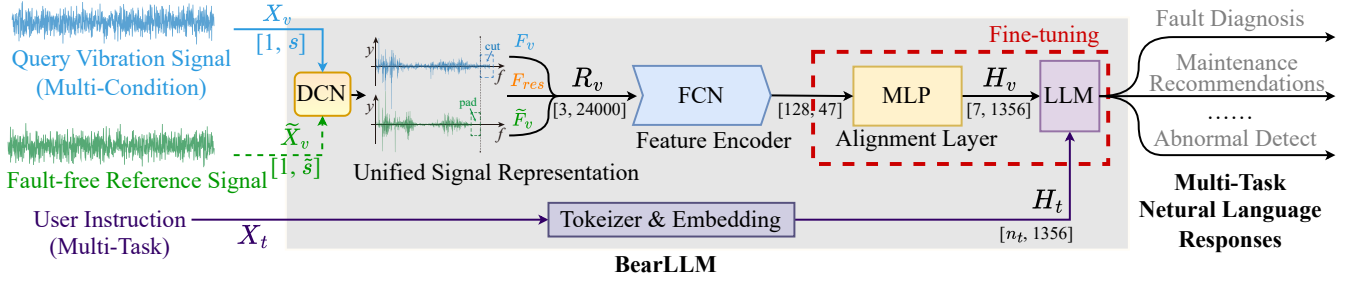


Figure 2: Architecture of our proposed BearLLM. Given a query vibration signal segment  $X_v$  and user instruction  $X_t$  as input, the model retrieves a fault-free vibration signal segment  $\tilde{X}_v$  with similar working conditions from the database as a reference. Two vibration signals are converted into a unified representation through DCN. A feature encoder identifies fault-related residuals between the two signals. The alignment layer converts these features into the word embedding  $H_v$ . Finally, an LLM is utilized with the user text embedding  $H_t$  to generate multi-task natural language responses, where  $n_t$  represents the length of the encoded text embedding.


Vibration Signal $X_v$	
Fault Label $L_v$	3 (severe inner ring)
Working Condition $C$	12 (DIRG, 100rpm, 500N, sensor#1)
Prompt Text $X_t$	... provide targeted <b>maintenance recommendations</b> based on the fault location and severity.
Response Text $L_t$	The bearing is in a <b>severe faulty state of inner ring</b> fault. <b>Maintenance recommendation: Replace</b> the bearing immediately to prevent further damage ...

Figure 3: Sample case of our MBHM dataset, includes vibration signal  $X_v$ , fault label  $L_v$ , working condition  $C$ , the specific task prompt text  $X_t$ , and the response text  $L_t$ .

(Case Western Reserve University 2008), DIRG (Daga et al. 2019), HIT (Hou et al. 2023a), IMS (Qiu et al. 2006), JNU (Jiangnan University 2012), JUST (Ren 2023), MFPT (Eric 2012), PU (Lessmeier et al. 2016), XJTU (Wang et al. 2020a). For each vibration signal, we have four different tasks, i.e., anomaly detection, fault diagnosis, maintenance recommendations, and potential risk analysis by generating text responses using ChatGPT (OpenAI et al. 2024). Detailed methodologies for dataset construction are provided in Appendix A.3. Our MBHM dataset contains the following features:

- **Multi-modal:** Each vibration signal is paired with four text prompts and responses, supporting the training and development of multimodal multi-task models.
- **Multiple working conditions:** Our dataset covers a wider range of working conditions, more accurately modeling real-world industrial production scenarios.

## 4 Method

In this section, we propose BearLLM, a novel multimodal model that unifies multiple bearing-related tasks. To handle various working conditions across multiple datasets, we introduce a prior knowledge-enhanced unified vibration signal representation in Section 4.1. The unified vibration signal is fed to a fault classification network to extract features in Section 4.2. We convert and align the extracted features into word embedding, and finally concatenate these with text embedding as input to an LLM in Section 4.3.

### 4.1 Prior Knowledge-Enhanced Unified Vibration Signal Representation

BearLLM aims to manage multiple bearing-related tasks across hundreds of working conditions. The basis for this is to build a unified vibration signal representation, involving adaptively sampling the vibration signal segments based on the sensor sampling rate, incorporating the frequency domain to unify input dimensions, and using a fault-free reference signal to calculate residual as auxiliary inputs to improve data utilization efficiency.

**Adaptive Sampling** To monitor various mechanical devices across different working conditions and industrial scenarios, vibration sensors are deployed with varying designations and sampling rates. However, most fault diagnosis methods (Zhu, Chen, and Tang 2023; Dong et al. 2024) use fixed-length signal segments in the time domain as inputs, where the fault frequency components in the inputs deviate from their original intrinsic values and vary with the sampling rate, hindering accurate fault diagnosis. Instead of sampling fixed-length signal segments, we adaptively sample vibration signals as variable-length but fixed-duration segments using prior knowledge of the sensor sampling rate. We extract the  $m$ -th query signal segment  $X_v \in \mathbb{R}^{1 \times s}$  from the original signal  $X_o$  by

$$X_v = X_o[ms, (m+1)s], \quad (1)$$

where  $s$  denotes the sampling rate of the sensor and controls the length of the  $X_v$ .

**Frequency-Domain Input Alignment** After adaptive sampling, each query segment ( $X_v$ ) has an equal duration, and the frequencies of  $X_v$  are aligned. However, varying lengths of  $X_v$  (due to different sampling rates) result in different numbers of frequency components, making them unsuitable for input to the network. We design a discrete cosine normalization (DCN) that consists of converting the vibration signal to the frequency domain using the discrete cosine transform (DCT), unifying the number  $n_f$  of frequency components using a pad or cut, and standardizing the amplitude using the normalization  $\mathcal{N}$ . The normalized frequency

representation  $F_v \in \mathbb{R}^{1 \times n_f}$  is obtained by

$$F_v = \begin{cases} \mathcal{N}(\text{DCT}(X_v)[0, n_f]), & \text{if } s \geq n_f \\ \mathcal{N}(\text{DCT}(X_v) \cup [0]_{n_f-s}), & \text{if } s < n_f \end{cases} \quad (2)$$

Signals with sampling rates below  $n_f$  are zero-padded, while those exceeding  $n_f$  are cut. To balance computational resources and fault classification accuracy, we empirically set  $n_f = 24000$  (more detail in Tab. 3). To enhance training stability, the amplitude of the frequency sequence is normalized to  $[-1, 1]$ ,

$$\mathcal{N}(x) = \beta \frac{\sqrt{\bar{n}x}}{\|x\|_2}, \quad (3)$$

where  $\beta$  is a scaling factor and is set to 0.01 by statistically analyzing our MBHM dataset.

**Fault-Free Reference Signal** To eliminate distributional differences from different inputs under various operating conditions, we introduce fault-free signals as reference signals. 1) In practical use, the reference signal segment  $\tilde{X}_v$  can be collected and saved when the equipment is working properly, such as after factory acceptance or maintenance; 2) in training on the MBHM dataset,  $\tilde{X}_v$  is acquired by

$$\tilde{X}_v \sim \{X_v^* | (X_v^*, L_v^*, C^*, X_t^*, L_t^*) \in \text{MBHM}, L_v^* = 0, C^* = C\}. \quad (4)$$

This indicates that  $\tilde{X}_v$  is selected when a signal ( $X_v^*$ ) of our MBHM dataset is fault-free (i.e.,  $L_v^* = 0$ ) and has the same working conditions as  $X_v$  (i.e.,  $C^* = C$ ).

We combine the query frequency signal ( $F_v$ ), the fault-free frequency signal ( $\tilde{F}_v$ ), and the residual frequency signal ( $F_{res} = F_v - \tilde{F}_v$ ) as unified vibration signal representation,

$$R_v = [F_v, \tilde{F}_v, F_{res}]. \quad (5)$$

## 4.2 Feature Extraction

### Algorithm 1: Training Algorithm

**Input:**  $\theta_E, \theta_C, \theta_A, \theta_L$  the weights of feature encoder, linear classification layer, alignment layer, and LLM in our BearLLM;  $X_v, L_v, X_t, L_t$  the vibration signal, fault label, prompt text, and response text from MBHM dataset

**Output:** The optimal parameters of BearLLM  $\theta_E^*, \theta_A^*, \theta_L^*$

#### Step 1: Pre-training FCN

**for**  $e \leftarrow 1$  to 50 epoches **do**

$R_v \leftarrow X_v$   $\triangleleft$  get unified representation by Eq. 5

$P = \text{FCN}(R_v)$

$\theta_E^*, \theta_C^* \leftarrow \text{argmin}_{\theta_E, \theta_C} -\nabla_{\theta_E, \theta_C} (\text{CE}(P, L_v))$   $\triangleleft$  cross-entropy

**end for**

#### Step 2: Fine-tuning BearLLM

Init.  $\theta_A$  by Eq. 7

**for**  $e \leftarrow 1$  to 20 epoches **do**

$Y = \text{BearLLM}(X_t, X_v)$

$\theta_A^*, \theta_L^* \leftarrow \text{PEFT}(Y, L_t)$

**end for**

**return**  $\theta_E^*, \theta_A^*, \theta_L^*$

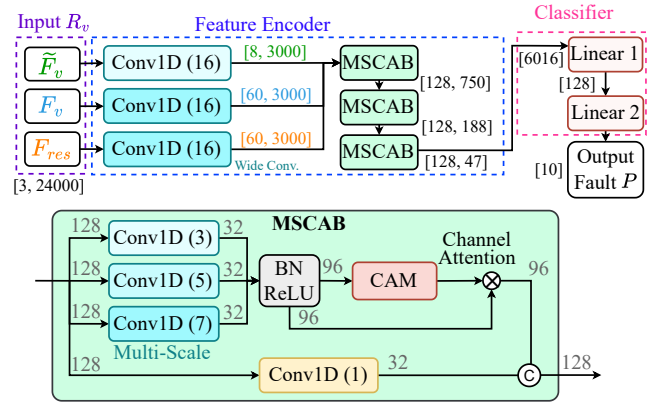


Figure 4: Structure of our proposed FCN. In the feature encoder, three wide convolutions are first used to extract main features, followed by three MSCAB blocks to transform and fuse multi-scale features for fault classification. The pre-trained FCN is used to initialize the feature extractor and alignment layer of BearLLM.

To extract the features of vibration signals, we propose a fault diagnosis network (FCN) containing a feature encoder parameterized by  $\theta_E$  and a linear classification layer parameterized by  $\theta_C$ , as shown in Fig. 4. We extract features from the unified vibration signal representation ( $R_v$ ) using three separate convolutional layers with large kernels (Zhang et al. 2017) and no weight sharing. We then transform features by three multiscale channel attention blocks (MSCAB) where the multiscale features are fused using the channel attention module (CAM) (Woo et al. 2018). Finally, we use two linear layers for fault classification.

Our FCN takes unified representation ( $R_v$ ) as input and outputs the fault type ( $P$ ). The shape of  $P$  is  $[1, \gamma]$  and  $\gamma$  denotes the number of fault types. We use cross-entropy loss for training with fault label  $L_v$  as ground truth. The training procedure is described in Algo. 1. The well-trained feature encoder weights ( $\theta_E^*$ ) of FCN are then used and frozen in BearLLM (see Fig. 2), while the classifier weights ( $\theta_C^*$ ) of FCN are used to initialize the alignment layer.

## 4.3 Feature Alignment

We propose a feature alignment layer to embed vibration features into word embedding, which is an MLP consisting of three linear layers (i.e.,  $l_1, l_2, l_3$ ). The weights of alignment layer is  $\theta_A = [\theta_C^*, \theta_{l_3}]$ , where  $\theta_C^*$  is the weights of  $l_1$  &  $l_2$  (i.e., two linear classification layers in FCN) and  $\theta_{l_3}$  is the weights of  $l_3$ . We use  $l_3$  transforms the output  $P$  of  $l_2$  into the word embedding  $H_v = \text{reshape}(l_3(P))$ , i.e.,

$$P \in \mathbb{R}^{1 \times \gamma} \xrightarrow{l_3} \mathbb{R}^{1 \times \tau h} \xrightarrow{\text{reshape}} H_v \in \mathbb{R}^{\tau \times h}, \quad (6)$$

where  $\tau$  signifies the token length after transformed,  $h$  is the hidden size of the LLM.

The weight  $\theta_{l_3}$  of  $l_3$  is initialized from the textual descriptions  $K$  of all fault categories by

$$K \in \mathbb{T}^{\gamma \times 1} \xrightarrow{\mathbf{T}} \mathbb{R}^{\gamma \times \tau} \xrightarrow{\mathbf{E}} \mathbb{R}^{\gamma \times \tau \times h} \xrightarrow{\text{reshape}} \theta_{l_3} \in \mathbb{R}^{\gamma \times \tau h}, \quad (7)$$

Method		CWRU	DIRG	HIT	IMS	JNU	JUST	MFPT	PU	XJTU	MBHM
Train and Test on Individual Datasets											
Methods for specific conditions	WDCNN	71.60	41.23	94.42	96.03	55.71	49.67	75.00	75.84	95.99	-
	TCNN	81.56	39.31	91.03	91.48	70.37	48.83	87.50	72.76	94.36	-
	QCNN	80.55	46.58	96.04	95.46	51.06	48.67	87.50	78.46	96.49	-
Train on MBHM Dataset and Test on Individual Datasets											
Methods for specific conditions	WDCNN	34.22	25.01	72.56	93.60	40.20	41.38	37.50	56.82	88.98	65.79
	TCNN	15.98	14.98	54.05	91.87	29.91	26.43	25.00	51.69	85.00	57.20
	QCNN	48.01	27.54	78.38	94.23	43.99	42.56	62.50	63.06	90.26	67.87
	WDCNN+DCN	80.35	68.76	96.51	96.39	90.87	77.78	62.50	84.86	95.66	87.60
	TCNN+DCN	48.40	28.19	68.60	92.41	46.98	48.19	12.50	68.49	90.96	68.98
	QCNN+DCN	95.55	91.32	99.41	97.36	98.96	85.91	100.00	92.32	98.27	87.60
Methods for cross conditions	MagNet	12.07	15.80	42.50	91.96	29.09	27.31	25.00	30.68	85.23	55.54
	BearingFM	41.18	14.29	33.33	92.18	48.39	58.78	33.33	54.22	88.45	75.18
	MagNet+DCN	95.31	79.94	89.78	88.79	94.95	65.85	87.50	74.43	95.85	81.18
	BearingFM+FCN	81.16	20.69	80.00	97.74	81.18	83.26	25.00	95.31	95.17	90.07
	<b>Ours</b>	<b>100.00</b>	<b>99.72</b>	<b>99.90</b>	<b>99.39</b>	<b>99.44</b>	<b>98.16</b>	<b>100.00</b>	<b>99.41</b>	<b>98.79</b>	<b>99.02</b>
		(+108%)	(+262%)	(+28%)	(+6%)	(+106%)	(+67%)	(+60%)	(+58%)	(+10%)	(+32%)

Table 2: Accuracy comparison with existing methods. “+DCN” denotes the addition of DCN to the original method, while “+FCN” indicates the replacement of the network of the original method with FCN, “(+108%)” represents a relative improvement from 48.01% to 100%. Our approach not only surpasses the SOTA accuracy on the MBHM dataset but also achieves results superior to those obtained from models trained specifically for individual datasets. The DCN and FCN components demonstrate broad applicability across diverse scenarios.

where  $\mathbb{T}$  stands for the text domain.  $\mathbf{E}$  and  $\mathbf{T}$  indicate the embedding layer and tokenizer of the pre-trained LLM, respectively. Using a tokenizer  $\mathbf{T}$  and an embedding layer  $\mathbf{E}$ , we generate a word embedding from  $K$ , which is then reshaped into the weight matrix  $\theta_{l_3}$ . See Appendix C.3 for more details on initializing weights.

We use the pre-trained Qwen2-1.5B (Bai et al. 2023) as our LLM parameterized by  $\theta_L$ , achieving basic human-computer interaction. However, its knowledge of specific domains and generation quality still requires improvement. We used the existing LoRA technique (Hu et al. 2022) and a general pipeline PEFT (Mangrulkar et al. 2022) for simultaneous fine-tuning of the LLM and our proposed alignment layer, which is detailed in Algo. 1.

## 5 Experiments

### 5.1 Experimental Setup

We implemented the proposed method using PyTorch (Paszke et al. 2019). Both pre-training and fine-tuning are performed on a single Nvidia RTX 4090 GPU. For pre-training, comparison trials, and ablation experiments, we used AdamW (Loshchilov and Hutter 2017) as the optimizer, and the batch size was set to 1024 for up to 50 epochs of training. Fine-tuning was performed using the existing PEFT (Mangrulkar et al. 2022) library.

To evaluate the effectiveness of our method, we provide quantitative comparison results for fault diagnosis, ablation of key components, and a user study to assess the quality of language responses. We addressed potential label leakage by dividing the 9 public datasets into a 7:2:1 ratio individually. The training set for the MBHM dataset consists of the concatenated training sets of these individual datasets, ensuring no overlap with their corresponding test sets. Other tasks including anomaly detection, maintenance recommendations, and potential risk analysis can be found in Appendix D.

### 5.2 Comparison With Fault Diagnosis Methods

We compared BearLLM with the following fault diagnosis methods. BearFM (Lai et al. 2024) and MagNet (Shi et al. 2023) are intended for diagnosing faults under cross-working conditions, while WDCNN (Zhang et al. 2017), TCNN (Chen, Gryllias, and Li 2020), and QCNN (Liao et al. 2023) are aimed at handling specific working conditions. Detailed descriptions of these methods can be found in Appendix B. To ensure a fair comparison, we re-implement these methods and test them under the same setup in section 5.1. The results are displayed in Tab. 2.




Figure 5: Accuracy and learning rate trends during training for different models. (a) Replacing the network of BearingFM with FCN resulted in increased accuracy and accelerated convergence. (b) Incorporating DCN into QCNN significantly mitigated overfit. (c) Our proposed method exhibits the fastest convergence and highest accuracy.

Our DCN achieves greater accuracy compared to BearingFM (Lai et al. 2024) when used with the same FCN (see

Fig. 5 (a)). The reason for this enhancement is likely due to BearingFM using absolute values after the FFT of the envelope spectrum. This method captures only the amplitude and ignores crucial phase information. In contrast, DCN leverages real-number computations, which help to reduce potential information loss, and operates in less than 20% of the time required by the comparison method. Combining DCN with MagNet (Shi et al. 2023) and utilizing aligned data for fusion augmentation has noticeably improved performance on datasets with substantial distribution differences.

Reflected in Tab. 2, the three methods (WDCNN, TCNN, QCNN) lacking data augmentation or alignment indicate strong accuracy on some specific datasets. However, their capacity to manage massive distribution differences is restricted when trained on the MBHM dataset. Including DCN eases the marked overfitting in QCNN (Liao et al. 2023), leading to a substantial improvement in validation accuracy (see Fig. 5 (b)). Similarly, adding DCN to both WDCNN (Zhang et al. 2017) and TCNN (Chen, Gryllias, and Li 2020) led to higher accuracy. Among all the methods tested, our proposed method achieves the highest accuracy and converges the fastest (within 20 epochs on the MBHM dataset as shown in Fig. 5 (c)).

### 5.3 Ablation Experiments And Generalization

$n_f$	Param.	FLOP	Accuracy
6,000	0.4013M	0.0177G	97.70%
12,000	0.5979M	0.0353G	98.32%
24,000	0.9747M	0.0704G	99.02%
48,000	1.7448M	0.1408G	99.20%

Table 3: Comparison of the number of parameters and FLOP of FCN under different  $n_f$  settings as well as the accuracy on the MBHM dataset.

The tests were carried out using four different  $n_f$  settings in DCN (see Eq. 2), as depicted in Tab. 3. As vibration information primarily resides in the low-frequency range, cut is unlikely to significantly impact accuracy. By increasing the number of frequency components, the distortion due to cut can be minimized, which enhances the precision on the MBHM dataset; however, this also raises parameters and computation in FCN. To achieve a balance between accuracy and performance, we opt for 24,000 as the  $n_f$ .

Method	Accuracy	Generalization	
	MBHM	JUST	IMS
Ours	99.02%	90.22%	98.52%
w/o DCN	72.15%	37.00%	91.46%
w/o $\tilde{F}_v, F_{res}$	98.35%	87.52%	97.81%
w/o $F_{res}$	98.82%	87.83%	97.96%
w/o $\tilde{F}_v$	98.63%	90.34%	98.29%

Table 4: A comparison of accuracy and generalization for different ablation setups is presented.

Ablation studies were conducted to further validate the effectiveness of each component in our proposed method.

We evaluated the performance by directly using raw time-domain vibration signals (fixed-length segments) as input and removing fault-free channels and residual channels separately and together.

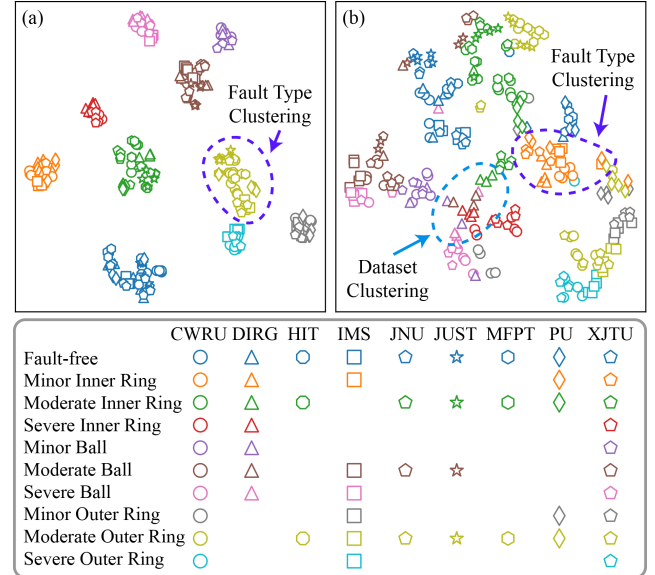


Figure 6: Visualization of output features with t-SNE. (a) Our method demonstrates clear inter-class separability. (b) Removing the fault-free and residual channels results in the signals from the same dataset exhibiting similar features.

Experimental results in Tab. 4 demonstrate significant accuracy and generalization drops when using time-domain signals only, further highlighting the efficacy of DCN. Applying the t-SNE, we compared the visualization of output with and without fault-free and residual channels. The blue box in Fig. 6 (b) shows how signal segments from the same dataset cluster closely in the feature space. This indicates that the model first identifies the dataset type before refining fault classification. Conversely, our proposed method, as shown in Fig. 6 (a), reduces inter-dataset differences. The model targets the residual between the query signal segments and the fault-free signal segments, creating a unified feature representation across the varying working conditions, and improving the generalization.

We evaluate the generalization ability of our proposed method using zero-shot settings. Among the publicly available datasets employed, JUST (Ren 2023) and IMS (Qiu et al. 2006) are the largest. We trained on the MBHM(w/o JUST&IMS) dataset, comprising only 35% of the MBHM training data, and performed zero-shot tests on the JUST and IMS datasets separately. On the JUST dataset, our method achieves an accuracy of 90.22% without any fine-tuning. In contrast, the method without fault-free and residual channels achieves an accuracy of only 87.54%.

Fig. 7 (a,b) illustrates a comparison of confusion matrices for zero-shot testing on the IMS dataset (Qiu et al. 2006), with and without fault-free and residual channels. Given that the IMS dataset is unbalanced (most samples are fault-free), the overall accuracy drops slightly from 98.52% to 97.81%.

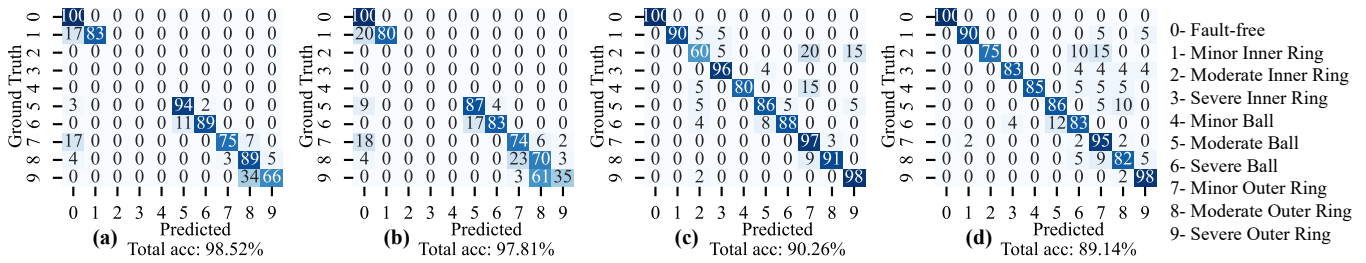


Figure 7: Confusion matrices for zero-shot performance in various scenarios. (a) Our method trained on MBHM(w/o IMS&JUST) and tested on IMS shows relatively reliable accuracy. (b) The method without fault-free and residual channels trained on MBHM(w/o IMS&JUST) and tested on IMS displays lower accuracy and a tendency to underestimate severity. (c) Our method trained on MBHM(w/o CWRU) and tested on CWRU confirms generalization. (d) Our method trained on MBHM(w/o CWRU&XJTU) and tested on CWRU further verifies the generalization of the unified representation.

However, the method without two auxiliary channels tends to grossly underestimate the severity. For example, 61% of severe outer ring faults are classified as moderate, and 23% of moderate outer ring faults are identified as minor.

The CWRU (Case Western Reserve University 2008) and XJTU (Wang et al. 2020a) datasets are the only ones that include all ten types of faults. To confirm the potential to create a unified representation, we trained our model on the MBHM(w/o CWRU) and MBHM(w/o CWRU&XJTU) datasets, respectively. We then performed zero-shot testing on the commonly used CWRU dataset, with the results of the confusion matrices displayed in Fig. 7 (c,d). Our method achieves remarkable accuracies of 90.26% and 89.14% on the untrained CWRU dataset for each setting, respectively. This result is even better than some methods trained on CWRU, which shows the generalization of our unified representation method and does not depend on any specific complete dataset for training.

#### 5.4 User Study

Task	A	B	C	D
FCN	9%	10%	-	-
w/o fine-tune	44%	32%	46%	45%
BearLLM	<b>47%</b>	<b>58%</b>	<b>54%</b>	<b>55%</b>

Table 5: Voting results from user study. Tasks A-D corresponds to anomaly detection, fault diagnosis, maintenance recommendations, and potential risk analysis. The fine-tuned BearLLM was the most favored across all tasks.

Tab. 5 summarizes the outcomes of four different tasks, with users choosing the best outputs from FCN, untuned BearLLM, and fine-tuned BearLLM in blind trials. Notably, in simpler tasks, few users chose the fault code output, while most preferred the natural language output. Fig. 8 illustrates examples of outputs before and after fine-tuning. Appendix D provides further comparisons for various tasks. Fine-tuning did not significantly affect the output of the simple anomaly detection task. In the fault diagnosis task, the model without fine-tuning sometimes missed information on fault severity, an issue that was resolved with fine-tuning. For the two more complex tasks, the fine-tuned model produced more accurate and detailed responses. Our method ad-

dresses the challenge faced by non-experts in utilizing maintenance systems due to their complexity.

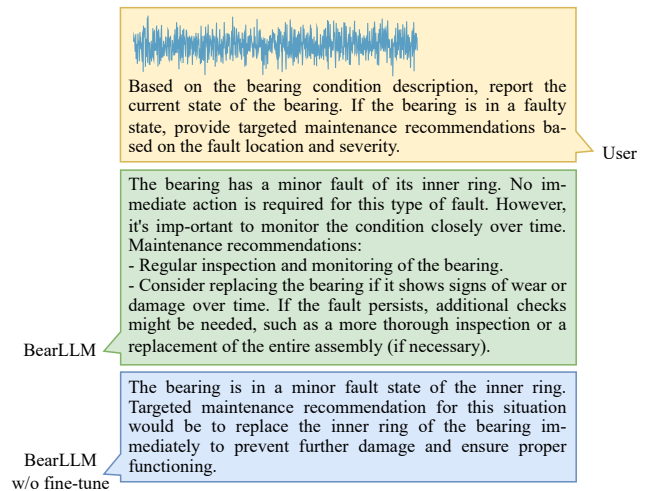


Figure 8: Examples of inputs and outputs of BearLLM. Vibration signals and task requirements are provided as user input, resulting in relevant natural language text output. The fine-tuned BearLLM exhibits improved response quality.

## 6 Conclusion

We propose BearLLM, a novel multimodal bearing health management framework that is the first attempt to unify multiple bearing-related tasks using LLMs, including anomaly detection, fault diagnosis, maintenance recommendations, and potential risk analysis. To build this unified framework, we introduce a prior knowledge-enhanced vibration signal representation for hundreds of different working conditions and construct the first large-scale multimodal bearing health management (MBHM) dataset. Experimental results on nine public fault diagnosis datasets show that BearLLM outperforms state-of-the-art methods, even surpassing those specifically trained on individual datasets. In addition, our frequency domain input alignment and feature extraction modules are plug-and-play, significantly improving the performance of other fault diagnosis models. We hope our work can inspire future research on building more capable industrial multimodal models.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (62073312); Applied Basic Research Program of Liaoning Province (2023JH2/101300228, 2023JH2/101300143); Natural Science Foundation of Liaoning Province (2022-MS-033). The funders have no role in study design, data collection, data analysis, data interpretation, or writing of the manuscript.

## References

- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; and Dang, K. 2023. Qwen Technical Report. *arXiv preprint arXiv:2309.16609*.
- Case Western Reserve University. 2008. CWRU Bearing Dataset. <https://engineering.case.edu/bearingdatacenter/download-data-file>. Accessed: 2024-12-26.
- Chaleshtori, A. E.; and Aghaie, A. 2024. A Novel Bearing Fault Diagnosis Approach Using the Gaussian Mixture Model and the Weighted Principal Component Analysis. *Reliability Engineering & System Safety*, 242: 109720.
- Chen, L.; Li, Q.; Shen, C.; Zhu, J.; Wang, D.; and Xia, M. 2022a. Adversarial Domain-Invariant Generalization: A Generic Domain-Regressive Framework for Bearing Fault Diagnosis Under Unseen Conditions. *IEEE Transactions on Industrial Informatics*, 18(3): 1790–1800.
- Chen, L.; Li, Q.; Shen, C.; Zhu, J.; Wang, D.; and Xia, M. 2022b. Adversarial Domain-Invariant Generalization: A Generic Domain-Regressive Framework for Bearing Fault Diagnosis Under Unseen Conditions. *IEEE Transactions on Industrial Informatics*, 18(3): 1790–1800.
- Chen, Z.; Gryllias, K.; and Li, W. 2020. Intelligent Fault Diagnosis for Rotary Machinery Using Transferable Convolutional Neural Network. *IEEE Transactions on Industrial Informatics*, 16(1): 339–349.
- Daga, A. P.; Fasana, A.; Marchesiello, S.; and Garibaldi, L. 2019. The Politecnico Di Torino Rolling Bearing Test Rig: Description and Analysis of Open Access Data. *Mechanical Systems and Signal Processing*, 120: 252–273.
- Dong, Y.; Jiang, H.; Yao, R.; Mu, M.; and Yang, Q. 2024. Rolling Bearing Intelligent Fault Diagnosis towards Variable Speed and Imbalanced Samples Using Multiscale Dynamic Supervised Contrast Learning. *Reliability Engineering & System Safety*, 243: 109805.
- Eckroth, J.; Gipson, M.; Boden, J.; Hough, L.; Elliott, J.; and Quintana, J. 2023. Answering Natural Language Questions with OpenAI’s GPT in the Petroleum Industry. In *SPE Annual Technical Conference and Exhibition*.
- Eric, B. 2012. MFPT Bearing Dataset. <https://www.mfpt.org/fault-data-sets>. Accessed: 2024-12-26.
- Hou, L.; Yi, H.; Jin, Y.; Gui, M.; Sui, L.; Zhang, J.; and Chen, Y. 2023a. Inter-Shaft Bearing Fault Diagnosis Based on Aero-Engine System: A Benchmarking Dataset Study. *Journal of Dynamics, Monitoring and Diagnostics*.
- Hou, Y.; Wang, J.; Chen, Z.; Ma, J.; and Li, T. 2023b. Diagnosisformer: An Efficient Rolling Bearing Fault Diagnosis Method Based on Improved Transformer. *Engineering Applications of Artificial Intelligence*, 124: 106507.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Huo, C.; Jiang, Q.; Shen, Y.; Zhu, Q.; and Zhang, Q. 2023. Enhanced Transfer Learning Method for Rolling Bearing Fault Diagnosis Based on Linear Superposition Network. *Engineering Applications of Artificial Intelligence*, 121: 105970.
- Jia, S.; Li, Y.; Wang, X.; Sun, D.; and Deng, Z. 2023. Deep Causal Factorization Network: A Novel Domain Generalization Method for Cross-Machine Bearing Fault Diagnosis. *Mechanical Systems and Signal Processing*, 192: 110228.
- Jiangnan University. 2012. JNU Bearing Dataset. <https://github.com/ClarkGableWang/JNU-Bearing-Dataset>. Accessed: 2024-12-26.
- Lai, Z.; Yang, C.; Lan, S.; Wang, L.; Shen, W.; and Zhu, L. 2024. BearingFM: Towards a Foundation Model for Bearing Fault Diagnosis by Domain Knowledge and Contrastive Learning. *International Journal of Production Economics*, 109319.
- Lessmeier, C.; Kimotho, J. K.; Zimmer, D.; and Sextro, W. 2016. Condition Monitoring of Bearing Damage in Electromechanical Drive Systems by Using Motor Current Signals of Electric Motors: A Benchmark Data Set for Data-Driven Classification. *PHM Society European Conference*, 3(1).
- Li, X.; Zhang, W.; Ma, H.; Luo, Z.; and Li, X. 2020. Domain Generalization in Rotating Machinery Fault Diagnostics Using Deep Neural Networks. *Neurocomputing*, 403: 409–420.
- Li, Y.-F.; Wang, H.; and Sun, M. 2024. ChatGPT-like Large-Scale Foundation Models for Prognostics and Health Management: A Survey and Roadmaps. *Reliability Engineering & System Safety*, 243: 109850.
- Liao, J.-X.; Dong, H.-C.; Sun, Z.-Q.; Sun, J.; Zhang, S.; and Fan, F.-L. 2023. Attention-Embedded Quadratic Network (Qtention) for Effective and Interpretable Bearing Fault Diagnosis. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–13.
- Lin, J.; Shao, H.; Zhou, X.; Cai, B.; and Liu, B. 2023. Generalized MAML for Few-Shot Cross-Domain Fault Diagnosis of Bearing Driven by Heterogeneous Signals. *Expert Systems with Applications*, 230: 120696.
- Liu, P.; Qian, L.; Zhao, X.; and Tao, B. 2024. Joint Knowledge Graph and Large Language Model for Fault Diagnosis and Its Application in Aviation Assembly. *IEEE Transactions on Industrial Informatics*, 20(6): 8160–8169.
- Loshchilov, I.; and Hutter, F. 2017. Fixing Weight Decay Regularization in Adam. *CoRR*, abs/1711.05101.
- Ma, L.; Jiang, B.; Xiao, L.; and Lu, N. 2023. Digital Twin-Assisted Enhanced Meta-Transfer Learning for Rolling Bearing Fault Diagnosis. *Mechanical Systems and Signal Processing*, 200: 110490.
- Mangrulkar, S.; Gugger, S.; Debut, L.; Belkada, Y.; Paul, S.; and Bossan, B. 2022. PEFT: State-of-the-art Parameter-Efficient Fine-Tuning methods. <https://github.com/huggingface/peft>. Accessed: 2024-12-26.

- Meta. 2024. Llama 3 Model Card. [https://github.com/meta-llama/llama3/blob/main/MODEL\\_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md). Accessed: 2024-12-26.
- Ni, Q.; Ji, J. C.; Feng, K.; Zhang, Y.; Lin, D.; and Zheng, J. 2024. Data-Driven Bearing Health Management Using a Novel Multi-Scale Fused Feature and Gated Recurrent Unit. *Reliability Engineering & System Safety*, 242: 109753.
- Omri, N.; Al Masry, Z.; Mairrot, N.; Giampiccolo, S.; and Zerhouni, N. 2021. Towards an Adapted PHM Approach: Data Quality Requirements Methodology for Fault Detection Applications. *Computers in Industry*, 127: 103414.
- OpenAI; Achiam, J.; Adler, S.; Agarwal, S.; and Ahmad, L. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; and et.al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Peng, F.; Zheng, L.; Peng, Y.; Fang, C.; and Meng, X. 2022. Digital Twin for Rolling Bearings: A Review of Current Simulation and PHM Techniques. *Measurement*, 201: 111728.
- Qiu, H.; Lee, J.; Lin, J.; and Yu, G. 2006. Wavelet Filter-Based Weak Signature Detection Method and Its Application on Rolling Element Bearing Prognostics. *Journal of Sound and Vibration*, 289(4): 1066–1090.
- Ren, X. 2023. JUST Slewing Bearing Datasets. *Mendeley Data*, 1.
- Ruan, D.; Wang, J.; Yan, J.; and Gühmann, C. 2023. CNN Parameter Design Based on Fault Signal Analysis and Its Application in Bearing Fault Diagnosis. *Advanced Engineering Informatics*, 55: 101877.
- Shi, Y.; Deng, A.; Deng, M.; Xu, M.; Liu, Y.; Ding, X.; and Bian, W. 2023. Domain Augmentation Generalization Network for Real-Time Fault Diagnosis under Unseen Working Conditions. *Reliability Engineering & System Safety*, 235: 109188.
- Trabelsi, I.; Zolghadri, M.; Zeddini, B.; Barkallah, M.; and Haddar, M. 2020. FMECA-Based Risk Assessment Approach for Proactive Obsolescence Management. In Nyffenegger, F.; Ríos, J.; Rivest, L.; and Bouras, A., eds., *Product Lifecycle Management Enabling Smart X*, 215–226. Cham: Springer International Publishing. ISBN 978-3-030-62807-9.
- Wan, L.; Li, Y.; Chen, K.; Gong, K.; and Li, C. 2022. A Novel Deep Convolution Multi-Adversarial Domain Adaptation Model for Rolling Bearing Fault Diagnosis. *Measurement*, 191: 110752.
- Wang, B.; Lei, Y.; Li, N.; and Li, N. 2020a. A Hybrid Prognostics Approach for Estimating Remaining Useful Life of Rolling Element Bearings. *IEEE Transactions on Reliability*, 69(1): 401–412.
- Wang, H.; Xu, J.; Yan, R.; and Gao, R. X. 2020b. A New Intelligent Bearing Fault Diagnosis Method Using SDP Representation and SE-CNN. *IEEE Transactions on Instrumentation and Measurement*, 69(5): 2377–2389.
- Wen, H.; Guo, W.; and Li, X. 2023. A Novel Deep Clustering Network Using Multi-Representation Autoencoder and Adversarial Learning for Large Cross-Domain Fault Diagnosis of Rolling Bearings. *Expert Systems with Applications*, 225: 120066.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–19.
- Wu, Y.; Zhao, R.; Ma, H.; He, Q.; Du, S.; and Wu, J. 2022. Adversarial Domain Adaptation Convolutional Neural Network for Intelligent Recognition of Bearing Faults. *Measurement*, 195: 111150.
- Xiao, Y.; Shao, H.; Han, S.; Huo, Z.; and Wan, J. 2022. Novel Joint Transfer Network for Unsupervised Bearing Fault Diagnosis From Simulation Domain to Experimental Domain. *IEEE/ASME Transactions on Mechatronics*, 27(6): 5254–5263.
- Zhang, W.; Peng, G.; Li, C.; Chen, Y.; and Zhang, Z. 2017. A New Deep Learning Model for Fault Diagnosis with Good Anti-Noise and Domain Adaptation Ability on Raw Vibration Signals. *Sensors*, 17(2): 425.
- Zhang, Y.; Ren, Z.; Zhou, S.; Feng, K.; Yu, K.; and Liu, Z. 2022. Supervised Contrastive Learning-Based Domain Adaptation Network for Intelligent Unsupervised Fault Diagnosis of Rolling Bearing. *IEEE/ASME Transactions on Mechatronics*, 27(6): 5371–5380.
- Zheng, H.; Yang, Y.; Yin, J.; Li, Y.; Wang, R.; and Xu, M. 2021. Deep Domain Generalization Combining A Priori Diagnosis Knowledge Toward Cross-Domain Fault Diagnosis of Rolling Bearing. *IEEE Transactions on Instrumentation and Measurement*, 70: 1–11.
- Zhu, Z.; Chen, G.; and Tang, G. 2023. Domain Adaptation With Multi-Adversarial Learning for Open-Set Cross-Domain Intelligent Bearing Fault Diagnosis. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–11.
- Zio, E. 2022. Prognostics and Health Management (PHM): Where Are We and Where Do We (Need to) Go in Theory and Practice. *Reliability Engineering & System Safety*, 218: 108119.