

# Edge AI - Enabled Smart Sensors for Predictive Maintenance of Marine Vessels

Igor Makienko<sup>1</sup>, Michael Grebshtein<sup>1</sup>, and Eli Gildish<sup>1</sup>

<sup>1</sup>*RSL Electronics, 70 Round Hill Road, Poughkeepsie, NY 12603, USA*  
*igor@rsl-electronics.com*

## ABSTRACT

Traditional Predictive Maintenance deployment on marine vessels requires continuous acquisition and transmission of large volumes of sensor data, resulting in high complexity, cost, and cybersecurity exposure that limit large-scale adoption.

This work presents RSL Smart Sensors, an Edge AI - enabled sensing platform that performs distributed on-sensor intelligence to significantly reduce data transmission, operational costs, and cyber risks. Sensor data are processed locally, with only health indicators and diagnostic insights transmitted to the cloud for decision support.

A novel machine learning algorithm is introduced to estimate machine operating conditions directly from vibration data under strict power constraints, without access to external operating parameters. Validation on marine auxiliary generator turbochargers demonstrates high accuracy and confirms the feasibility of high-performance predictive maintenance at the edge under limited power and information conditions.

## 1. INTRODUCTION

### 1.1. Predictive Maintenance of Maritime Vessels

Vibration-based condition monitoring is widely adopted for predictive maintenance of auxiliary equipment in the maritime industry due to its sensitivity to early-stage faults, as demonstrated by Kalafatelis et al. (2025), and the growing need to mitigate increasing vessel vibrations (Park et al., 2024).

However, both research and deployment challenges remain. From a research perspective, high accuracy in failure prognosis is required despite limited datasets, a small number of labeled failure cases, and varying operating conditions. From a deployment perspective, the required infrastructure is

costly for this industry, with significant communication and cloud-processing expenses associated with handling large volumes of data (Kalafatelis et al., 2025).

In recent years, the challenges of limited labeled data and varying operating conditions have been addressed through the integration of data-driven feature learning with conventional spectral analysis. For example, He et al. (2024) employed sparse autoencoders to extract latent representations from vibration signals for monitoring marine centrifugal pumps. Hu et al. (2025) applied deep transfer learning for turbocharger fault diagnosis by leveraging simulated data. Ouyang et al. (2024) introduced classification methods to improve fault discrimination under changing operating conditions of auxiliary diesel generators. Additionally, multimodal data fusion techniques have been proposed by Lv et al. (2025), combining vibration data with auxiliary sensor inputs to enhance diagnostic reliability in multi-machine environments affected by signal interference.

Despite these advances, deep learning-based approaches typically require expensive infrastructure for transmitting and processing raw vibration data within cloud-based architectures, which slows the adoption of predictive maintenance solutions. High-frequency vibration sampling generates large data volumes that are difficult to transmit over maritime satellite links and costly to store and process centrally (Liang et al., 2025). Cloud-centric frameworks further introduce substantial communication and computational overhead, compounded by latency and connectivity limitations (Chit et al., 2023).

To address these issues, recent studies have explored edge-based vibration monitoring as an alternative to cloud-based approaches in maritime applications. Edge computing enables real-time data reduction and local processing, making it well-suited for bandwidth-constrained environments (Michala et al., 2021). Comparative analyses of edge and cloud architectures indicate that edge processing can reduce both deployment and operational costs while improving system reliability in industrial IoT settings (Lu et al., 2023).

---

Igor Makienko et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

In this paper, a novel Edge AI-enabled sensing platform is introduced that eliminates the need for expensive supporting infrastructure and removes reliance on cloud-based computation. Compared to existing edge-processing solutions, typically transmitting large volumes of vibration data from sensors to local gateways, the proposed system significantly reduces deployment complexity and infrastructure costs by enabling fully autonomous, sensor-level intelligence.

The primary challenge in implementing such distributed intelligence lies in achieving high-accuracy failure prognosis under strict constraints, including extremely limited computational resources (to support years of maintenance-free operation) and incomplete knowledge of continuously varying operating conditions.

To address this, the paper presents a preprocessing method implemented directly on smart sensors and demonstrates its application to monitoring auxiliary diesel turbochargers. The proposed approach enables estimation of continuously varying machine rotating speed and performs blind stochastic-periodic signal separation, while maintaining minimal computational requirements.

## 1.2. Smart Sensor - Based Edge Processing

The system consists of two segments: onboard and onshore.

The onboard segment includes the following components:

- Smart sensors for recording and processing vibration and temperature data
- An onboard tablet for collecting data from smart sensors and transmitting it to the onshore cloud via the vessel's internet communication system

The onshore segment comprises the Vessel Predictive Maintenance Platform, which provides data visualization, analysis, and maintenance recommendations.

The smart sensors consist of a power source, sensing module, memory, data processing unit, and communication module. Data is processed locally on the sensors, and only health indicators and diagnostic insights are transmitted to the cloud for decision support. In this approach, the entire data analysis pipeline (including preprocessing, feature extraction, and anomaly detection) is executed directly on smart sensors installed on the monitored equipment. Only component health indicators and diagnostic insights are sent to the cloud, where AI-powered failure prognosis and maintenance recommendations are generated.

This architecture enables rapid deployment and cost-effective operation by eliminating the need for extensive data transmission, processing, and storage infrastructure.

The system data flow and top-level architecture are shown in Figure 1:



Figure 1. System data flow: Only health indicators are transmitted from the ship to the onshore Predictive Maintenance Platform via the vessel's internet connection.

The sensor module operates in two primary modes: initial data acquisition and operational mode. During the initial phase, data is collected for learning and calibration purposes, after which the sensor configures itself for regular operation. The transition to operational mode can also be triggered manually via a remote platform by transmitting the appropriate configuration from the onshore system to the tablet.

Examples of smart sensors, staff operations, and handheld tablet interfaces are shown in Figure 2.



Figure 2. An example of smart sensors, staff operations, and tablet usage.

The smart sensor is designed to provide real-time failure prognosis for the following components:

- Bearing outer and inner races, as well as rolling elements
- Gears, including both localized and distributed wear
- Shafts, including unbalance and misalignment
- Bolt loosening
- Vibration-related anomalies

To achieve this functionality, several research challenges must be addressed. This study focuses on the initial evaluation of the sensor's capabilities, including:

1. Automatic estimation of machine rotational speed based on vibration data
2. Vibration preprocessing using periodic–stochastic signal separation: the periodic component is used for gear and shaft fault prognosis, while the stochastic component is used for bearing and structural fault detection

A pilot field test was conducted on a turbocharger of an auxiliary power diesel generator aboard a cargo ship to evaluate these capabilities and address the identified challenges.

### 1.3. Problem Statement

Any observation of vibration signals can be modeled as a discrete-time stochastic process  $y[n]$  observed over a window  $N$ :

$$y[n] = x[n] + \eta[n], \quad n \in [0, M - 1] \quad (1)$$

Where:

- $x[n]$  is a deterministic periodic process:  $x[n] = x[n + P_0]$  for some period  $P_0$  (in samples).  $P = f_s/P_0$ , where  $f_s$  is the sampling frequency. The signal is composed of sinusoids with fundamental frequency  $F_0$ , its harmonics and other frequencies are not changing with  $F_0$ .
- In practice  $\eta[n]$  is an additive color noise such that  $\eta[n] = \sum_{j=1}^J b_j \varepsilon[n - j]$  and  $\varepsilon[n]$  is uncorrelated innovation process.

The objective is to estimate the deterministic periodic component and its fundamental frequency ( $\hat{x}[n], \hat{F}_0$ ) while minimizing computational complexity.

The estimated components are then used for feature extraction as follows: the machine’s operational state is characterized by the fundamental frequency  $F_0$ ; the periodic component  $\hat{x}[n]$  is utilized for shaft and gear prognosis; and the residual signal  $\eta[n]$  is utilized for bearing and structural health feature extraction.

### 1.4. Related Works

The estimation of  $F_0$  from vibrations and estimation of  $\hat{x}[n]$  constitute two distinct tasks, each requiring a separate literature review.

Peeteres et al. (2019) and Liu et al. (2023) provide a comprehensive review of existing methods for machine instantaneous speed estimation. These methods can be broadly categorized into three groups: spectrum-based

approaches, including the Short-Time Fourier Transform (STFT) and its variations for estimating high-speed fluctuations; cepstrum-based methods; and related hybrid techniques. These approaches estimate the fundamental frequency  $F_0$  to provide highly accurate input for angular resampling of vibration signals, thereby compensating for speed-related variability.

In contrast, smart sensors typically operate under limited power constraints and have limited prior information about expected speed variations. In such cases,  $F_0$  is primarily used for identifying machine operating modes or as a contextual parameter in the modeling of condition indicators. One of the earliest methods that does not require supplementary information is based on the Short-Time Scale Transform (STST), proposed by Combet et al. (2009) for relative instantaneous speed estimation. However, this method is computationally intensive and therefore unsuitable for our application. Gildish et al. (2024) proposed an alternative approach, estimating relative  $F_0$  is robust optimization applied to spectral peak statistics. This method requires minimal computation and appears well-suited for implementation in smart sensors.

Numerous methods for periodic–stochastic signal separation have been proposed, as reviewed by Tiboni et al. (2023) and Liu et al. (2023). These methods can be classified into time-domain, frequency-domain, cepstrum-based, and machine learning approaches. Among these, only methods that require minimal computational resources and do not rely on prior knowledge of  $F_0$  are relevant to our application.

Among the well-known blind approaches are Autoregressive (AR) modeling and Self-Adaptive Noise Cancellation (SANC), which predict periodic signal components using short-term historical data without requiring knowledge of  $F_0$  by Ho (1999). This approach was further improved by Peeters et al. (2025) using multi-delay prediction to reduce prediction variance. However, these methods require long, stationary vibration recordings to achieve high accuracy in the frequency domain, making them less suitable for scenarios with significantly varying operating conditions.

More recently, Gildish et al. (2024) and Makienko et al. (2024) addressed this limitation by employing multi-layer dilated convolutional neural networks (CNNs), which increase the receptive field without increasing the number of trainable parameters. Nevertheless, a key drawback of these methods is that model parameters must be re-estimated for each new observation  $y[n]$  making them impractical for deployment under strict power constraints.

This limitation was recently addressed by Gildish et al. (2026), who proposed reusing fixed model parameters across observations by rescaling the time domain according to the estimated  $F_0$ . The present paper adopts this method for estimating  $x[n]$  with minor modifications tailored for power-limited smart sensor applications.

## 1.5. Contribution

This paper presents the proposed smart sensor concept and a methodology that combines the vibration signal pre-processing approach of Gildish et al. (2026) with the rotational frequency estimation method of Gildish et al. (2023). The pre-processing approach assumes prior knowledge of the rotational frequency, while the estimation method estimates it directly from vibrations. As both methods contribute to the overall estimation uncertainty, their combined error must be quantified under real operating conditions of smart sensors. The methodology is validated through a pilot field experiment performed on the turbocharger of an auxiliary diesel generator installed aboard a cargo vessel. The results demonstrate the sensor's capability to accurately estimate machine rotational frequency and to perform periodic–stochastic signal separation without requiring external information, even under conditions of significant rotational speed variations and strict power constraints.

This paper is structured as follows. Section 2 describes the methods used for rotational frequency estimation from vibration signals and for periodic–stochastic signal separation. Section 3 presents the experimental results. Section 4 summarizes the conclusions and outlines directions for future work.

## 2. METHODS

The study focuses on two tasks: estimating machine rotational frequency from vibration signals and separating periodic and stochastic components. We assume that each vibration observation is sufficiently short such that the rotational frequency remains stable within a single observation.

### 2.1. Rotational Frequency Estimation from Vibrations

As demonstrated by Gildish et al. (2024), when the fundamental frequency  $F_0$  changes by scaling factor  $s$  relative to a reference value (e.g.  $F_0 = sF_{ref}$ ). Assume that in observations  $y_1[n]$  and  $y_2[n]$  the  $F_0$  was changed by scaling factors  $s_1$  and  $s_2$  respectively, then the observations can be presented as follows:

$$\begin{aligned} y_1[n] &= x[n \cdot s_1] + \eta_1[n], \\ y_2[n] &= x[n \cdot s_2] + \eta_2[n] \end{aligned} \quad (2)$$

Their corresponding power spectral densities (PSDs) are given by:

$$\begin{aligned} Y_2(f) &= \left| \frac{1}{s_1} X\left(\frac{f}{s_1}\right) \right|^2 + |\Theta_1(f)|^2, \\ Y_2(f) &= \left| \frac{1}{s_2} X\left(\frac{f}{s_2}\right) \right|^2 + |\Theta_2(f)|^2 \end{aligned} \quad (3)$$

By accumulating  $N$  observations, the scaling factors can be estimated using the quantiles of the spectral peak distributions and solving the following eigenvalue problem:

$$\begin{aligned} \mathbf{Z} &= \text{diag}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{A}, \\ \mathbf{Z} \mathbf{s} &= N \mathbf{s}, \\ \mathbf{A} &= (\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N) \end{aligned} \quad (4)$$

where  $\mathbf{f}_i = F_{M_i}^{-1}(\mathbf{p})$  denotes the inverse cumulative distribution function (CDF) of the spectral peak frequencies in signal  $y_i[n]$ , evaluated at a predefined probability vector  $\mathbf{p}$  of length  $P$ .

The matrix  $\mathbf{A}$  is formed from the vectors  $\mathbf{f}_i$  each having dimensions defined by  $P$  and  $N$ . As a result, matrix  $\mathbf{Z}$  consists exclusively of exact pairwise ratios. Therefore, its dominant eigenvector recovers, up to a scaling factor, the original absolute values that generated these ratios. The solution for the problem is given by the eigenvector  $\mathbf{s}_{max}$  corresponding to the largest eigenvalue  $\lambda_{max}$  of matrix  $\mathbf{Z}$ :

$$(\mathbf{Z} - \lambda_{max} \mathbf{I}) \mathbf{s}_{max} = 0 \quad (5)$$

The fundamental frequency across all observations is then estimated as  $F_0 = sF_{ref}$ , based on the estimated scaling factor  $s$  relative to the reference observation. After the training stage is completed, each new scaling factor  $s_{N+i}$  is easily estimated by updating the training matrix  $\mathbf{Z}$  to be  $N + 1 \times N + 1$  with an additional row and column and the scaling factor is equal to:

$$s_{N+i} = \frac{\mathbf{r}_{N+1}^T \{0, \dots, N\} \mathbf{s}_{max}}{N + 1 - d_{N+1}} \quad (6)$$

where  $\mathbf{r}_{N+1}^T \{0, \dots, N\}$  is the  $N + 1$  row of the updated matrix  $\mathbf{Z}$  excluding the last element, and  $d_{N+1}$  is its last element.

The algorithm's computational complexity is dominated by spectrum estimation, which requires  $O(M \log_2 M)$  where  $M$  is the signal length. The subsequent local peak detection and eigenvalue computation for estimating  $F_0$  require  $O(P)$  calculations, where  $P$  is the length of the vector  $\mathbf{p}$ . Since spectrum estimation is typically required for feature extraction, the overall computational cost remains low, making the method suitable for power-constrained applications.

### 2.2. Vibrations Pre-Processing: Periodic-Stochastic Separation

This section presents the proposed methodology, whereby a DCNN is trained exclusively on a reference observation and subsequently applied to observations with different fundamental frequencies. To ensure compatibility between the training and inference data, the time axis of each observation is resampled to align its fundamental frequency with that of the reference observation. As a result, the trained model can be reused across varying operating conditions

without requiring retraining. The method used in this paper was developed by Gildish et al. (2026). It performs separation of periodic and stochastic signal components by predicting samples in vibration measurements using a linear combination of past samples which is feasible only in periodic deterministic processes, in contrast to stochastic processes.

The estimation procedure consists of two stages: training and inference. During the training stage, a linear deep convolutional neural network (DCNN) is employed instead of traditional Time Synchronous Averaging (TSA), which requires measurements of the rotational frequency, and instead of autoregressive (AR) - based methods, which require very large window size to achieve high frequency resolution. The DCNN is trained using a single time-series observation from a given run and then kept fixed. During the inference stage, the trained model is applied to all observations, first resampled to be synchronized with the reference observation used during training.

For periodic signals, the optimal predictor in terms of mean squared error is a linear combination of past samples. Therefore, the predictor for the values of  $y[n]$  is equivalent to the estimator of the periodic signal  $\hat{x}[n]$ , as follows:

$$\hat{y}[n + \Delta] = \hat{x}[n + \Delta] = \sum_{k=1}^K w_k \cdot y[n - k] \quad (7)$$

The predictor  $\hat{x}[n]$  estimates the value of  $y[n]$  by exploiting its previous  $K$  samples, given a specific value of  $F_0$ . The time delay  $\Delta$  ensures that the stochastic component remains uncorrelated, as expressed by  $\mathbb{E}\{\eta[n]\eta[n - k]\} = 0$  for all  $k > \Delta$ .

The algorithm employs a linear one-dimensional DCNN to obtain an optimal linear predictor for the periodic signal, as shown in (7). The estimation is carried out such that each output sample  $y[n]$  is predicted using the preceding  $K$  samples  $y[n - 1], \dots, y[n - K]$ .

The network parameters are:

$$\Theta = \{W^{(1)}, \dots, W^{(L)}\}$$

where  $W^{(\ell)}$  are the convolution weights of layer  $\ell$ ,

The dilation of layer  $\ell$  is defined as:

$$d_\ell = 2^{\ell-1} \quad (8)$$

Let  $h_r^{(\ell)}[n]$  denote the  $r$ -th channel of layer  $\ell$ . The input is

$$h^{(0)}[n] = y[n - 1], \dots, y[n - K]$$

For a kernel of length  $K_\ell$ , the dilated convolution of layer  $\ell$  is calculated as follows:

$$h_r^{(\ell)}[n] = \sum_{c=1}^{C_{\ell-1}} \sum_{k=0}^{K_{\ell-1}} w_{rck}^{(\ell)} h_c^{(\ell-1)}[n - kd_\ell] \quad (9)$$

Where  $w_{rck}^{(\ell)}$  are weights of the  $\ell$  layer corresponding to channel  $r$ .

An example of a 3-layer DCNN is shown in Figure 3

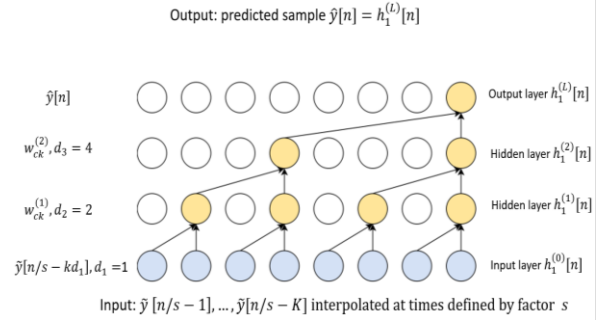


Figure 3. An example of a single-channel DCNN with two hidden layers and a scaling factor  $s$

When  $F_0$  is scaled by factor  $s$  the predictor changes as follows given the model in equation (2):

$$\hat{x}[n/s + \Delta] = \sum_{k=1}^K w_k \cdot \tilde{y}[n/s - k] \quad (10)$$

This implies that the model weights  $w_k$  remain unchanged when  $F_0$  varies. To account for this, the time axis of the input signal  $y[n]$  is scaled by a factor  $s$ , resulting in a resampled signal  $\tilde{y}[n/s]$ , interpolated at time instances  $n/s$ .

The DCNN model, whose parameters are learned during training from a single observation, can then be applied to denoise new observations using these fixed weights. However, before denoising, each new observation must undergo a resampling procedure using the scaling factor  $s$ , which is estimated by the algorithm described in Section 2.1.

The resampling process is implemented using upsampling ( $U$ ) and downsampling ( $L$ ), such that the ratio  $U/L$  equals the scaling factor  $s$ . This computation is performed using polyphase filtering, which reduces computational complexity by carrying out the filtering in the downsampled domain. The upsampling and downsampling factors must satisfy  $s = U/L$ , ensuring the required scaling condition:

$$\begin{aligned} \text{int}(U, L > 0), \text{argmin}(U + L), \\ \text{s.t. } |s - U/L| \leq \varepsilon \end{aligned} \quad (11)$$

In practice, the precision of  $s$  can be determined by fixing  $L$  and computing  $U = sL$ . For example, if  $L = 5$  is held constant, the precision of  $s$  is  $1/L = 0.2$ . Any value of  $s$  is then rounded to this precision; for instance,  $s = 1.61856$  becomes 1.6, and consequently  $U = sL = 1.6 \cdot 5 = 8$ . The

choice of precision is guided by the desired trade-off between accuracy and computational simplicity.

The re-sampling using polyphase is shown schematically in Figure 4:




Figure 4. Vibration re-sampling according to scale factor  $s = U/L$

### 2.3. Algorithm Summary

Initialization:

- Define training size  $N$ , Probability vector length  $P$ , downsampling factor  $L$ , DCNN architecture (depth and kernel size)

Training Stage:

- Collect  $N$  observations
- Estimate fundamental frequencies  $F_0$  for each observation (Section 2.1), forming  $N \times N$  matrix  $\mathbf{Z}$  and  $N \times 1$  vector  $\mathbf{s}_{max}$
- Select a representative observation and train the DCNN using the defined architecture

Inference Stage:

- For each new observation, estimate the scale factor  $s$  and  $F_0$  using Equation (6)
- Compute  $L$  and  $U$  from  $s$  and re-sample the signal using Equation (11).
- Estimate periodic component using the trained DCNN
- Proceed to feature extraction: periodic components for gears/shafts, stochastic – for bearings and structural faults

## 3. EXPERIMENTAL SETUP

### 3.1. Goals

- Evaluation of a rotational frequency estimation method based on smart sensors installed on auxiliary power diesel turbochargers, which presents a significant challenge due to extremely high rotational speeds and substantial frequency variations during operation.
- Assessment of a periodic-stochastic separation method utilizing the estimated machine frequency.

- Comparison of the proposed method with industry-standard periodic-stochastic separation techniques, including Time-Synchronous Averaging (TSA) and AR-based models, highlighting the associated trade-offs.

### 3.2. Installation Setup

The pilot study utilizes six months of operational data from a diesel-powered turbocharger installed on a commercial cargo ship. Vibration data were recorded periodically every two hours during operation using two parallel systems: a single smart sensor and a standard vibration acquisition system, which also interfaces with the turbocharger's rotational speed for reference. Both systems recorded vibration signals at a sampling frequency of 24 kHz. The rotational speed interrupts were measured differentially at a 50 MHz sampling rate to ensure high measurement accuracy.

The method was evaluated using a total of 1,250 recordings. Each recording consists of a 1-second vibration measurement, corresponding to 24,000 samples.

## 4. EXPERIMENT RESULTS

### 4.1. Vibrations Data Recordings

The spectrograms of all observations, along with the corresponding vibration RMS values, are shown in Figure 5. The measured rotational speed is shown in black on the spectrogram for reference. The rotational frequency changes rapidly during operation, making the case particularly challenging.

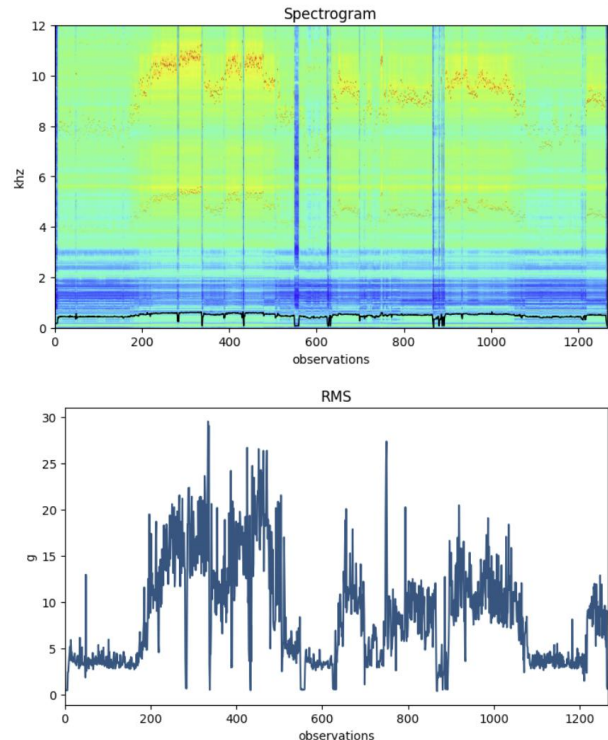


Figure 5. Vibration spectrogram (top) and vibration RMS (bottom) over half a year of turbocharger operation. The measured rotational frequency is indicated in black on the spectrogram.

#### 4.2. Rotational Frequency Estimation from Vibrations

The rotational frequency estimated using the method described in Section 2.1 was evaluated vs. actual rotational speed measured by the reference vibration recording system. The results are shown in Figure 6.




Figure 6. Actual rotational frequency compared to the algorithm's estimate (top), and the residual error between actual and estimated values (bottom) over half a year of turbocharger operation.

The error metrics ( $RMSE = 16.18$  Hz and  $R^2 = 0.97$ ) indicate high estimation accuracy, making the approach suitable for smart sensor applications such as machine state estimation and periodic-stochastic signal separation. The algorithm parameters were selected as follows: the number of observations used for training was  $N = 100$ , and the length of the probability vector defining the inverse CDF was  $P = 30$ . After training, a subset of inverse CDF vectors  $f_1, f_2, \dots, f_{100}$  is stored, and the rotational frequency in new observations is subsequently estimated by using Equation (6).

#### 4.3. Vibrations Pre-Processing: Periodic-Stochastic Separation

The vibration pre-processing using periodic-stochastic separation, described in Section 2.2 was evaluated in field tests against industry-standard methods: AR-based periodic signal modeling and Time-Synchronous Averaging (TSA), which relies on measured rotational frequency for cycle-synchronous averaging.

Both the AR approach and the proposed method operate in a similar manner, predicting input samples as a linear combination of past samples. However, the proposed method offers two key advantages. First, during training, it achieves comparable frequency resolution while using fewer optimization parameters. Second, during inference, it significantly reduces computational cost, as it does not require re-training of model parameters, unlike the AR approach.

The following algorithm settings were used:

##### AR:


- Window length: 300 samples

##### DCNN:

- Number of signals for training: 1
- Architecture: 5 layers, kernel size = 10 samples (50 parameters), providing a receptive field comparable to AR
- Re-sampling during inference:  $L = 5$ ,  $U = sL$
- The estimated  $F_0$  was used to compute the scaling factor  $s$

Both AR and DCNN were evaluated against the industry-standard TSA method, using actual rotational frequency measurements. Prediction errors were calculated by comparing the outputs of AR and DCNN to TSA, which served as a pseudo ground truth.

Qualitative results are presented in Figure 7 and Figure 8 for the AR and DCNN methods, respectively.



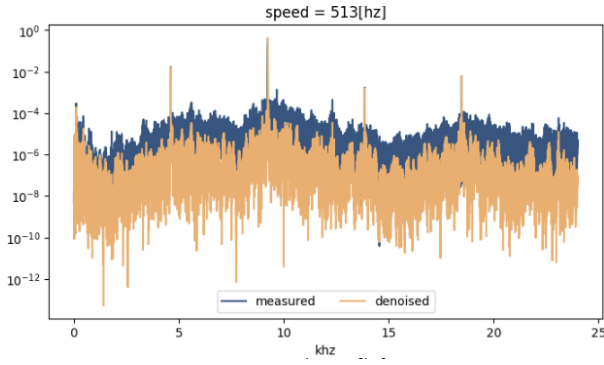


Figure 7. Example of AR-based periodic–stochastic separation: time-domain periodic signal estimation (top) and frequency-domain representation (bottom). The blue line represents the measured vibrations, and the orange line shows the estimated signal.




Figure 8. Example of DCCN-based periodic–stochastic separation: time-domain estimation of the periodic signal (top) and frequency-domain representation (bottom). The blue line represents the measured vibrations, and the orange line shows the estimated signal.

The comparison graph of percentage errors in signal norm for AR, DCNN, and TSA is presented below. The errors are sorted by rotational frequency. The results demonstrate high accuracy, with a maximum average error of approximately 7% of the signal energy. Both methods perform better as the rotational frequency increases, indicating improved prediction accuracy. This is because a larger number of periods fall within the prediction window at higher frequencies.

The AR method outperforms the DCNN for most frequencies. However, as noted above, the DCNN is more practical, since its computational complexity is significantly reduced during operation.




Figure 9. RMSE error between AR (blue), DCNN (orange) vs. TSA as industry standard sorted by rotational frequency.

### 5. CONCLUSION AND FUTURE WORK

The study presents a novel smart sensor–based concept for maritime predictive maintenance, with the potential to significantly reduce deployment, infrastructure, and cloud-related costs. This is achieved by enabling on-sensor processing of vibration data and transmitting only health indicators and alerts.

The results demonstrate the capability of the sensors to estimate machine rotational frequency from vibration signals and to perform periodic–stochastic vibration separation without prior knowledge of external parameters, while operating under strict power constraints.

Ease of deployment and cost-effective operation at scale are key advantages of the proposed approach. The method, along with pilot field test results, demonstrates the ability to perform high-accuracy preprocessing of vibration measurements within limited power budgets.

To further evaluate benefits of the proposed concept, future research should focus on developing methods for on-sensor feature extraction and failure prognosis. This would eliminate the need to transfer large volumes of raw data to the cloud and help lower barriers to the adoption of predictive maintenance in the maritime industry. Additionally, the evaluation dataset should be substantially expanded, and future assessments should include non-stationary operating conditions.

### REFERENCES

- Chit, T. W., Toro, C., Lim, H. C., & Muthu, R. (2023). Scalable remote cloud data center for vessel equipment predictive maintenance service-as-a-product (SaaS). *Procedia Computer Science*, 225, 2826-2833.
- Combet, F., & Zimroz, R. (2009). A new method for the estimation of the instantaneous speed relative fluctuation in a vibration signal based on the short time scale transform. *Mechanical Systems and Signal Processing*, 23(4), 1382-1397.
- Gildish, E., Grebshtein, M., Aperstein, Y., & Makienko, I. (2023, October). Vibration Signal Decomposition using Dilated CNN. In *Annual Conference of the PHM Society* (Vol. 15, No. 1).
- Gildish, E., Grebshtein, M., & Makienko, I. (2024). Vibration-based rotation speed estimation for Industrial IoT. *Internet of Things*, 25, 101024.
- Gildish, E., Grebshtein, M., & Makienko, I. (2026). Dilated CNNs for Periodic Signal Processing: A Low-Complexity Approach. *arXiv preprint arXiv: 2604.21651*
- He, Y., Yao, Y., & Ou, H. (2024). Status recognition of marine centrifugal pumps based on a stacked sparse auto-encoder. *Applied Sciences*, 14(4), 1371.
- Ho, Dominique. "Bearing diagnostics and self-adaptive noise cancellation." *PhD diss., UNSW Sydney*, 1999.
- Hu, L., Liu, L., Yang, J., Hu, H., Zheng, C., & Yu, Y. (2025). Fault diagnosis based on deep transfer learning for marine turbocharger. *International Journal of Mechanical Sciences*, 300, 110444.
- Kalafatelis, A. S., Nomikos, N., Giannopoulos, A., Alexandridis, G., Karditsa, A., & Trakadas, P. (2025). Towards predictive maintenance in the maritime industry: A component-based overview. *Journal of Marine Science and Engineering*, 13(3), 425.
- Liang, Q., Knutsen, K. E., Vanem, E., Aesoy, V., & Zhang, H. (2024). A review of maritime equipment prognostics health management from a classification society perspective. *Ocean Engineering*, 301, 117619.
- Liu, D., Cui, L., & Wang, H. (2023). Rotating machinery fault diagnosis under time-varying speeds: A review. *IEEE Sensors Journal*, 23(24), 29969-29990.
- Lu, S., Lu, J., An, K., Wang, X., & He, Q. (2023). Edge computing on IoT for machine signal processing and fault diagnosis: A review. *IEEE Internet of Things Journal*, 10(13), 11093-11116.
- Lv, Y., Hao, J., Tang, M., & Wu, J. (2025). Multimodal data fusion-based intelligent fault diagnosis for ship rotating machinery: Status quo and perspectives. *Engineering Applications of Artificial Intelligence*, 160, 111767.
- Makienko, I., Grebshtein, M., & Gildish, E. (2024). Estimating vibration sources for industrial IoT using dilated CNN and deconvolution. *Internet of Things*, 27, 101303.
- Michala, A. L., Vourganas, I., & Coraddu, A. (2021). Vibration edge computing in maritime IoT. *ACM Transactions on Internet of Things*, 3(1), 1-18.
- Ouyang, H., Li, W., Gao, F., Huang, K., & Xiao, P. (2024). Research on fault diagnosis of ship diesel generator system based on IVY-RF. *Energies*, 17(22), 5799.
- Park, M. H., Yeo, S., Choi, J. H., & Lee, W. J. (2024). Review of noise and vibration reduction technologies in marine machinery: Operational insights and engineering experience. *Applied Ocean Research*, 152, 104195.
- Peeters, C., Antoni, J., & Helsen, J. (2025). A multi-delay extension of the Discrete/Random Separation method. *Mechanical Systems and Signal Processing*, 236, 112959.
- Peeters, C., Leclere, Q., Antoni, J., Lindahl, P., Donnal, J., Leeb, S., & Helsen, J. (2019). Review and comparison of tachless instantaneous speed estimation methods on experimental vibration data. *Mechanical Systems and Signal Processing*, 129, 407-436.
- Tiboni, M., Remino, C., Bussola R., & Amici, C. (2022). A review on vibration-based condition monitoring of rotating machinery. *Applied Sciences*, 12(3), 972.

## BIOGRAPHIES

**Igor Makienko** earned his B.Sc., and M.Sc. degrees from the Faculty of Electrical Engineering at the Technion – Israel Institute of Technology, specializing in Machine Learning and Signal Processing. Currently, he serves as the Head of the PHM group at RSL Electronics Ltd, where he has accumulated over 20 years of experience developing AI-powered Predictive Maintenance Platforms for Aerospace, Energy, Maritime, and Transportation sectors. His current research focuses on integrating edge AI into PHM applications, enabling processing at the sensor level.

**Michael Grebshtein** received his Ph.D. degree in Aerospace Engineering from the Technion – Israel Institute of Technology. He is currently Lead diagnostic expert in RSL Electronics Ltd working on innovative PHM solutions including AI-based data analysis and physical models development.

**Eli Gildish** is a data scientist with nearly 20 years of experience in applied research, software and algorithm development in the fields of Machine Learning, Artificial Intelligence, Signal Processing and Computer Vision. He is a Senior Algorithm Developer at RSL Electronics and his current research focuses on application of deep learning to anomaly detection, condition monitoring, diagnostics, and prognostics. Eli studied M. Sc. in Applied Mathematics and B. Sc. in Aerospace Engineering at the Technion – Israel Institute of Technology.