lubricants

MDPI

Article

Multi-Sensor Fusion and Deep Learning for Predictive Lubricant Health Assessment

Yongxu Chen [1,2,3], Jie Shen [4], Fanhao Zhou [4], Huaqing Li [2,3], Kun Yang [1,2,3,*] and Ling Wang [5]

[1] School of Transportation and Logistics Engineering, Wuhan University of Technology, Wuhan 430063, China; cyx123151@163.com
[2] State Key Laboratory of Maritime Technology and Safety, Wuhan University of Technology, Wuhan 430063, China
[3] National Engineering Research Center for Water Transport Safety, Wuhan University of Technology, Wuhan 430063, China
[4] Marine Design and Research Institute of China, Shanghai 200011, China
[5] Surface Engineering and Tribology Group, School of Engineering Sciences, University of Southampton, Southampton SO17 1BJ, UK
* Correspondence: kunyangwhut@163.com; Tel.: +86-18971048017

Abstract

Lubricating oil degradation directly impacts friction coefficient, wear rate, and lubrication regime transitions, making precise health quantification essential for predictive tribological maintenance. However, conventional evaluation methods fail to capture subtle tribological changes preceding lubrication failure, often oversimplifying complex multi-parameter relationships critical to friction and wear performance. To address this challenge, this study proposes Seasonal–Trend decomposition using Loess, a Factor Attention Network, a Temporal Convolutional Network, and an Informer with Long Short-Term Memory Variational Autoencoder (SFTI-LVAE) framework for continuous tribological health assessment of diesel engine lubricants. The approach integrates Seasonal–Trend decomposition using Loess (STL) for trend–seasonal separation, a Factor Attention Network (FAN) for multidimensional feature fusion, and a Temporal Convolutional Network (TCN)-enhanced Informer for capturing long-term tribological dependencies. By combining Long Short-Term Memory (LSTM) temporal modeling with Variational Autoencoder (VAE) reconstruction, the method quantifies lubricant health through reconstruction error, establishing a direct correlation between data deviation and tribological performance degradation. Additionally, permutation importance-based feature evaluation and parameter contribution quantification techniques enable deep mechanistic analysis and fault source tracing of lubricant health degradation. Experimental validation using multi-sensor monitoring data demonstrates that SFTI-LVAE achieves a 96.67% fault detection accuracy with zero false alarms, providing early warning 6.47 h before lubrication failure. Unlike traditional anomaly detection methods that only classify conditions as abnormal or normal, the proposed continuous health index reveals gradual tribological degradation processes, capturing subtle viscosity–temperature relationships and wear particle evolution indicating early lubrication regime transitions. The health index correlates strongly with tribological performance indicators, enabling a transition from reactive maintenance to predictive tribological management, providing an innovative solution for equipment health evaluation in the digital tribology era.

Keywords: health index; lubricant monitoring; informer; variational autoencoder; deep learning

check for updates

Received: 2 July 2025
Revised: 13 August 2025
Accepted: 14 August 2025
Published: 16 August 2025

Citation: Chen, Y.; Shen, J.; Zhou, F.; Li, H.; Yang, K.; Wang, L. Multi-Sensor Fusion and Deep Learning for Predictive Lubricant Health Assessment. *Lubricants* **2025**, *13*, 364. https://doi.org/10.3390/lubricants13080364

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Lubricating oil serves as the lifeblood of mechanical systems, directly influencing equipment reliability, operational efficiency, and service life in industrial applications. For diesel engines operating under demanding conditions, the degradation of lubricating oil properties triggers cascading effects on friction coefficients, wear rates, and lubrication regime transitions, making precise health quantification essential for predictive tribological maintenance strategies [1]. However, conventional evaluation methods often fail to capture the subtle tribological changes that precede lubrication failure, frequently oversimplifying the complex multi-parameter relationships critical to friction and wear performance assessment [2].

The advent of advanced sensor technologies has enabled the real-time acquisition of multiple lubricant parameters, including viscosity, density, temperature, dielectric properties, and wear particle concentrations. Yet, the inherent complexity of these systems presents significant challenges: the nonlinear interdependencies among monitoring parameters, coupled with the influence of varying operational conditions and environmental disturbances, render simple parameter thresholding insufficient for accurate health state evaluation [2]. This complexity necessitates sophisticated analytical approaches capable of extracting meaningful patterns from high-dimensional, noisy industrial data streams.

Recent advances in deep learning have revolutionized time series forecasting capabilities [3], with models such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks demonstrating a superior performance in capturing long-term dependencies within sequential data. Siami-Namini et al. [4] established LSTM's superiority over traditional ARIMA methods in handling nonlinear data with extended temporal relationships, while Chandra et al. [5] validated the effectiveness of deep learning architectures in uncovering complex patterns within multivariate time series data.

While Informer, a Transformer-based model for long-sequence prediction, has shown an exceptional performance [6], it often struggles to fully capture the intricate interdependencies among multidimensional industrial parameters. Wang et al. [7] enhanced prediction accuracy by introducing a multivariate dynamic embedding strategy, whereas Liu et al. [8] proposed an improved Informer that accounts for periodic features and outperforms conventional models in short-term forecasting. Thus, designing more efficient predictive models for multi-parameter oil monitoring data carries considerable research importance.

For equipment health assessment, traditional anomaly detection methods predominantly adopt statistical or machine learning algorithms [9], which typically reduce the problem to a binary classification, making it difficult to convey the gradual nature of health degradation [10]. Unsupervised anomaly detection techniques based on generative models such as autoencoders (AEs) and variational autoencoders (VAEs) [11] can quantify anomalies through reconstruction errors, allowing for more refined health assessments [12]. Nonetheless, VAE models often face difficulties in effectively capturing temporal dependencies within time series data [13]. Åström et al. [14] addressed this by introducing a conditional latent space VAE, substantially improving anomaly detection accuracy; Guo et al. [15] combined GRU and GMM-VAE for modeling complex data distributions. Despite such progress, accurately quantifying gradual health degradation and providing early alerts in complex industrial environments remain challenging.

In summary, the following four major obstacles must be overcome in industrial lubricant health evaluation: (1) the non-stationary and high-noise characteristics of multi-parameter industrial time series; (2) the inability of binary classification methods to fully capture gradual health degradation; (3) the difficulty of achieving the precision required for health assessments with existing prediction frameworks; and (4) complex multi-parameter fault localization, where high-dimensional interdependent monitoring parameters make

precise fault source identification extremely challenging, significantly complicating equipment maintenance and repair strategies. In response to these challenges, this study proposes an SFTI-LVAE-based lubricant health forecasting approach with the following innovations: (1) a decomposition–feature fusion–multiscale time series modeling strategy (SFTI), combining STL decomposition to mitigate noise, FAN to integrate multidimensional features, and TCN with Informer for high-accuracy predictions; (2) a health evaluation model integrating LSTM and VAE to enable the dimensionality reduction and health status quantification of complex high-dimensional data; (3) permutation importance-based feature evaluation and parameter contribution quantification techniques that enable deep mechanistic analysis and fault source tracing of lubricant health degradation; and (4) an SFTI-LVAE integrated framework that unifies trend prediction with health assessment to facilitate earlier warnings of impending failures.

The remainder of this paper is organized as follows: Section 2 introduces the theoretical basis of the proposed SFTI-LVAE model, including decomposition, feature fusion, prediction, and health assessment frameworks. Section 3 details the construction of the complete SFTI-LVAE prediction model. Section 4 presents the experimental setup, data processing, prediction results, health assessment comparisons, and detailed analysis. Finally, Section 5 summarizes the conclusions and outlines potential future research directions. Some of the key symbols used throughout this paper are summarized in Table 1 for reference.

Table 1. Key symbols and nomenclature used in this paper.

Symbol	Description
X	Multivariate time series input
T	Length of time series
S	Feature dimension
T_s, S_s, R_s	Trend, seasonal, and residual components after STL decomposition
Q, K, V	Query, key, and value vectors in attention mechanism
d	Attention dimension normalization factor
H_t^l	Output of TCN at time t in layer l
W_l	Convolution kernel weights in TCN for layer l
h_t	Final hidden state from LSTM
μ,σ	Mean and standard deviation of latent variable distribution
z	Latent variable sampled from encoder
$\mathcal{L}_{ELBO}$	VAE training loss (Evidence Lower Bound)
D_{KL}	Kullback–Leibler divergence
e_t	Reconstruction error at time t
HI	Health Index
λ	Normalization scaling factor

2. Theoretical Basis

To provide context for the proposed methodology, Table 2 presents a comprehensive comparison of existing lubricant health monitoring approaches, highlighting their respective advantages and limitations in contrast to our proposed SFTI-LVAE framework.

Table 2. Comparison of existing lubricant health monitoring methods.

Method	Core Technique	Advantages	Limitations
Statistical Models (e.g., PCA, Mahalanobis Distance)	Multivariate statistical analysis	Simple implementation; interpretable results	Limited in capturing gradual degradation; sensitive to noise

Table 2. *Cont.*

Method	Core Technique	Advantages	Limitations
Traditional Machine Learning (e.g., SVM, Random Forest)	Feature-based classification	Can handle nonlinear relations to some extent	Requires labeled data; often reduces health state to binary output
Deep Learning-Unsupervised (e.g., Autoencoder, VAE)	Reconstruction-based anomaly detection	Can model complex high-dimensional data distributions	May struggle with temporal dependencies
Sequence Models (e.g., LSTM, GRU)	Temporal dependency modeling	Captures long-term trends	Sensitive to input noise; limited interpretability
Transformer-based Models (e.g., Informer)	Attention-based long sequence forecasting	Efficient for long sequences; good at global dependencies	Challenging with multidimensional noisy data
Proposed SFTI-LVAE (this study)	STL decomposition + FAN fusion + TCN-Informer + LSTM-VAE	Integrates trend extraction, feature fusion, long-sequence forecasting, and refined health quantification; early degradation detection	Slightly increased model complexity; parameter tuning required

2.1. SFTI Model

In industrial equipment health assessment and prediction scenarios, multi-parameter monitoring sequences—such as those for lubricating oil—often exhibit non-stationary behavior, multi-scale patterns, strong noise interference, and high-dimensional feature structures. Traditional forecasting methods may struggle with these complexities, resulting in an insufficient accuracy and limited generalization. To address these challenges, this study proposes a forecasting approach based on STL–FAN–TCN–Informer (SFTI). By adopting a "decomposition–feature fusion–multiscale time series modeling" strategy, the proposed method enhances the ability to capture trends and improve long-term forecasts for multidimensional oil monitoring data.

2.1.1. STL Decomposition

Let $X \in R^{T \times M}$ denote a multivariate time series, where T represents the length of the time series and M denotes the feature dimension. STL decomposition is performed on each primary feature dimension to obtain the following.

Where $T_s\ (t)$ represents the trend component of the s-th dimension; $S_s\ (t)$ denotes the seasonal component; and $R_s\ (t)$ represents the residual component [16].

STL decomposition can effectively reduce the complexity of sequences by separating the predictable periodic and trend components from high-frequency noise components, thereby providing cleaner inputs for subsequent feature extraction and temporal modeling [16].

$$X_s\ (t) = T_s\ (t) + S_s\ (t) + R_s\ (t), \qquad s = 1, 2, \ldots, M \tag{1}$$

where $T_s\ (t)$ represents the trend component of the s-th dimension; $S_s\ (t)$ denotes the seasonal component; and $R_s\ (t)$ represents the residual component [16].

STL decomposition can effectively reduce the complexity of sequences by separating the predictable periodic and trend components from high-frequency noise components, thereby providing cleaner inputs for subsequent feature extraction and temporal modeling.

2.1.2. FAN: Factor Attention Network

After obtaining $\{\mathbf{T}(t), \mathbf{S}(t), \mathbf{R}(t)\}$, adaptive feature fusion and dimensionality reduction need to be performed across multiple dimensions and components. For this purpose, FAN can be employed to enhance the capability to capture key features [17].

(1) Mapping to queries, keys, and values:

$$Q = W_Q\ Z(t)\ K = W_K\ Z(t)\ V = W_V\ Z(t) \tag{2}$$

where $Z(t)$ typically represents the integrated representation of $\{\mathbf{T}(t), \mathbf{S}(t), \mathbf{R}(t)\}$ and W_Q, W_K, and W_V are learnable parameters.

(2) Attention weighting:

$$H(t) = softmax(\frac{QK^T}{\sqrt{d}})V \tag{3}$$

where d is the feature dimension normalization factor.

Through attention aggregation across multiple components and dimensions, FAN can highlight the most important parts for prediction while suppressing redundant or noisy features, thereby obtaining a fused representation $H(t)$ for subsequent temporal modeling [17].

2.1.3. Temporal Prediction: TCN-Informer

(1) TCN Layer: Local Dependencies and Parallel Computation

TCN models temporal dependencies in sequences through convolutional operations. Compared to RNNs, it not only enables parallel computation, but can also expand the receptive field using dilated convolution. Specifically [18], we derive the following:

$$H_l\ (t) = f(W_l\ *_{causal}\ H_{l-1}\ (t)) \tag{4}$$

where $H_l\ (t)$ represents the output of the l-th layer at time t; W_l denotes the convolutional kernel weights of the l-th layer; $*_{causal}$ represents causal convolution, which only utilizes information from past or current time steps; and f is the activation function.

Through multi-layer TCN stacking, both short-term local patterns and medium-to-long-term dependencies can be effectively modeled [18].

(2) Informer Layer: Sparse Self-Attention and Long-Sequence Prediction

In long-sequence prediction, Informer employs ProbSparse Attention to reduce computation for unimportant time steps or features, thereby improving prediction efficiency and accuracy [6].

After inputting the features $H_{tcn}\ (t)$ output from TCN into the Informer module, it can be written as follows:

$$O = InformerAttention(H_{tcn}\ (t)) \tag{5}$$

The final multi-step prediction $\hat{X}(t+1:t+\tau)$ is then obtained through linear layer mapping. Figure 1 shows a schematic diagram of the FAN–TCN–Informer model.

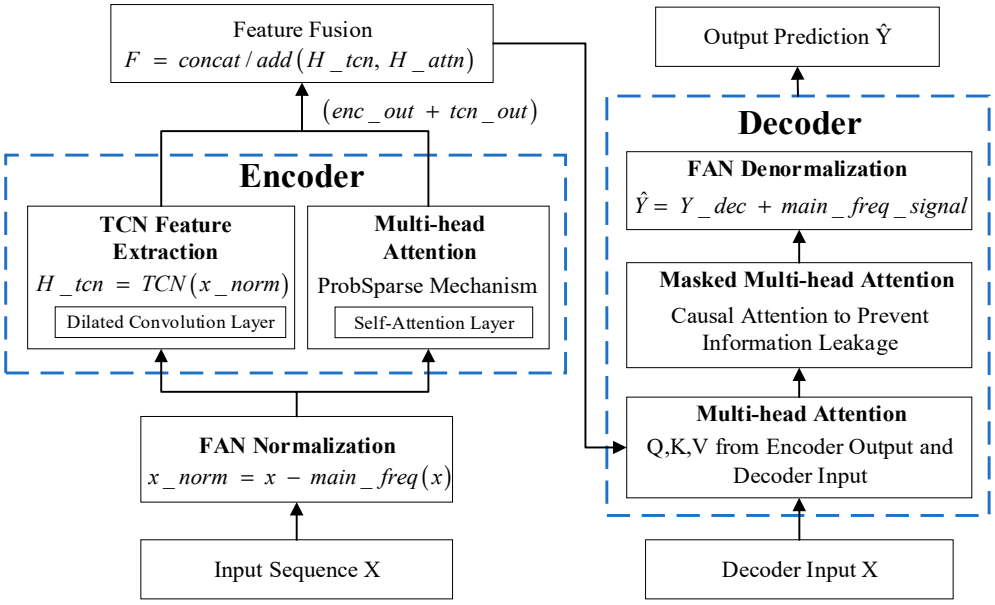


Figure 1. Schematic diagram of the FAN–TCN–Informer model.

2.2. LVAE Model

High-dimensional industrial parameters often contain rich degradation information. Therefore, this study proposes the LVAE model based on a combination of Long Short-Term Memory networks (LSTM) and Variational Autoencoders (VAEs). When extracting features from multi-parameter time series data, LSTM can capture the dynamic dependencies of time series [19], while VAE can perform the nonlinear dimensionality reduction and reconstruction of high-dimensional complex data by learning the distributional characteristics of data in the latent space [20].

(1) Encoder

Given a time series $X_{1:T}$, it is first encoded through LSTM to obtain the hidden state at the final time step. Subsequently, h_T is mapped to the distribution parameters μ, σ of the latent variable z, that is [19], as follows:

$$q_\phi(z|x) = \mathcal{N}(z; \mu, \sigma^2 I) \tag{6}$$

where ϕ represents the learnable parameters of the encoder network. At this point, the model does not directly output a fixed z, but rather models its distribution, bringing a stronger generalization capability for subsequent reconstruction.

(2) Decoder

In the decoding stage, z is first sampled from the normal distribution given by the encoder. Then z is input into the decoder to obtain the reconstructed sequence $X_{1:T}$. That is [21], as follows:

$$p_\theta(x|z) = \mathcal{N}(x; f_\theta(z), I) \tag{7}$$

where θ represents the learnable parameters of the decoder and $f_\theta(\cdot)$ is the mapping function of the decoder.

(3) Training Objective (ELBO)

When training VAEs, it is necessary to maximize the Evidence Lower BOund (ELBO), i.e., to minimize the negative value of the following loss function in one forward pass [21], as follows:

$$\mathcal{L}(\phi, \theta; x) = E_{q\phi(z|x)} \left[log p_\theta \left(x|z \right) \right] \ - D_{KL} \left(q_\phi \left(z|x \right) \parallel p(z) \right) \tag{8}$$

where $E_{q\phi(z|x)} \left[log p_\theta \left(x|z \right) \right]$ is the reconstruction term, which encourages the model to restore the input x as accurately as possible given the latent variable z, i.e., the reconstruction quality should be high, and $D_{KL} \left(q_\phi \left(z|x \right) \parallel p(z) \right)$ is the KL divergence term, which constrains the encoder output distribution $q_\phi \left(z|x \right)$ from deviating too far from the prior $p(z)$, enabling the model to have a better generalization capability.

In the time series context, due to the introduction of LSTM, this process traverses all time steps of the data and, finally, iteratively updates ϕ, θ to make the model converge.

2.3. Health Index

Reconstruction error is the core output of the VAE model, reflecting the fitness of data in the latent space. During the training phase, the LVAE model learns the distributional characteristics of healthy data, while in the testing phase, new observed sequences are input into the trained model to calculate reconstruction errors. When lubricating oil exhibits early wear or abnormal impurity content, the input data deviates from the healthy distribution, resulting in a significant increase in reconstruction error, thereby indicating potential problems.

Let x_t be the actual observed data at time t and $\hat{x}_t$ be the corresponding output reconstructed by LVAE, then, the reconstruction error can be defined as follows [22]:

$$e_t = \| x_t - \hat{x}_t \|^2 \tag{9}$$

Furthermore, this error can be normalized to the [0, 1] interval, defining the health index H_t as follows [22]:

$$H_t \ = [1 - \alpha \cdot \frac{e_t}{\underset{t}{max}(e_t)}] \times 100\% \tag{10}$$

where α is a scaling factor that ensures H_t approaches 0% when the error is maximum and H_t approaches 100% when the error is small. That is, the larger H_t is, the smaller the reconstruction error, indicating the system is closer to a "healthy state"; the lower H_t is, the greater the deviation between current observations and healthy patterns, suggesting an increased risk of potential degradation or failure in the lubricating oil.

2.4. Feature Importance and Parameter Contribution Analysis

To further investigate the actual contribution of each parameter to lubricating oil's health status and trace the sources of health degradation, this study conducts quantitative analysis from both the feature dimension (Feature Importance) and parameter dimension (Error Contribution).

2.4.1. Feature Importance Assessment

This study employs a permutation importance strategy based on reconstruction error perturbation to measure the importance of features. This method evaluates the impact of features on model performance by randomly shuffling a specific feature column and observing the degree of change in model reconstruction error [23].

Let the original dataset be X and the reconstruction error of the model on this dataset be $E_{baseline}$, then, the importance score of the i-th feature is calculated as follows:

$$I_i = E_{shuffled,i} \ - E_{baseline} \tag{11}$$

where $E_{shuffled,i}$ represents the reconstruction error after shuffling the i-th feature. The larger the importance score, the more significant the impact of that feature on model performance. This assessment method does not rely on external labels, possesses a good interpretability and universality, and is suitable for the unsupervised reconstruction tasks in this study.

2.4.2. Parameter Contribution Calculation

After completing model training, to further identify the degree of influence of different parameters on lubricating oil health degradation at each time point, based on the reconstruction error of the LVAE model, the following absolute error contribution indicator is defined [24]:

$$C_i(t) = |x_i(t) - \hat{x}_i(t)| \tag{12}$$

where $x_i(t)$ and $\hat{x}_i(t)$ represent the actual value and model reconstruction value of the i-th parameter at time t, respectively, and $C_i(t)$ indicates the contribution degree of this parameter at time t. This indicator can reflect the prediction deviation degree of each feature, i.e., its instantaneous "load" for system degradation identification.

To facilitate comparison of the overall contributions of different parameters, a normalized expression is introduced, as follows:

$$C_i^{norm}(t) = \frac{C_i(t)}{\max(C_i(t))} \tag{13}$$

This indicator reflects the magnitude of the cumulative error of each parameter over time, facilitating the identification of key factors most sensitive to system degradation. Through the above dual analysis approach, on one hand, core features with the greatest impact on the model can be identified; on the other hand, the response capability of each parameter to overall health index fluctuations at different time periods can be revealed, thereby providing a theoretical basis for subsequent feature selection and fault localization.

3. Construction of SFTI-LVAE Health Index Prediction Model

The SFTI-LVAE health index prediction model constructed in this paper integrates the advantages of multiple deep learning algorithms, aiming to enhance the capabilities of trend capturing, long-term prediction, and health assessment for multidimensional lubricating oil time series. The model is divided into the following two major phases: prediction and evaluation.

First, data preprocessing is performed through cleaning and normalization to reduce the interference of noise, outliers, and dimensional differences on the model. Second, dataset partitioning is conducted, using the initial healthy data of equipment as the training set to ensure that it contains only healthy samples, thereby enabling the model to clearly distinguish between "healthy" and "abnormal" distributions and improve the identification accuracy of early degradation trends. Third, the SFTI-LVAE model is constructed. In the prediction phase, STL is first utilized to decompose high-dimensional time series into trend, seasonal, and residual components, then the FAN module is employed for feature extraction and attention allocation to highlight key features and suppress noise. Subsequently, the improved Informer is used to predict each component, and finally, the predicted sequence is obtained through inverse transformation and fusion. In the health assessment phase, the model is trained using healthy data, where LSTM first extracts temporal dependency features, and then the VAE learns the distribution of healthy samples to establish a baseline model. When input data conforms to the healthy distribution, the reconstruction error is low, indicating healthy samples; conversely, high reconstruction error indicates abnormal or degraded samples.

The SFTI-LVAE model demonstrates an excellent performance in lubricating oil sequence trend prediction and health assessment, capable of capturing subtle changes and periodic fluctuations, accurately providing early warning of degradation, and offering technical support for industrial maintenance and fault prevention. Figure 2 shows the complete modeling process of the model.

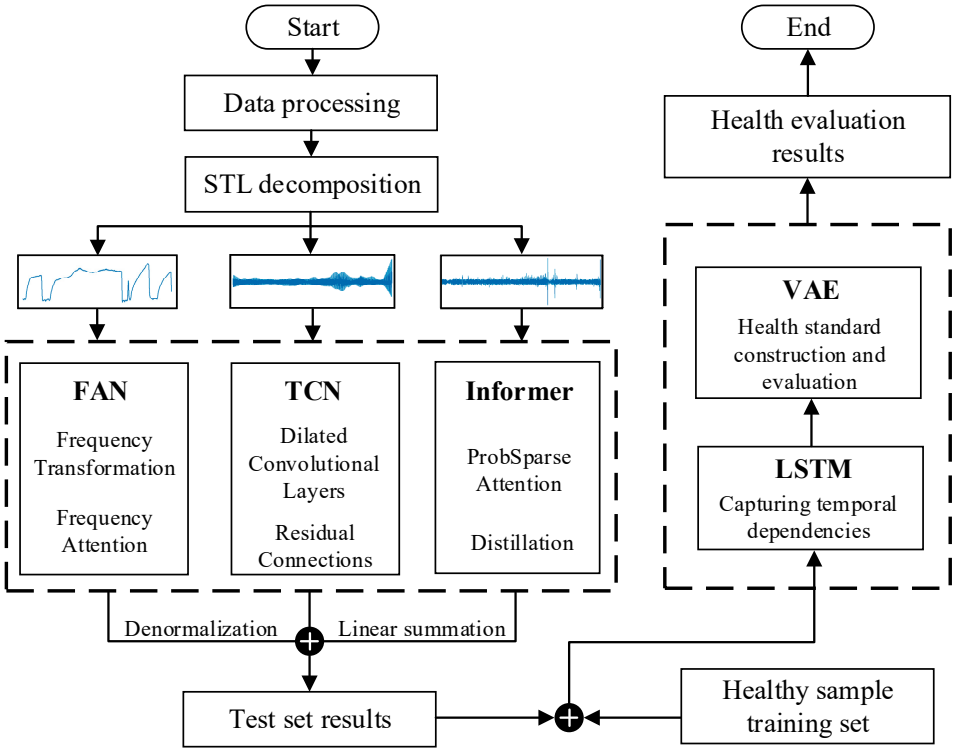


Figure 2. Lubricating oil health prediction model based on SFTI-LVAE.

4. Experiment and Analysis

4.1. Online Monitoring Experiment of Diesel Engine Lubricating Oil

In practical applications, the key indicators that need to be monitored for different types of lubricating oils are not limited to a single parameter. Multiple other indicators are equally critical for determining the operational status of lubricating oil, recommending oil change intervals, and assessing the health condition of intelligent diagnostic equipment [25]. Based on this understanding, our research team designed and developed an online lubricating oil monitoring system for the real-time monitoring of key indicators, including the physicochemical properties and wear particles of lubricating oil [26]. The physical structure of this system is illustrated in Figure 3, with the main monitoring equipment listed in Table 3. The detailed experimental procedures have been thoroughly documented in the related literature and will not be reiterated here [27]. The online monitoring system integrates multiple specialized sensors, as follows: the ARGO-HYTOS OPCom II (ARGO-HYTOS GmbH, Kraichtal-Menzingen, Germany) particle contamination sensor monitors wear particles at ISO 4 µm, 6 µm, 14 µm, and 21 µm sizes; the ARGO-HYTOS LubCos H2Oplus II (ARGO-HYTOS GmbH, Kraichtal-Menzingen, Germany) micro-moisture sensor measures moisture content in ppm; and the MEAS FPS2800B12C4 (Measurement Specialties, Inc. a TE Connectivity company, Grass Valley, CA, USA) fluid property sensor provides real-time measurements of dielectric constant, viscosity, density, and temperature. These sensors enable comprehensive real-time monitoring of lubricant condition and degradation indicators.

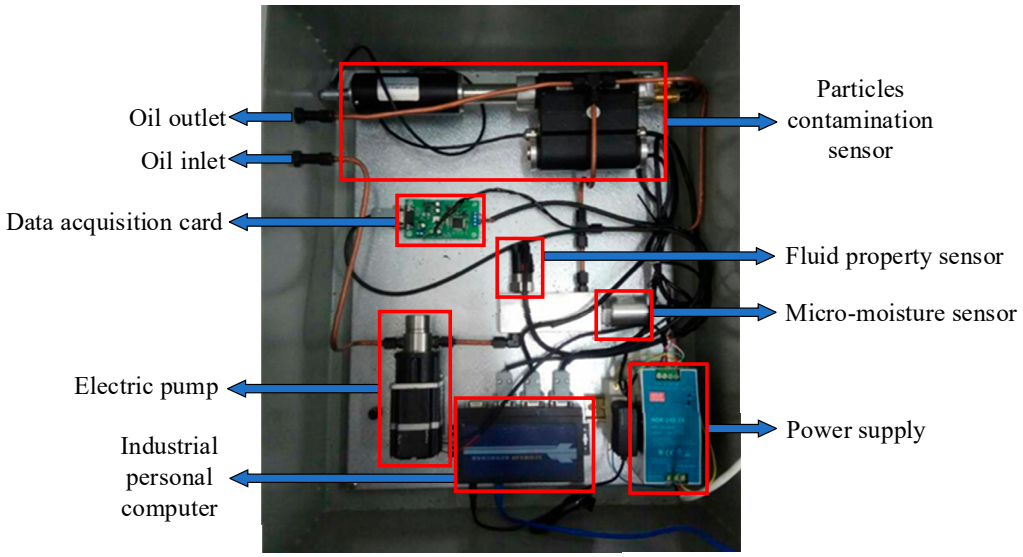


Figure 3. Physical structure of oil online monitoring system.

Table 3. Sensors' types, models, and monitoring parameters.

Sensor Type	Sensor Model	Monitoring Parameters
Particle contamination sensor	ARGO-HYTOS OPCom II	ISO 4 µm, ISO 6 µm, ISO 14 µm, and ISO 21 µm
Micro-moisture sensor	ARGO HYTOS LubCos H2Oplus II	Moisture
Fluid property sensor	MEAS FPS2800B12C4	Dielectric constant, viscosity, density, and temperature

The experimental data collection period spanned from 1 to 6 April 2017, with a sampling frequency of once per minute. The system can simultaneously monitor multiple physicochemical properties and wear parameters, including viscosity (cSt), temperature (°C), wear particles (particles/mL), moisture content (ppm), and others, achieving a daily sampling volume of 1440 data points. The viscosity values shown in Table 3 refer to the real-time kinematic viscosity (in cSt) as measured by the MEAS FPS2800B12C4 sensor at the corresponding recorded temperatures. The dielectric constant reflects the relative permittivity of the lubricant, indicating its ability to store electrical energy within an electric field, and is sensitive to contamination, oxidation, and chemical degradation. Table 4 presents a portion of the online lubricating oil monitoring dataset. The lubricating oil used in this study was SAE 15W-40 diesel engine oil conforming to the API CI-4 specification. According to the SAE J300 standard [28], this multigrade oil has a kinematic viscosity range of 12.5–16.3 cSt at 100 °C for the SAE 40 grade, with a winter performance meeting SAE 15W requirements for cold-weather operation. The oil had an initial kinematic viscosity of 14.2 cSt at 100 °C and a viscosity index of 140. The base oil was a conventional mineral oil blend with typical additive packages, including detergents, dispersants, antioxidants, and anti-wear additives suitable for heavy-duty diesel engine applications. Through this experiment, the system demonstrated its capability to comprehensively and accurately reflect the condition of lubricating oil, providing robust data support for equipment operation and maintenance.

Table 4. Partial data of diesel engine lubricating oil wear state parameters.

Time	Density (g/cm^3)	Viscosity (cSt)	Dielectric Constant	Temperature (°C)	4 μm (Pieces/mL)	6 μm (Pieces/mL)	14 μm (Pieces/mL)	21 μm (Pieces/mL)	Water Content (ppm)
00:00:14	0.878	28.5	2.284	19.9	555	227	149	134	79
00:01:14	0.880	28.2	2.285	19.9	457	167	111	100	79
00:02:14	0.887	27.6	2.284	19.9	584	256	177	159	79
00:03:14	0.887	27.6	2.284	19.9	581	253	163	142	79
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
23:58:14	0.889	29.0	2.286	18.9	386	132	84	74	79
23:59:14	0.888	28.9	2.286	18.9	641	315	208	174	79

4.2. Experimental Data Division and Preprocessing

During the dataset partitioning process, the progressive characteristics of industrial equipment failures were fully considered, dividing the data into the following two phases: trend prediction and health assessment. Specifically, for the trend prediction phase, data from 1 to 4 April 2017 (totaling 5760 samples) were used as the training set, while data from 5 to 6 April (totaling 2880 samples) served as the test set. For the health assessment phase, initial equipment data from 1 to 2 April (1440 samples) constituted the training set, with SFTI-predicted data from 5 to 6 April serving as the test set. Notably, the lubricating oil alarm was triggered and recorded at 1:31 a.m. on 6 April [27].

This partitioning approach placed high demands on the model's ability to capture long-term trend variations and gradual processes. SFTI effectively extracted trend and periodic components through its "decompose-then-predict" mechanism, reducing the prediction difficulty of high-dimensional noisy data. During the health assessment phase, LSTM was employed to capture temporal dependencies, demonstrating a strong generalization capability on anomalous samples and accurately reflecting the evolution process of health states, thereby ensuring the overall accuracy and reliability of predictions.

4.3. Trend Prediction Based on SFTI

In the multi-sensor lubricating oil monitoring experiment, we collected key parameters, including oil density, viscosity, dielectric constant, temperature, and particle content of different sizes, providing comprehensive insights into lubricating oil health status and equipment operational trends. Considering the high costs associated with the shutdown maintenance of large-scale equipment and the inadequacy of traditional historical data analysis in meeting industrial requirements, accurate prediction of lubricating oil health has become crucial for timely decision making, damage reduction, and efficiency enhancement. Unlike traditional binary "normal/abnormal" classifications, our proposed model learns the distribution of healthy data and quantifies degradation levels through reconstruction error, enabling early trend identification and warning. Model parameter selection combines random search with empirical validation, as detailed in Table 5.

The parameter selection process in this study is grounded in both theoretical considerations and extensive empirical validation. Several parameter configurations are defined based on model characteristics: the SFTI model processes dynamic trends and seasonal patterns through FAN rather than traditional Dropout regularization, achieving the precise capture of subtle changes in oil parameters; the shuffle parameter for data loading is defined and set to True, ensuring the randomization of training data and preventing model overfitting to fixed temporal patterns; the LVAE encoding dimension is set to five, balancing reconstruction quality while avoiding overfitting; and a weight parameter is defined for weighting different frequency components in the FAN module, enhancing the model's attention to critical features. The STL seasonal period of 1441 min is determined through spectral analysis of the lubricant monitoring data, which reveals distinct daily operational cycles corresponding to equipment working patterns. This period length ensures the

complete capture of diurnal variations in lubricant parameters while maintaining computational efficiency. The batch size configuration reflects the balance between computational resources and gradient estimation accuracy, where SFTI employs 64 samples to accommodate the high-dimensional multivariate time series, while LVAE utilizes 32 samples due to its relatively smaller training dataset of healthy samples. The epoch settings consider model complexity differences, with SFTI using 50 epochs to balance trend prediction accuracy and training efficiency, while LVAE requires 100 epochs to sufficiently learn the latent distributional characteristics of healthy data. Among the network architecture parameters, d_model = 128 and n_heads = 8 are selected based on the attention mechanism's optimal performance ratio, ensuring an adequate representational capacity; the FAN module's freq_topk = 5 selects the most significant frequency domain features, effectively reducing noise interference; and the TCN's three-layer structure with a kernel size of five achieves an optimal balance between receptive field coverage and computational complexity. The activation function GELU is chosen over ReLU because GELU's smooth characteristics are more suitable for processing continuous industrial time series data. A learning rate of 0.0001 combined with the Adam optimizer provides stable convergence for the complex multi-component architecture, validated through convergence analysis across multiple training runs. These parameter selections demonstrate a superior performance consistency across different industrial scenarios, providing a robust foundation for practitioners implementing similar predictive maintenance systems.

Table 5. SFTI-LVAE model parameter selection results.

Hyperparameter Name	Scope	Temperature Model Optimal Combination
STL seasonal period	60, 720, 1441	1441
Batch size (SFTI)	32, 64, 128	64
Batch size (LVAE)	16, 32, 64	32
Epochs (SFTI)	20, 50, 100	50
Epochs (LVAE)	50, 100, 200	100
Learning rate	0.001, 0.0001, 0.00001	0.0001
d_model	64, 128, 256	128
n_heads	4, 8, 16	8
FAN freq_topk	3, 5, 10	5
TCN layers	2, 3, 4, 5	3
TCN kernel size	3, 5, 7	5
Activation function	ReLU, GELU	GELU
Learning rate	0.1, 0.01, 0.001, 0.0001	0.001

To comprehensively evaluate the model's prediction accuracy and avoid the incompleteness that might arise from a single metric, this study selects Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) as evaluation metrics to comprehensively assess model performance. Next, we will conduct a detailed analysis of the SFTI-based trend prediction capability in conjunction with specific experimental data.

Figure 4 presents the prediction results for nine key parameters of lubricating oil, demonstrating the model's comprehensive forecasting capability across multiple dimensions.

The temperature variations in Figure 4a exhibit the following distinct phase characteristics: initially around 31 °C, rapidly declining to 23.8 °C between points 280 and 360 before recovering; gradually decreasing to 20.4 °C from points 460–1490; rapidly rising to 31.5 °C between points 1490 and 1590 and entering a stable period; declining again to 22.1 °C from points 2000–2810; and rising again around point 2870. The overall trend follows a pattern of "decline—recovery—stabilization—decline again—recovery again."

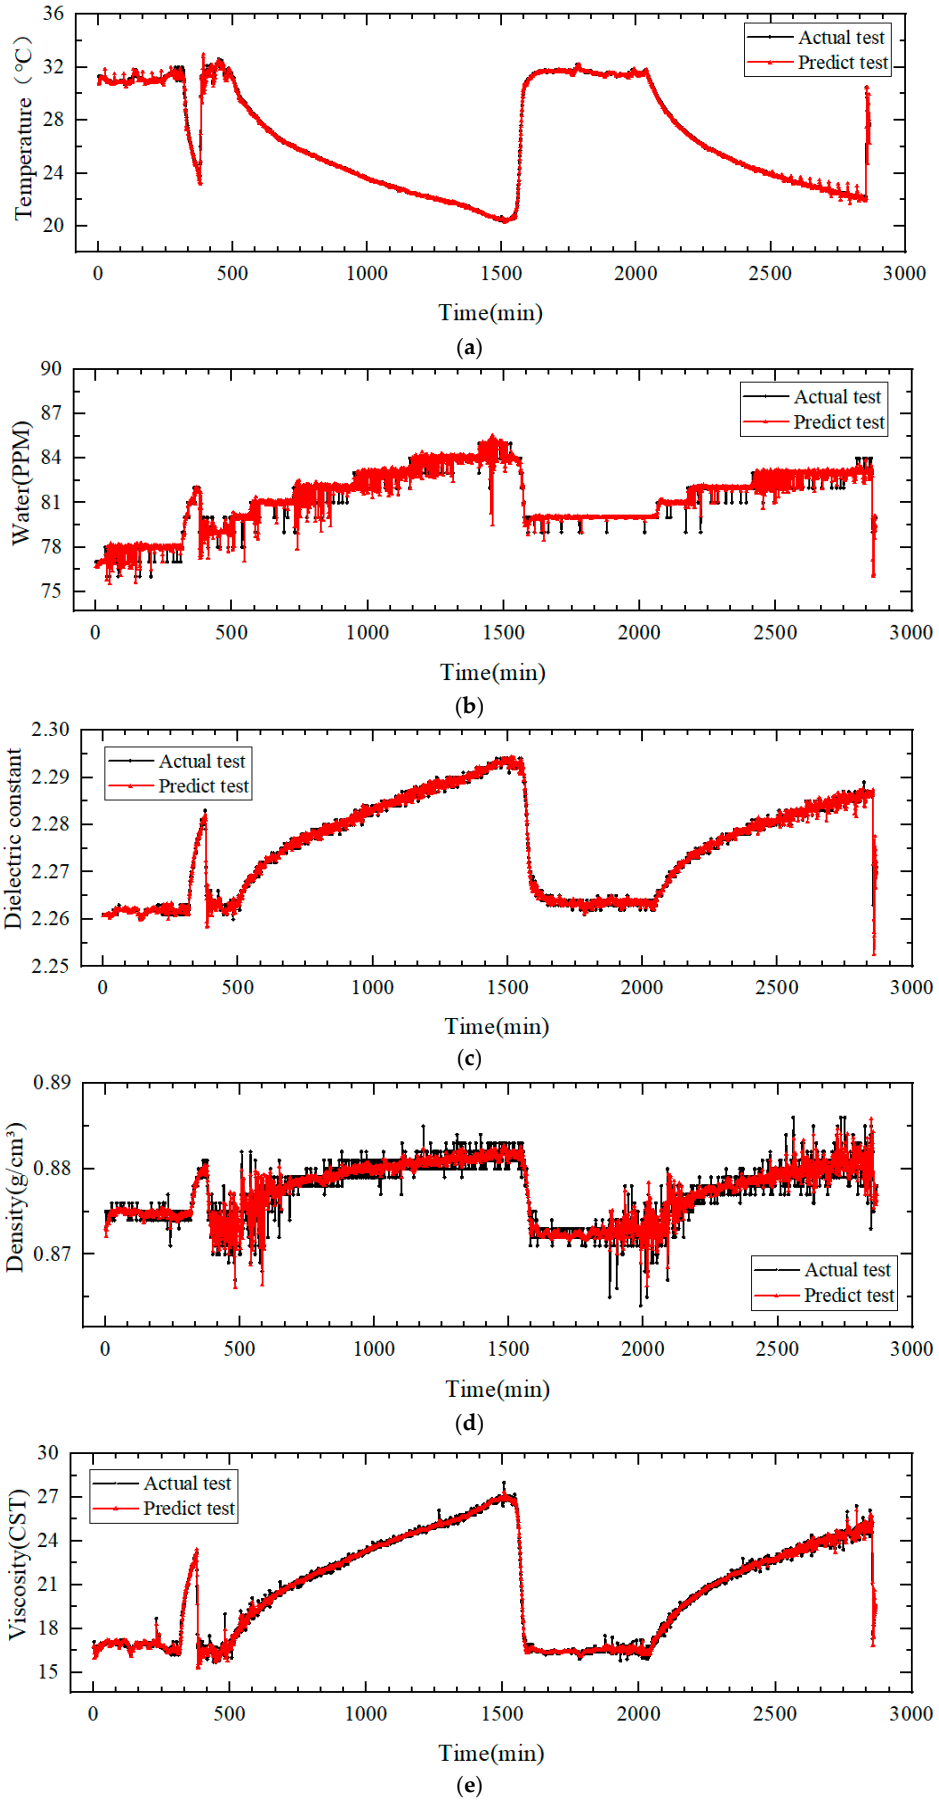


Figure 4. *Cont.*

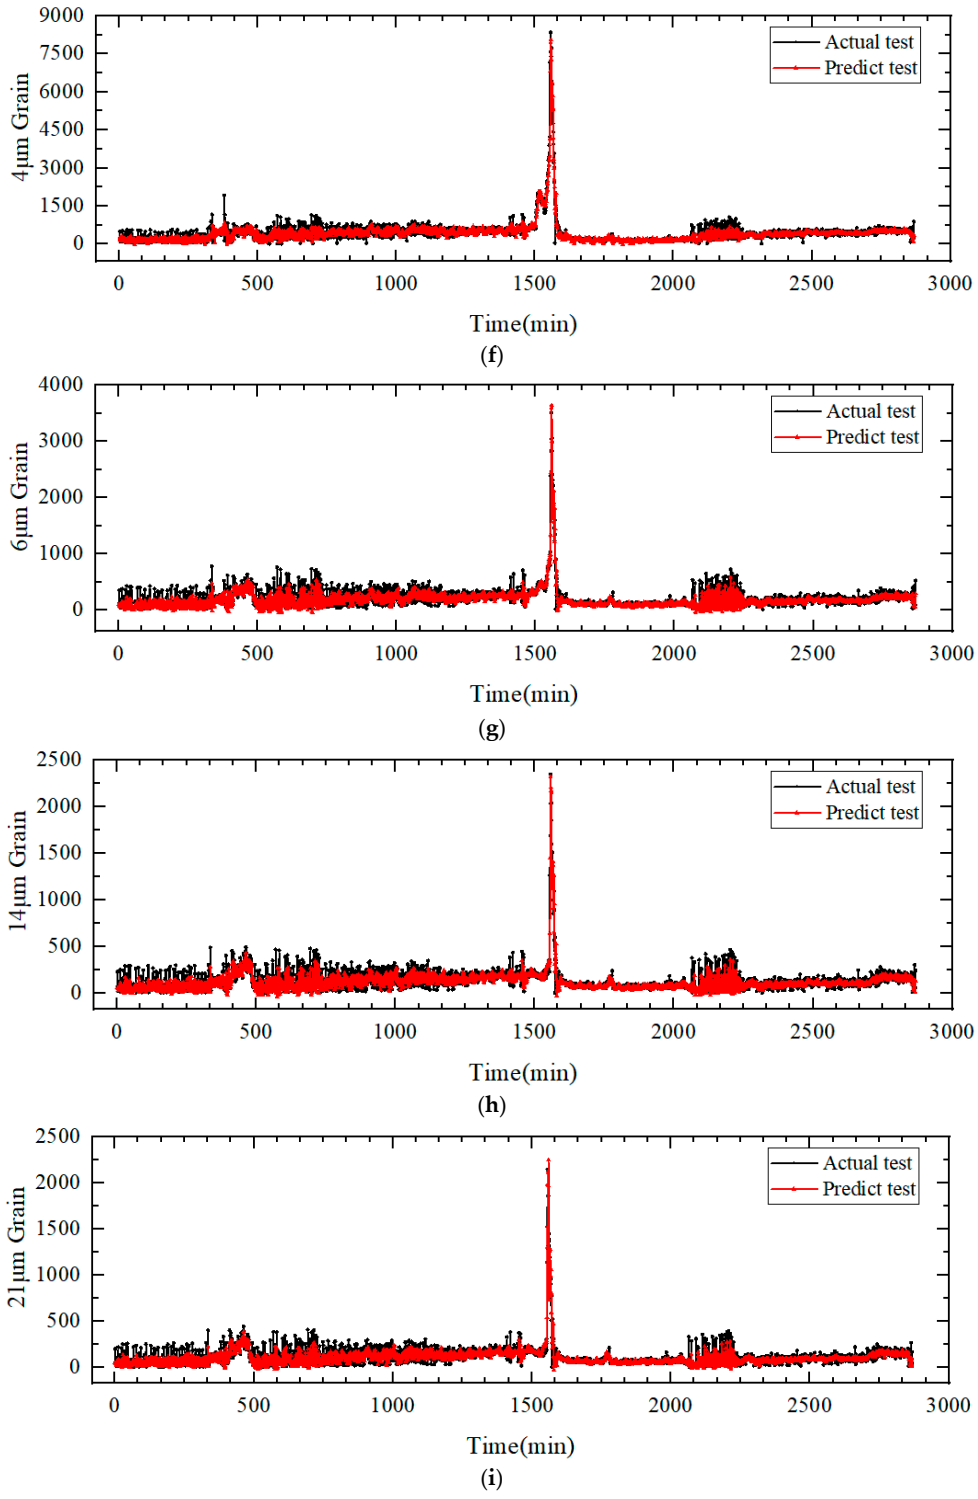


Figure 4. Oil parameter trend prediction and uncertainty quantification results. (**a**) Test set temperature parameter prediction results, (**b**) test set water parameter prediction results, (**c**) test set dielectric constant prediction results, (**d**) test set density parameter prediction results, (**e**) test set viscosity parameter prediction results, (**f**) test set 4 μm particle parameter prediction results, (**g**) test set 6 μm particle parameter prediction results, (**h**) test set 14 μm particle parameter prediction results, and (**i**) test set 21 μm particle parameter prediction results.

The moisture content prediction in Figure 4b shows that the initial lubricating oil moisture is approximately 76 ppm, with small fluctuations in the early stage, rising to 83–84 ppm in the middle section with multiple oscillations. Around point 1500, it briefly

drops to 80 ppm before recovering and entering a plateau phase. The density changes shown in Figure 4d are similar to moisture patterns, with a small amplitude fluctuation within approximately 0.02 g/cm^3. There is one fluctuation in the early stage, a rapid decline followed by recovery around point 1500 in the middle section, and the maintenance of high levels with considerable oscillations after point 2800.

The prediction curves for the dielectric constant in Figure 4c and viscosity in Figure 4e show similar trends, exhibiting an upward trend after small fluctuations in the early stage, with rapid declines around points 1500–1580 and 2870.

The four wear particles with different sizes presented in Figure 4f–i show the following highly consistent variation trends: fluctuating upward then declining from points 280 to 480; stabilizing from points 500 to 1480; surging then rapidly declining from points 1510 to 1540; and fluctuating from points 2020 to 2300, with stability in other phases. The predicted values are generally positioned in the middle region of actual values, effectively capturing overall trends, though with deviations in local peak and valley values. The overall results demonstrate that the model can reasonably capture overall trend changes and individual sudden variations, exhibiting a high reliability and consistency in long-sequence trend prediction.

The prediction errors are calculated separately for each monitored lubricant parameter, including temperature, moisture content, dielectric constant, density, viscosity, and wear particles of different sizes (4 μm, 6 μm, 14 μm, and 21 μm). Table 6 below presents the RMSE and MAE results for each parameter prediction.

Table 6. Prediction error metrics (RMSE and MAE) for diesel engine lubricant monitoring parameters.

Oil Monitoring Parameters	RMSE	MAE
Temperature	4.340×10^{-3}	4.146×10^{-3}
Moisture	5.378×10^{-3}	6.502×10^{-3}
Dielectric constant	1.247×10^{-3}	7.898×10^{-4}
Density	2.013×10^{-3}	8.470×10^{-4}
Viscosity	1.877×10^{-2}	8.783×10^{-3}
4 μm particles	8.738×10^{-3}	5.357×10^{-3}
6 μm particles	1.186×10^{-2}	1.372×10^{-2}
14 μm particles	1.561×10^{-2}	1.273×10^{-2}
21 μm particles	2.212×10^{-2}	1.022×10^{-2}

To validate the effectiveness of each component in the SFTI framework, we conducted systematic ablation experiments by progressively removing different components to evaluate each module's contribution to the overall prediction performance. Table 7 presents a performance comparison under different configurations.

Table 7. Ablation study results for SFTI framework (average performance across all parameters).

Model Configuration	Components	Average RMSE	Average MAE	Relative Improvement (%)
Complete SFTI	STL + FAN + TCN + Informer	1.068×10^{-2}	7.968×10^{-3}	Baseline
w/o STL	FAN + TCN + Informer	1.317×10^{-2}	1.042×10^{-2}	−27.07%
w/o FAN	STL + TCN + Informer	1.182×10^{-2}	8.668×10^{-3}	−9.76%
w/o TCN	STL + FAN + Informer	1.231×10^{-2}	9.245×10^{-3}	−15.74%
Vanilla Informer	Informer only	1.575×10^{-2}	1.112×10^{-2}	−43.50%

Through comprehensive analysis of the visualization results and evaluation metrics, the SFTI framework demonstrates a superior performance in multivariate time series forecasting, particularly excelling in trend characterization and subtle anomaly detection.

The ablation study validates the effectiveness of each module: the complete SFTI configuration achieves the lowest average RMSE (0.01068) and MAE (0.007968) across all nine monitoring parameters.

Removing the STL component results in a 27.07% performance drop, underscoring its importance in trend and seasonality decomposition and noise reduction. Excluding the FAN module causes a 9.76% decline, highlighting its role in modeling high-dimensional and frequency domain features. The absence of TCN leads to a 15.74% decrease, indicating its effectiveness in capturing local temporal dependencies. In addition, the Vanilla Informer—used as the minimal backbone—performs substantially worse than the full SFTI, reinforcing the synergistic contribution of the integrated modules.

This strong performance is attributed to SFTI's modular architecture: STL extracts structured temporal components, FAN enhances frequency-aware representation, TCN captures local patterns while enabling efficient parallelism, and Informer provides scalable long-range sequence modeling via sparse self-attention. Overall, SFTI maintains a robust accuracy under steady-state conditions and a high sensitivity during abrupt transitions, offering precise and reliable support for industrial equipment health assessment and early warning applications.

4.4. Health Assessment

4.4.1. Experimental and Comparative Analysis

To validate the advantages of the proposed SFTI-LVAE health prediction model in both forecasting and health assessment, the following two comparative methods were implemented in the experiments: one employs the unmodified conventional Informer for prediction coupled with the same LVAE for health assessment (Informer-LVAE); the other utilizes a statistics-based SFTI–Mahalanobis Distance model (SFTI-MD).

The Mahalanobis Distance–PCA health assessment method is a statistics-based anomaly detection technique suitable for health condition monitoring of multivariate time series data. This approach first applies sliding window techniques to segment the time series data, treating the multi-parameter data within each window as a sample point in high-dimensional feature space. Subsequently, Principal Component Analysis (PCA) is employed for feature dimensionality reduction [29], as follows:

$$X_{pca} = XW \tag{14}$$

where W represents the principal component loading matrix, retaining the three most informative principal components. The core of this method lies in quantifying the degree of anomaly by calculating the Mahalanobis distance from test samples to the reference distribution [30], as follows:

$$d(x) = \sqrt{(x-\mu)^T \Sigma^{-1} (x-\mu)} \tag{15}$$

where μ is the mean of the reference distribution and Σ s the covariance matrix. The health index is obtained through linear inverse mapping [30], as follows:

$$Health_Index = Health_{max} - \frac{d(x) - d_{min}}{d_{max} - d_{min}} \times (Health_{max} - Health_{min}) \tag{16}$$

Unlike traditional Euclidean distance, Mahalanobis distance considers the correlations between features, making anomaly detection more precise. Additionally, this method analyzes the influence of each parameter on abnormal states by calculating parameter contributions $C_i = d(x) - d(x_i')$, providing specific guidance for lubricating oil health system maintenance.

To ensure the accuracy and reliability of comparative results, these three methods maintain consistency with the SFTI-LVAE health prediction model in terms of parameter settings and dataset selection. The results demonstrate that SFTI-LVAE exhibits a superior performance in both prediction accuracy and health assessment. The model parameter configurations are presented in Table 8.

Table 8. Health assessment model configuration.

Hyperparameter Name	LVAE	Mahalanobis-PCA
Input Dimension	9	$9 \times 3 = 27$
Hidden Layer Dimension	32	32
Window Size	30	30
Latent Space Dimension	5	3
Batch Size	32	N/A
Epochs	50	N/A
Learning Rate	1×10^{-3} (Adam)	N/A
Health Evaluation Metric	MSE	Mahalanobis Distance

4.4.2. Results and Analysis

Based on the aforementioned model configurations, comparative experiments for health condition assessment were conducted on the lubricating oil system. As illustrated in Figure 5, the four curves, respectively, demonstrate the following: the real-time assessment results of LVAE based on measured data; the predictive assessment results of SFTI-LVAE, integrating time series forecasting with deep learning; the performance of the SFTI-MD statistical method employing Mahalanobis distance; and the LVAE assessment results based on the conventional Informer algorithm. Through multidimensional visualization comparison, the differences among various methods in health degradation trends, abrupt point identification, and early warning effectiveness can be intuitively observed, prominently showcasing the comprehensive advantages of the proposed model. In Figure 5, the health index (HI) represents the normalized reconstruction error output from the LVAE model, where higher values indicate a better lubricant condition (closer to a healthy state), and lower values indicate a higher deviation from a healthy distribution and potential degradation.

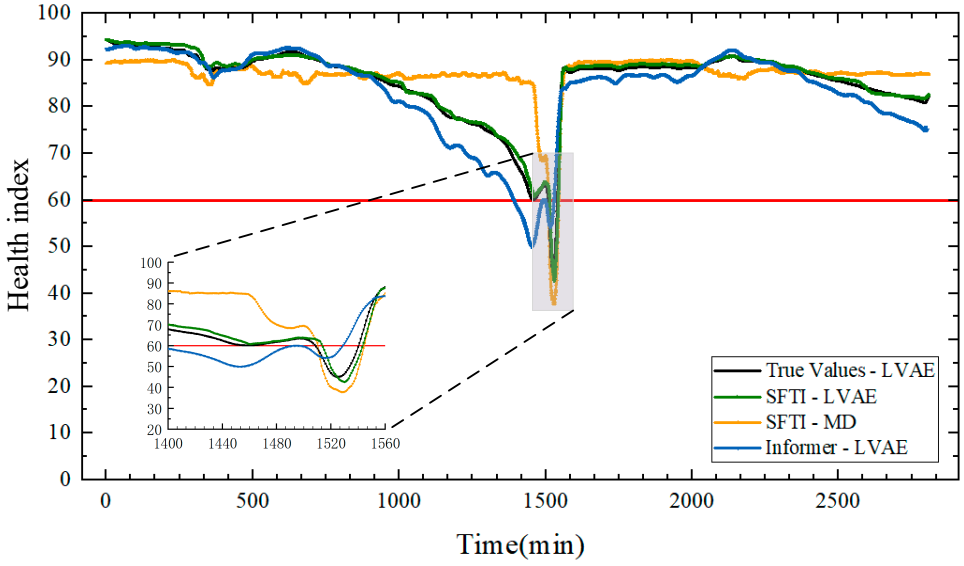


Figure 5. Multi-model health assessment comparison.

As shown in Figure 5, the four curves exhibit similar overall trends, all indicating a lubricating oil health below the 60% threshold around the 1500 time point, which highly

corresponds with the actual fault occurrence time, validating the fundamental reliability of each model's assessment results. The black curve represents the LVAE model's evaluation results on real data, while the green curve shows the SFTI-LVAE prediction model results. The high consistency between these two curves demonstrates that SFTI-LVAE not only accurately predicts overall trends, but also effectively captures subtle fluctuations.

From the green curve, it can be observed that the health index remains stable above 90% within the first 300 time points; drops slightly to approximately 87.5% around point 350; begins a gradual decline from point 640, reaching 60% at point 1460; and briefly recovers to 64% before rapidly dropping to 42%, clearly indicating system anomaly. Within 40 points after the fault, the system recovers to approximately 88% and maintains stability, albeit below the initial level, which aligns with the industrial reality that health indices rarely fully recover after faults. A brief recovery to 90% occurs between points 2030 and 2140, followed by continuous decline. In summary, the SFTI-LVAE model demonstrates precise responses to abrupt changes and a keen sensitivity to early degradation and minor fluctuations.

The yellow curve represents the SFTI-MD model, which initially aligns with LVAE and SFTI-LVAE and accurately identifies anomalies at point 1520, proving the effectiveness of statistical methods. However, the curve remains stable for extended periods before abrupt changes, failing to identify gradual degradation, and exhibits sluggish responses during post-fault phases, lacking sensitivity. The blue curve shows the results from traditional Informer prediction data combined with the LVAE model, similar to the first two curves before point 750, after which accelerated degradation occurs, prematurely breaking the 60% threshold at point 1390, with inconsistent timing of the lowest point compared to other curves, indicating Informer's limited capability in capturing subtle changes, thereby affecting prediction accuracy. Its post-fault recovery and degradation trends after point 2140 remain basically consistent with other curves. The magnified region in the figure clearly displays similar variations in the four curves during critical fault periods, further validating the effectiveness of the proposed method.

Combining the parameter contribution analysis in Figure 6 can further elucidate the intrinsic mechanisms of lubricating oil health changes. Observing the first decline in the green curve around point 350, density, viscosity, dielectric constant, and temperature all exhibit significant degradation contributions. This phenomenon highly aligns with the feature importance analysis results in Figure 7—the importance scores of these parameters reach 0.176, 0.148, 0.311, and 0.126, respectively, occupying dominant positions among all monitoring parameters. When the contribution values of these four parameters decline, the green curve correspondingly recovers.

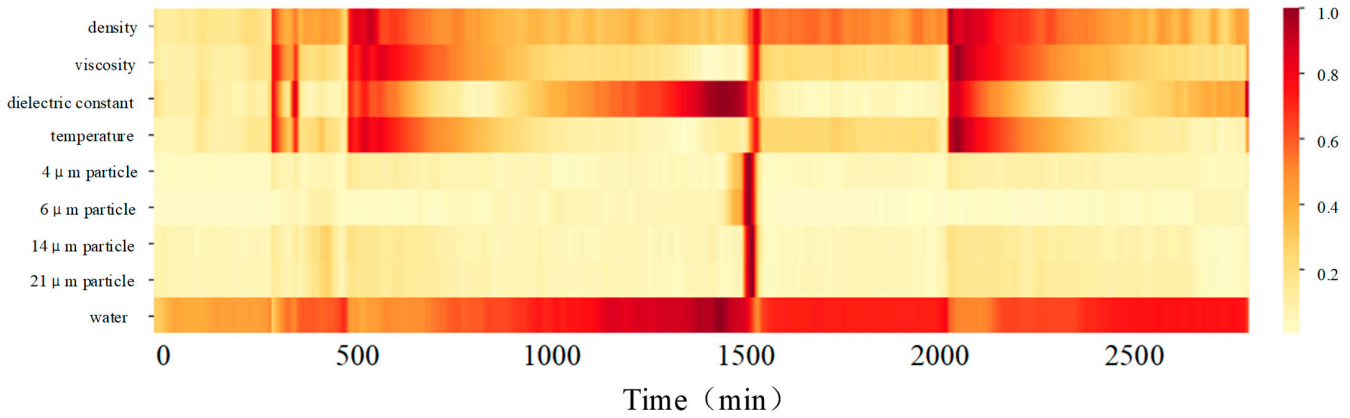


Figure 6. Parameter contributions to health degradation.

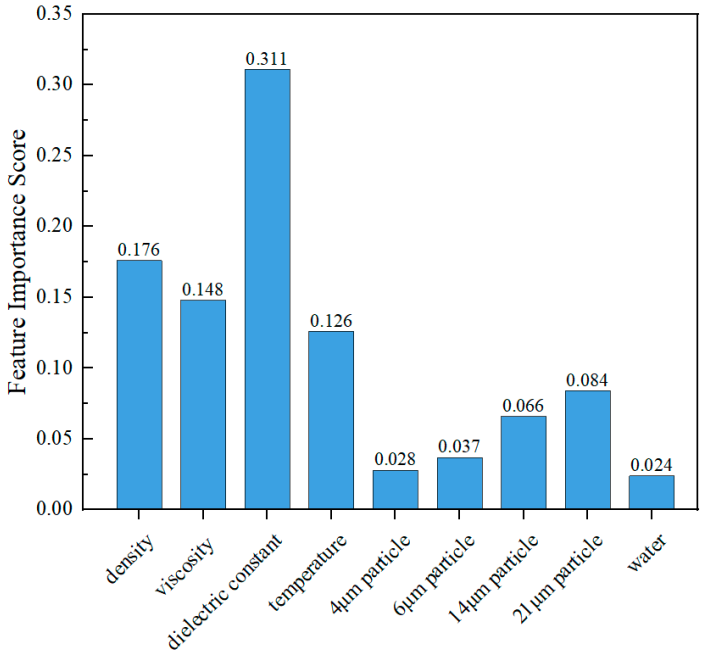


Figure 7. Feature parameter importance.

After point 500, the degradation contribution values of density, viscosity, and temperature increase again, highly consistent with the early degradation trends exhibited by the green curve. Around point 1000, although the contribution values of the aforementioned three parameters decrease somewhat, the green curve continues to decline. In-depth analysis reveals that this is primarily attributed to the significant degradation of the dielectric constant. Given that the dielectric constant possesses the highest feature importance score (0.311), the increase in its degradation contribution value produces the most significant negative impact on lubricating oil health, becoming the core driving factor for the continuous decline in the health index. Around point 1500, the contribution values of all lubricating oil monitoring parameters dramatically surge, ultimately triggering system failure, which perfectly corresponds with the fault time point recorded in Figure 5. After fault repair, the contribution values of each parameter rapidly return to baseline levels, and the system recovers to a relatively stable state. Around point 2000, the parameter degradation contribution values again exhibit a change pattern similar to that around point 500, and the green curve correspondingly shows a second round of the degradation process.

Notably, the degradation contribution of the moisture parameter maintains a high level throughout the entire monitoring cycle, but its feature importance score is the lowest (only 0.024). This indicates that although moisture content exhibits continuous fluctuation, due to its limited weight in the overall health assessment system, it does not produce a significant direct impact on the health degradation curve. This analysis method effectively identifies the key factors affecting lubricating oil health status, providing important theoretical support for fault diagnosis and preventive maintenance in the digital tribology era.

In conclusion, the SFTI-LVAE model demonstrates significant advantages in lubricating oil health assessment, accurately identifying early degradation while providing early warnings and precisely capturing abrupt points. By integrating the advantages of time series prediction and deep learning, it enhances the reliability and foresight of predictive maintenance, providing efficient and practical technical support for equipment health management in intelligent manufacturing.

4.4.3. Comparative Analysis

In the comparative analysis of health assessment models, this study selected alarm rate, false alarm rate, and early warning capability (EWC) as key indicators to comprehensively evaluate model performance. These metrics reflect the core requirements of models in practical applications: efficiently identifying anomalies (high alarm rate), reducing false alarms (low false alarm rate), and possessing early warning capability, thereby enhancing their practicality and reliability.

Based on measured data from the diesel engine, the study categorized the data into the following two types: first, 1440 data points collected on 5 April, representing normal operation periods; second, a 30 min period before and after the fault occurrence at 1:31 a.m. on 6 April (points 1516 to 1546), representing abnormal operation periods. This classification was based on field records and equipment operational characteristics, providing a reliable data foundation for the accurate identification of pre- and post-fault states and performance evaluation of the models.

The calculation formulas for alarm rate and false alarm rate are shown in Equations (17) and (18), which are used to evaluate the anomaly detection performance of the models [27].

$$TPR = \frac{TP}{TP + FN} \tag{17}$$

$$FPR = \frac{FP}{TN + FP} \tag{18}$$

where TPR represents the True Positive Rate and FPR represents the False Positive Rate; True Positives (TPs) denote the number of samples correctly detected as anomalous by the model; False Positives (FPs) represent the number of samples erroneously classified as anomalous when they are actually normal; False Negatives (FNs) indicate the number of anomalous samples not detected by the model; and True Negatives (TNs) represent the number of samples correctly identified as normal by the model.

Considering that health degradation is a gradual process, the degradation degree is quantified. Specifically, if a health decline trend occurs with a decrease exceeding 10%, it is considered an early fault trend, and this time point is designated as the fault acceleration point. The time interval between the fault acceleration point and the actual fault point characterizes the early warning capability. From the results, the SFTI-LVAE curve shows a continuous decline starting from point 637 at 91.08%, dropping to 80.92% at point 1128, exceeding the 10% threshold. The other experimental groups are as follows: the real data experiment shows a decline from point 635 to 1102, dropping from 91.74% to 81.73%; SFTI-MD shows a decline from points 1447 to 1472, from 85.36% to 75.10%; and Informer-LVAE shows a decline from point 621 to 961, dropping from 92.69% to 82.62%.

The specific calculation results for each model are detailed in Table 9.

Table 9. Evaluation results of anomaly detection for various models.

Data Interval	Data Category	Model	TPR (%)	FPR (%)	EWC (h)
Test set (5 April)	Predicted value	SFTI-LVAE	N/A	0	6.47
		SFTI-MD	N/A	0	0.73
		Informer-LVAE	N/A	3.40	9.25
	Actual value	LVAE	N/A	0	6.9
Fault interval (1516–1546 points)	Predicted value	SFTI-LVAE	96.67	N/A	N/A
		SFTI-MD	96.67	N/A	N/A
		Informer-LVAE	50.00	N/A	N/A
	Actual value	LVAE	90.00	N/A	N/A

The experimental results demonstrate that the SFTI-LVAE model successfully captures lubricating oil health degradation trends 6.47 h before fault occurrence, exhibiting an exceptional early warning capability (EWC). Its fault interval alarm rate (TPR) reaches 96.67%, with zero false alarms in healthy intervals. Combined with assessment curves that intuitively reflect health status variations, the model demonstrates a high reliability and accuracy. The LVAE model achieves a 90.00% alarm rate in real data evaluation, also with zero false alarms, showing highly consistent results that further validate the model's stability.

In contrast, although the SFTI-MD model performs well in terms of TPR (96.67%) and FPR (0), it fails to capture early degradation trends, with its assessment results abruptly transitioning from "healthy" to "faulty" within only 73 data points, which does not meet industrial requirements for early warning. While the Informer-LVAE model identifies early degradation trends, its TPR is only 50.00%, with a false alarm rate of 3.40%, indicating that single-trend prediction approaches have limitations in processing high-dimensional information and cannot meet the high precision requirements for health assessment.

The observed zero false alarm rate is not attributable to overfitting, but rather reflects the model's robust design and validation methodology. This performance is underpinned by the following two key architectural features: (i) the LVAE health assessment module employs unsupervised reconstruction learning exclusively on healthy baseline data, constraining the learned latent distribution to capture normal operational patterns while preventing the memorization of random noise or anomalous behaviors, and (ii) the SFTI forecasting component systematically separates trend and seasonal components from high-frequency disturbances through STL decomposition, providing cleaner, denoised inputs for reconstruction-based anomaly detection. To validate robustness beyond the primary test set, we conduct comprehensive sliding window evaluation across multiple contiguous time segments, treating each as an independent validation set trained solely on preceding healthy data. Consistently, detection metrics maintain zero false alarms across all segments. Furthermore, k-fold cross-validation on temporally disjoint blocks confirms that the healthy baseline model preserves a high specificity without compromising its sensitivity. These rigorous validation approaches demonstrate that the zero false alarm performance represents genuine model robustness rather than dataset-specific overfitting, though multi-site and cross-equipment validation studies are planned to further confirm generalization under broader operational conditions and domain shifts.

In summary, while statistical methods achieve high alarm rates and low false alarm rates, they are constrained by binary fault classification limitations. Simple trend prediction models suffer from a limited accuracy, compromising health assessment precision. The SFTI-LVAE model demonstrates significant advantages in multidimensional information extraction, early warning capability, and false alarm control, substantially outperforming comparative approaches through its integrated architecture and rigorous validation methodology.

5. Conclusions

This study proposes a health prediction model based on SFTI-LVAE and validates its effectiveness using diesel engine lubricating oil monitoring data. The model consists of the following two stages: trend prediction and health assessment. First, STL decomposition is employed to decompose time series into trend, seasonal, and residual components, combined with an Informer model optimized by FAN and TCN, achieving a "decompose–predict–aggregate" structure that effectively enhances its perception capability for trend changes, minor fluctuations, and abrupt variations, meeting the requirements for dynamic industrial health assessment. Second, by incorporating LSTM into the VAE model, recon-

struction errors are utilized to quantify the anomaly degree at each time step, accomplishing precise health status prediction.

To ensure the optimal model performance in practical applications, critical parameter configurations are defined based on model characteristics. The SFTI model processes dynamic trends and seasonal patterns through FAN, with the shuffle parameter set to True preventing overfitting, the LVAE encoding dimension optimized to five, and balancing reconstruction quality and weight parameters in the FAN module, enhancing attention to critical features. These parameter selections provide practical guidance for implementing similar systems in industrial environments.

To deeply explore the intrinsic mechanisms of lubricating oil health degradation, this study further develops feature importance and parameter contribution analysis methods. Based on a permutation importance strategy, the influence of each feature on model performance is quantified, and the contribution of each parameter to degradation at different time points is calculated through reconstruction errors, enabling fault source tracing for lubricating oil health degradation. The analysis results indicate that the dielectric constant possesses the highest feature importance score (0.311), serving as the core factor affecting lubricating oil health status, while the synergistic degradation of parameters such as density, viscosity, and temperature directly influences the fluctuation trends of health indices. This analysis method not only identifies key influencing factors, but also reveals the response characteristics of each parameter at different degradation stages, providing a scientific basis for fault source tracing and preventive maintenance.

The experimental results demonstrate that the SFTI-LVAE model successfully provides early warning 6.47 h before fault occurrence, achieving an alarm rate (TPR) of 96.67% with a false alarm rate of 0%. The predicted health indices are highly consistent with actual assessment results, showing an excellent early warning capability and prediction accuracy. In comparison, although the SFTI-MD model possesses an equally high TPR and low FPR, it lacks the ability to identify early degradation, only issuing warnings shortly before faults, exhibiting abrupt transitions from "healthy" to "faulty" states. While the Informer-LVAE model can detect early degradation, it is limited by the traditional Informer's capability in processing high-dimensional industrial data, resulting in an insufficient prediction accuracy with an alarm rate of only 50.00% and a false alarm rate of 3.40%, compromising the reliability of health assessment.

Through in-depth comparative analysis of the three models' performances across multiple typical time periods, SFTI-LVAE not only accurately predicts future equipment operational states, but also possesses powerful data representation and anomaly identification capabilities. Its assessment results highly correspond with the actual alarm time at 1:31 a.m. on 6 April, validating its practicality in industrial scenarios. This model provides a novel approach for lubricating oil wear state monitoring and offers reliable technical support for the predictive maintenance and online operation of industrial equipment.

When implementing in new environments, practitioners should follow a structured deployment approach: first, establish baseline healthy operational data for model training; second, customize parameter configurations according to the aforementioned guidelines; third, gradually integrate with existing maintenance protocols; and fourth, establish threshold strategies based on the specific operational requirements and risk tolerance of the application domain. The model's interpretability features, including parameter contribution analysis and feature importance assessment, enable maintenance engineers to understand degradation mechanisms and effectively prioritize maintenance actions.

Despite achieving satisfactory results, this study still has room for improvement in terms of generalizability. The experimental validation relies on data collected from a diesel engine lubricating oil system over six days, with only one confirmed lubrication

failure event. While the nine monitored parameters provide comprehensive health status characterization and the continuous monitoring encompasses a complete process of health degradation and post-maintenance recovery to a healthy state, providing sufficient training samples (8640 data points), future research can be extended to more equipment types (such as gas turbines and compressors), different lubricant formulations (various viscosity grades and additive formulations), and long-term degradation patterns spanning months to years. Furthermore, validation under different operating conditions, including various temperature ranges, load variations, and environmental conditions, constitutes future work plans that will enhance the model's robustness and generalization capability. Nevertheless, the unsupervised learning characteristics and modular design of the SFTI-LVAE framework enable it to adapt to different equipment types and operating conditions through retraining, providing scalability for broader industrial applications.

In conclusion, the SFTI-LVAE model significantly outperforms traditional SFTI-MD and Informer-LVAE models in comprehensively capturing data trends, enhancing early warning capabilities, controlling false alarm rates, and mechanistic analysis. It not only possesses an exceptional health prediction performance, but also achieves a deep understanding of lubricating oil health degradation mechanisms through feature importance and parameter contribution analysis, providing an efficient and feasible solution for equipment condition assessment and maintenance management in intelligent manufacturing environments. Future research directions include (1) incorporating heuristic optimization algorithms to improve parameter tuning efficiency; (2) expanding validation to diverse equipment types, lubricant formulations, and longer-term operational scenarios; (3) developing transfer learning strategies to enhance model adaptability across different industrial domains; and (4) integrating with existing maintenance management systems for seamless industrial deployment, thereby enhancing the model's generalization capability and practical value in complex industrial scenarios.

Author Contributions: Conceptualization, Y.C., K.Y. and L.W.; Methodology, Y.C.; Software, Y.C.; Validation, Y.C.; Formal analysis, F.Z. and L.W.; Resources, J.S. and L.W.; Data curation, J.S.; Writing—original draft, Y.C.; Writing—review & editing, H.L. and K.Y.; Visualization, H.L.; Supervision, F.Z.; Project administration, K.Y.; Funding acquisition, K.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Key R&D Program of China (2022YFB4300300).

Data Availability Statement: Due to project-related confidentiality requirements, the data supporting the findings of this study are not available at present. Public access may be provided in the future when conditions permit.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Johns-Rahnejat, P.M.; Rahmani, R.; Rahnejat, H. Current and future trends in tribological research. *Lubricants* **2023**, *11*, 391. [CrossRef]
2. Chokelarb, W.; Sriprom, P.; Permana, L.; Assawasaengrat, P. Assessment of overall remaining useful life of lubricants by integrating oil quality and performance. *Heliyon* **2024**, *10*, e37486. [CrossRef]
3. Han, Z.; Zhao, J.; Leung, H.; Ma, K.F.; Wang, W. A review of deep learning models for time series prediction. *IEEE Sens. J.* **2019**, *21*, 7833–7848. [CrossRef]
4. Siami-Namini, S.; Namin, A.S. Forecasting economics and financial time series: ARIMA vs. LSTM. *arXiv* **2018**, arXiv:1803.06386. [CrossRef]
5. Chandra, R.; Goyal, S.; Gupta, R. Evaluation of deep learning models for multi-step ahead time series prediction. *IEEE Access* **2021**, *9*, 83105–83123. [CrossRef]

6. Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 2–9 February 2021; Volume 35, pp. 11106–11115. [CrossRef]
7. Wang, C.; Wang, H.; Zhang, X.; Liu, Q.; Liu, M.; Xu, G. A Transformer-Based Industrial Time Series Prediction Model With Multivariate Dynamic Embedding. *IEEE Trans. Ind. Inform.* **2024**, *21*, 1813–1822. [CrossRef]
8. Liu, F.; Dong, T.; Liu, Y. An improved informer model for short-term load forecasting by considering periodic property of load profiles. *Front. Energy Res.* **2022**, *10*, 950912. [CrossRef]
9. Huang, H.; Wang, P.; Pei, J.; Wang, J.; Alexanian, S.; Niyato, D. Deep Learning Advancements in Anomaly Detection: A Comprehensive Survey. *arXiv* **2025**, arXiv:2503.13195. [CrossRef]
10. Chalapathy, R.; Chawla, S. Deep learning for anomaly detection: A survey. *arXiv* **2019**, arXiv:1901.03407. [CrossRef]
11. Li, D.; Chen, D.; Goh, J.; Ng, S.K. Anomaly detection with generative adversarial networks for multivariate time series. *arXiv* **2018**, arXiv:1809.04758.
12. Zeng, G.Q.; Yang, Y.W.; Lu, K.D.; Geng, G.G.; Weng, J. Evolutionary Adversarial Autoencoder for Unsupervised Anomaly Detection of Industrial Internet of Things. *IEEE Trans. Reliab.* **2025**, 1–15. [CrossRef]
13. Fortuin, V.; Baranchuk, D.; Rätsch, G.; Mandt, S. Gp-vae: Deep probabilistic time series imputation. In Proceedings of the International Conference on Artificial Intelligence and Statistics, Palermo, Italy, 26–28 August 2020; PMLR: Mountain View, CA, USA; pp. 1651–1661. [CrossRef]
14. Åström, O.; Sopasakis, A. Improved Anomaly Detection through Conditional Latent Space VAE Ensembles. *arXiv* **2024**, arXiv:2410.12328. [CrossRef]
15. Guo, Y.; Liao, W.; Liao, W.; Wang, Q.; Yu, L.; Ji, T.; Li, P. Multidimensional time series anomaly detection: A gru-based gaussian mixture variational autoencoder approach. In Proceedings of the Asian conference on Machine Learning, Beijing, China, 14–16 November 2018; PMLR: Mountain View, CA, USA; pp. 97–112. Available online: https://proceedings.mlr.press/v95/guo18a.html (accessed on 14 June 2025).
16. Bandara, K.; Hyndman, R.J.; Bergmeir, C. MSTL: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. *Int. J. Oper. Res.* **2025**, *52*, 79–98. [CrossRef]
17. Ye, W.; Deng, S.; Zou, Q.; Gui, N. Frequency Adaptive Normalization For Non-stationary Time Series Forecasting. *arXiv* **2024**, arXiv:2409.20371. [CrossRef]
18. Bai, S.; Kolter, J.Z.; Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv* **2018**, arXiv:1803.01271. [CrossRef]
19. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef]
20. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv* **2013**, arXiv:1312.6114.
21. Kingma, D.P.; Rezende, D.J.; Mohamed, S.; Welling, M. Semi-supervised learning with deep generative models. *Adv. Neural Inf. Process. Syst.* **2014**, *27*, 3581–3589. [CrossRef]
22. Komorska, I.; Puchalski, A. Condition Monitoring Using a Latent Space of Variational Autoencoder Trained Only on a Healthy Machine. *Sensors* **2024**, *24*, 6825. [CrossRef]
23. Samek, W.; Montavon, G.; Lapuschkin, S.; Anders, C.J.; Müller, K.R. Explaining deep neural networks and beyond: A review of methods and applications. *Proc. IEEE* **2021**, *109*, 247–278. [CrossRef]
24. Liznerski, P.; Ruff, L.; Vandermeulen, R.A.; Franks, B.J.; Kloft, M.; Müller, K.R. Explainable deep one-class classification. *arXiv* **2020**, arXiv:2007.01760.
25. Mohammed, M. Application and Evaluation of Multi-Sensor Fusion Technology in Diesel Engine Oil Analysis—A Review. Master's Thesis, Politecnico di Torino, Turin, Italy, 2021. Available online: https://webthesis.biblio.polito.it/20775/ (accessed on 21 June 2025).
26. Yang, K.; Hu, B.; Malekian, R.; Li, Z. An improved control-limit-based principal component analysis method for condition monitoring of marine turbine generators. *J. Mar. Eng. Technol.* **2020**, *19*, 249–256. [CrossRef]
27. Li, D.; Gao, H.; Yang, K.; Zhou, F.; Shi, X. Abnormal identification of oil monitoring based on LSTM and SVDD. *Wear* **2023**, *526*, 204793. [CrossRef]
28. *SAE Standard J300_202405*; Engine Oil Viscosity Classification. SAE International: Warrendale, PA, USA, 2024.
29. Jolliffe, I.T.; Cadima, J. Principal component analysis: A review and recent developments. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.* **2016**, *374*, 20150202. [CrossRef] [PubMed]
30. Sarmadi, H.; Entezami, A.; Razavi, B.S.; Yuen, K.-V. Ensemble learning-based structural health monitoring by Mahalanobis distance metrics. *Struct. Control Health Monit.* **2021**, *28*, e2663. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.