Operating Speed Effects on Vibration-Data-Driven Discrimination of Field-Used Lubricant Samples in an Internal Combustion Engine

Abstract

The concept of inferring lubricant condition from engine vibration is already established in the literature, particularly at fixed or narrow operating conditions. This paper presents a feasibility study for lubricant property prediction across varying operating speeds, investigating whether lubricant-related vibration signatures can be separated from operating-speed variability well enough for vibration to serve as a practical surrogate for direct lubricant condition sensing. Vibration was recorded from a petrol engine at six speeds under no load with five field-used lubricant samples, and discrimination between the samples, a necessary first step, was tested using engineered statistical descriptors and supervised machine-learning classification. Three feature representations, differing in whether lubricant information is assumed to sit in a few interpretable descriptors or to be spread across many, were evaluated under two strategies in which the test speed was either represented in or withheld from the training data. The samples were discriminable to a modest extent (accuracy up to 0.63 against a chance level of 0.20), indicating that exploitable lubricant-related information exists in a concentrated minority of the global feature set, although weak and diffuse across the full feature space. Discrimination varied with operating speed and thermal history of the engine, although per-speed estimates are descriptive. An auxiliary analysis identified the most contributing descriptors, which may be transferable unlike the trained models, and showed a descriptive tendency for confusion between samples of similar viscosity.

1 | Introduction

In internal combustion engines (ICEs), the lubrication system plays a vital role in reducing its frictional resistance and maximising its mechanical efficiency. It is often believed that the energy lost due to friction is what differentiates a good ICE from an average one [1]. In addition to friction reduction, an engine lubricant also performs several auxiliary tasks such as protecting the engine against wear, removing injurious impurities, holding gas leakages at an acceptable level and dissipation of frictional heat. During its service interval, an engine lubricant is subjected to a number of stresses, as well as to contamination by particles originating from wear and corrosion within the engine, combustion products, water and unburnt fuel [2]. After an engine lubricant has reached the end of its useful life upon being subjected to aforementioned degradation and depletion of additives present in the lubricant, its behaviour changes in an undesirable manner.

The conventional practice of engine lubrication maintenance typically involves a preventive action where the degraded lubricant is replaced after a fixed mileage or a time interval as per manufacturer's recommendation. Such practices can often lead to two detriments: over-maintenance or under-maintenance. Over-maintenance of engine lubricants can be quite disadvantageous economically and environmentally [3] whereas under-maintenance can result in more mechanical detriments such as increased component wear, reduced fuel efficiency or even engine failure [4] before next service interval. Although some high-end vehicle models have implemented predictive methods, most commercial automobile manufacturers still rely on conventional preventive maintenance albeit its drawbacks [5].

In monitoring internal and external health of ICEs, vibration analysis often provides a holistic view of their behaviour [6]. In practice, Engine Condition Monitoring (ECM) already uses vibration monitoring through a piezoelectric accelerometer mounted on the engine to detect engine knock [7]. Thus,

integration of a vibration-driven system to monitor lubricant condition is only practical given that such a method is feasible and reliable. This poses the question as to whether tribological differences between lubricants are measurable through a vibration signal, as a first step towards detecting their degradation.

In theory, it is established that lubrication has a direct effect on the damping mechanism of a vibratory system. Eq. 1-4 derives the generalised dynamical relationship between two surfaces under relative motion with a full-film lubricant layer in-between under the simplification that the lubricant behaves as a Newtonian fluid [8]. Such behaviour is mostly experienced under hydrodynamic lubrication regime when the sliding surfaces are sufficiently separated by engine lubricant while moving at a relatively high speed and withstanding a limited load. For example, motion between piston skirt and cylinder liner at high sliding velocities [1]. The same theory can be extended towards journal bearings, one of the most prominent lubricated parts of the engine which holds the crankshaft, a dominant source of vibration. The corresponding dynamical relationship is derived as shown in Eq. 5-7 from Petroff's Law [8]. Consequent theory proposes that as viscosity increases, the damping coefficient increases and in turn, attenuates vibration much better.

$$F = \tau A \tag{1}$$

$$F = \left(\mu \frac{v}{h}\right) A \tag{2}$$

$$F = \frac{\mu A}{h} v \tag{3}$$

The damping coefficient c can be expressed as:

$$c = \frac{F}{v} = \frac{\mu A}{h} \tag{4}$$

where F is the resistive shear force, τ is the shear stress, A is the area of contact, v is the relative velocity between two surfaces, h is thickness of lubricant layer, μ is dynamic viscosity of the lubricant.

$$T = (\tau A) R \tag{5}$$

$$T = \left(\frac{\mu R \omega}{d}\right)(2\pi R l) R = \frac{2\pi \mu R^3 l \omega}{d} \tag{6}$$

The rotational damping constant c_R can be expressed as:

$$c_R = \frac{T}{\omega} = \frac{2\pi \mu R^3 l}{d} \tag{7}$$

where T is the resistive torque acting on the crankshaft, R is the radius of the shaft, l is the length of the bearing, d is the bearing clearance.

Hydrodynamic lubrication is generally considered the ideal, low-friction form of lubrication [9]. However, for an engine, this does not always hold due to its cyclostationary dynamics. For instance, the piston ring and cylinder interface at the top and bottom dead centre crank positions where the piston speed is lowest typically experience boundary lubrication [1]. In this regime, bulk physical properties such as viscosity (Eq. 4,7) contribute weakly and mostly the dynamic behaviour is instead governed by the Coulomb damping defined by the surface properties and chemical properties of the surface film of the lubricant [9]. In contrast to reciprocating piston movement, the journal bearings undergo constant rotation where their damping depends on the operating speed and the load acting on the crankshaft. This behaviour was systematically studied by Stribeck [10], leading to the development of the classic

Stribeck curve (Fig. 1) which portrays the transition between boundary and hydrodynamic lubrication between two metal surfaces as a function of non-dimensional Hersey number H [11].

$$H = \frac{\mu N}{P} \tag{8}$$

for where N is the rotational speed and P is the pressure on the lubricated surface.

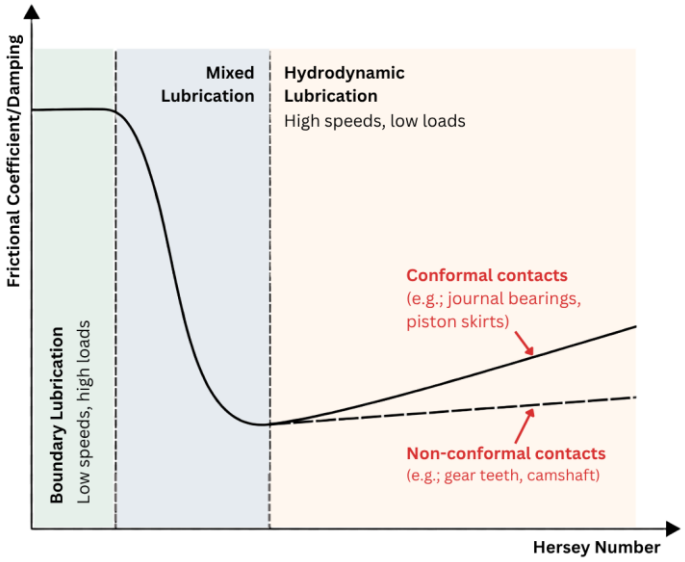


Figure 1: A typical Stribeck curve characterised by lubrication regimes and corresponding operating conditions *[12].*

Valve stems, rocker arms and crankshaft timing gears are also expected to experience boundary lubrication due to their slow relative movement [1]; although their magnitude of contribution is small compared to that of the pistons and journal bearings, it is not negligible when it comes to vibration measurement. Therefore, it is essential to develop a comprehensive understanding of the components affected by lubrication and the transitions between boundary and hydrodynamic lubrication within the engine, in order to extract lubricant properties from a vibration signal across a wide range of operating conditions in an ICE.

As the theory suggests, a number of past researchers have studied the effect of changes in engine lubricant properties on vibrations and concluded that there is an increase in vibration magnitude as lubricant viscosity decreases [13], [14], [15]. The preliminary study by Santana et al. [13] and subsequent research [16] has used laboratory controlled lubricant conditions for over 600 continuous hours to simulate lubricant degradation, and performed descriptive analyses on the trend of root mean square (RMS) value of crankshaft and engine block vibration separately. This was conducted at a fixed engine speed of 4500 revolutions of minute (RPM) under full dynamometer load. A similar study was done by Wu et al. [17] to analyse several broadband and band-wise time-domain statistical descriptors of vibration signal and concluded the feasibility of non-intrusive vibration monitoring of wear status of cylinder liner of a marine engine; however, this experiment was conducted at three coupled load-speed pairs using three samples of lubricants to establish correlations. Ji et al. [15] has studied how combustion status and engine block vibrations are coupled at controlled lubricant temperature conditions and concluded that frequency components higher than 20 kHz are more sensitive to lubricant damping by high lubricant viscosity at low temperatures for a fixed operating speed.

Although it has been established there is a direct correlation between vibration characteristics and lubricant viscosity, quite possibly via the damping mechanism, in reality, physical properties of a lubricant alone cannot quantify the damping effect. Chemical properties also govern the tribological

behaviour in boundary lubrication regime through the additive molecules in the lubricant that attach themselves on to the solid metal surface to form an absorbed film on the metal surface [1]. Total base number (TBN) is a quantity used to measure the alkalinity of engine lubricant, which reflects its ability to neutralise corrosive acidic byproducts of combustion via its additives. A study by Nagy et al. [18] has concluded that the degradation trend of TBN is statistically more consistent across different driving cycles compared to viscosity. A similar conclusion was drawn by Nassef et al. [19] regarding the statistical consistency of oxidation-related parameters such as TBN compared to wear metal indicators under varying thermodynamic conditions. A study by Omiya et al. [3] has also revealed that TBN has a strong correlation with mileage (degradation indicator of conventional preventive maintenance) compared to viscosity. Lukášik [20] has proposed a regression model built using in-service test data to predict remaining useful life (RUL) of engine lubricant using TBN and anti-wear additives measurements claiming them to be the best indication of lubricant's chemical lifespan. Even though prior research suggests that chemical properties are better indicators to quantify engine lubricant degradation, its effect on vibration may not be as discernible as the effect of changes in viscosity in hydrodynamic regime. Nevertheless, vibration may still reflect tribological changes attributable to chemical changes to some extent as found by previous research that has studied the variation of vibration characteristics with TBN [16], [17].

According to recent research on condition monitoring and predictive maintenance of the lubrication system in automobiles, the existing technology can be generally classified into two approaches. A study done by Omiya et al. [3] used an indirect method where the lubricant condition is estimated using vehicle information and driving data. Certain automobile manufacturers also use similar methods to monitor engine lubricant by tracking driving cycles, engine loads and temperature cycles [5]. Research by Chen et al. [21], [22] has proposed predictive algorithms that use more direct real-time measurements such as dielectric constant, viscosity, density, temperature, wear particulate content and moisture content fused from an online multi-sensor loop. They have identified dielectric constant as the most critical parameter for assessing lubricant health, which explains why it's a well-established technique in other industrial systems [5]. However, not every existing approach is without shortcomings despite their success. Indirect approach does not reflect real engine lubricant behaviour but relies on operational history and fails to estimate abnormalities such as contamination. Nonetheless, indirect estimation almost assumes no hardware cost compared to direct sensing of physical and physiochemical properties. Dedicated hardware used for direct sensing should be met with durability and robustness to withstand harsh engine environment, subsequently leading to high initial cost despite exceptionally robust and accurate outcome obtained in previous work through deep learning methods [21], [22].

This motivates examining whether engine vibration carries information that distinguishes lubricant samples, and how that information depends on operating speed. Most prior studies that have examined this relationship in automobile engines have done so at a fixed operating speed [13], [14], [15], [16], leaving open the extent to which lubricant-related vibration signatures persist, or can be reliably extracted, across the range of speeds an engine experiences in practice. Given that engine operating speed is itself a fundamental determinant of engine vibration, it is plausible that operating speed is a dominant source of variability in engine vibration that can obscure or interact with lubricant-related information in a vibration signal (Eq. 4,7,8). Moreover, the vibration signature of an engine can be described as quasi-deterministic due to the presence of combustion transients [23] and mechanical noise [24]. If the mechanical noise generated is directly impacted by the tribological conditions as concluded by Ji et al. [15], then operating speed and lubricant effects may not be cleanly separable through an analytical relationship. This motivates a data-driven approach capable of empirically quantifying how much of the variability in engineered vibration data is attributable to lubricant condition.

This paper investigates the relative influence of lubricant sample and engine operating speed on statistical descriptors extracted from vibration signals and examines how this operating-speed dependence affects the practical discriminability of different field-used lubricant samples through supervised feature representations. Machine-learning (ML) models were used as primary analytical tools to evaluate discriminative capability. Specifically, this study has three objectives: (1) quantification of relative influence of lubricant sample at different engine operating speeds on a broad set of engineered vibration descriptors; (2) evaluation of robustness of different feature representations formulated using vibration descriptors; and (3) identification of interpretable vibration descriptors that retain lubricant-discriminative information despite the dominant influence of engine operating speed. This paper is a feasibility study: discriminating five field-used lubricant samples is used throughout as the application through which these operating-speed and feature-representation effects are examined, as a necessary first step towards predicting lubricant properties and degradation in future work.

The study was conducted on an automobile petrol engine mounted in a laboratory environment at six different regulated operating speeds under no-load condition. The paper begins with a detailed overview of the experimental setup including equipment, materials and the data acquisition method in Section 2. The data preparation methodology in Section 3 begins by validating the acquired signal in Section 3.1, followed by the extraction of features, specifically statistical descriptors of the vibration signal in Section 3.2. Section 4 then presents a preliminary data analysis using the Analysis of Variance (ANOVA) test. Section 5 outlines the ML pipeline designed for the primary analysis. Lastly, the paper ends with the discussion and interpretation of results from the primary analysis in Section 6, concluding with remarks on the practical considerations to be considered when applying this novel method.

2 | Experimental Setup & Data Acquisition Method

The engine used in this study was a Nissan inline four-cylinder, four-stroke, naturally aspirated, spark-ignition engine with a displacement of 1597 cc, a bore of 76 mm, and a stroke of 88 mm [25]. The test rig was installed on an isolated laboratory bench with an open-loop cooling system as shown in Fig. 2. Wu et al. [17] found lubricant condition detection through vibration is most effective at the highest tested speed-load point. As vehicles rarely sustain extreme operating points in real-world use, a method validated where data can be practically collected is more representative of everyday operating conditions. Therefore, in this study, the engine is disconnected from dynamometer in order to target low to moderate operating speeds under the least mechanically pronounced condition, that is, no load.

Let e_x, e_y, e_z be a right-handed orthonormal coordinate frame such that,

$$e_x \parallel \text{crankshaft axis}, e_z \parallel \text{piston motion}, e_y = e_x \times e_z.$$

Due to the imbalance of reciprocating masses, the cranktrain contributes significantly to the overall engine vibration compared to crankshaft, camshaft, valvetrain and other auxiliary components [14], [26]. Consequently, these mechanical vibrations are predominantly confined to e_y and e_z directions. Conversely, vibrations along e_x primarily result from combustion pressures and torsional-axial coupling. Because the number of accelerometer sensors available was limited to two, e_y and $e_z$were strategically selected to cross-isolate distinct mechanical phenomena. Specifically, e_y sensor was mounted on the engine block (near the crankshaft) to capture crankshaft dynamics and piston slap vibrations while the e_z sensor was mounted on the engine head to capture valvetrain and camshaft vibrations. The sensor placement strategy was subsequently validated using spectral coherence in Section 3.1. Insight into the reinforcing structure of the engine block was drawn from an article

published by the manufacturer [25], which guided an intuitive understanding of the likely nodal points, antinodes, and optimal mounting locations attributed to engine block in e_y direction.

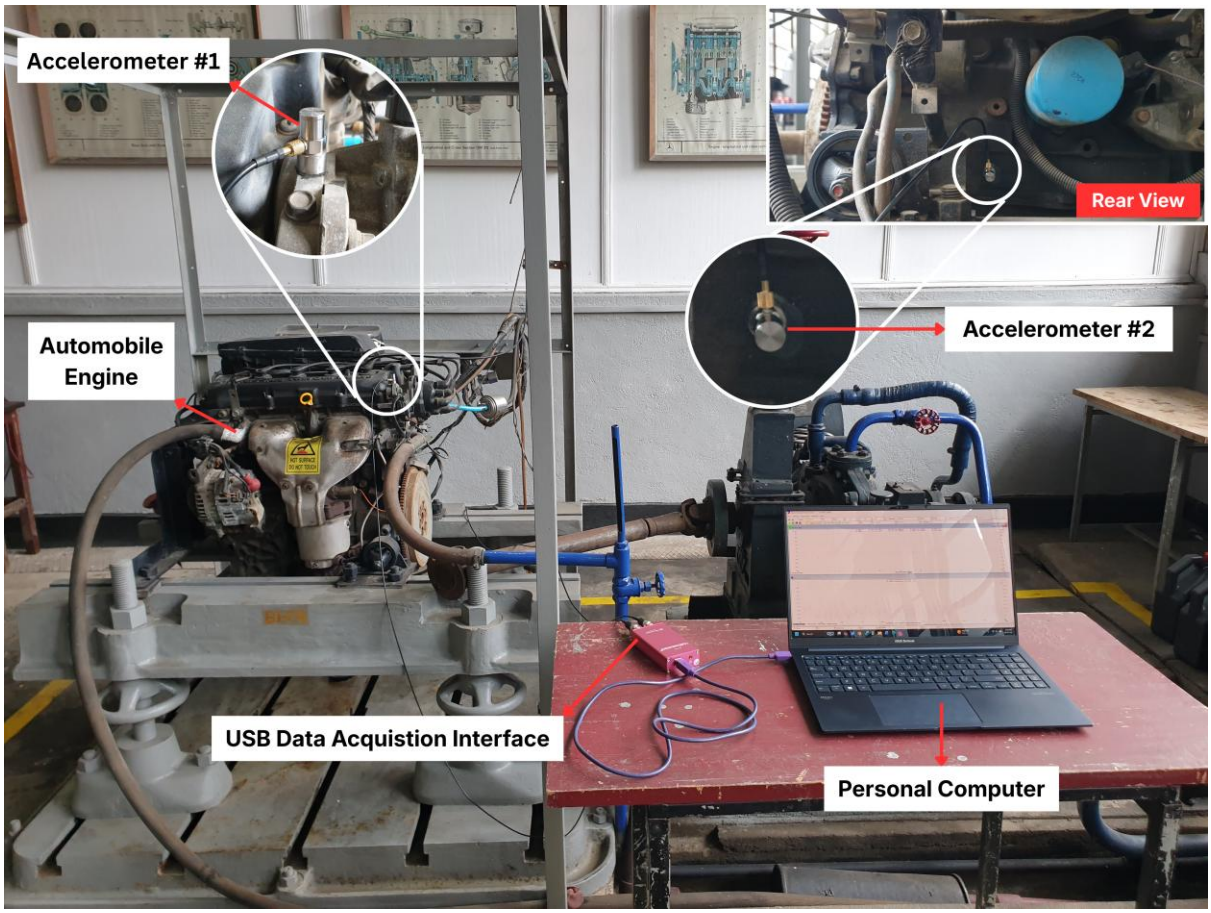


Figure 2: Experimental Setup including Data Acquisition System

Vibration data were collected using two CA-YD 181 integrated electronic piezo-electric (IEPE) accelerometers, featuring a nominal sensitivity of 10 ± 0.5 mV/g and frequency response up to 10 kHz. The accelerometers were interfaced with a 24-bit dual-channel VT-IEPE-2G05 USB data acquisition (DAQ) unit, which operates at 48 kHz sampling rate and 0.03 Hz-23 kHz bandwidth, satisfying the Nyquist-Shannon sampling theorem for full frequency range of the accelerometers. Signal conditioning, including constant-current power supply and high-pass filtering, was handled natively by the internal circuitry of the IEPE sensors and the DAQ interface. To ensure rigid coupling, the accelerometers were mounted via custom-made bases with a thin layer of epoxy adhesive. The alternative options were magnetic mounting, which exhibits poor transmissibility for frequencies above 3 kHz, and stud mounting, which requires destructive drilling of the lab test rig [27]. Note that the accelerometer mounting bases were not removed or remounted during the entire course of the experiment.

The vibration signatures were recorded for five distinct lubricant samples listed in Table 1 at six discrete engine speeds: 850, 1100, 1300, 1600, 1850 and 2100 RPM. The crankshaft rotational speed was measured through the flywheel using a Kimo CT50 optical tachometer. Each data acquisition interval lasted approximately 30 seconds. To ensure the dynamic stability of the initial cold start, the engine was first operated at 850 RPM 10 minutes prior to the first measurement. For subsequent speed increments, the system was allowed to stabilise for 5 minutes before the next measurement. Upon reaching the

maximum measured speed, 2100 RPM, the engine was throttled down to idle speed and turned off. To capture the influence of thermal history of the engine from a baseline hot-start condition, a second sweep of repeated measurements under identical operating speeds was taken immediately afterwards.

It is known that the temperature-viscosity relationship of an engine lubricant is a crucial factor that affects lubricant performance. The Society of Automotive Engineers (SAE) has established a viscosity grade system that categorises all engine lubricants based on their viscosity at specific temperatures for different engines and climates. The lubricants used for this experiment include only SAE 10W-30 and SAE 15W-40 samples as listed in Table 1. Testing was conducted at Peradeniya, Sri Lanka, where the ambient temperature was about 25-30 °C at the time of the experiment. Previous research has studied the effect of temperature on engine lubricants; these temperature controlled experiments have concluded its distinct effect on engine vibration [14], [15]. In this experiment, however, lubricant temperature was not treated as a control variable and was deliberately left uncontrolled, allowed to fluctuate naturally to mimic real-world engine behaviour. In such cases, the lubricant temperature depends on operating speed [28], ambient conditions [20] and time elapsed in operation [2]. Nonetheless, two sequences of measurements for hot- and cold-starts were taken across all operating speeds and lubricant samples as mentioned earlier.

Table 1 details the operational histories and viscosity information of the five lubricant samples used in this study. Each sample had been exposed to distinct degradation processes because each of them was sourced from in-service field operation across five different vehicles. Note that on-lubricant mileage is not a good representation of viscosity change, plausibly due to competing mechanisms that alter lubricant viscosity. Processes such as polymerisation and the accumulation of soot and wear particles can increase the lubricant viscosity, whereas fuel dilution and coolant contamination can reduce the lubricant viscosity during field operation [29]. All information except viscosity was documented by the lubricant change service provider. The original kinematic viscosity was retrieved from the lubricant manufacturer's technical data sheet, and the current viscosity was measured using a Redwood No. 1 Viscometer as per the IP 70 standard for comparison. Three measurements were taken for each lubricant sample at 40 °C, and the average Redwood seconds was calculated before converting them to kinematic viscosity. To minimise cross contamination of lubricant samples within the laboratory engine test rig, precautions were taken; the existing lubricant was allowed to drain completely from the sump before refilling it with the next sample. The lubricant samples were run in the order given below in Table 1, from A to E over a course of two consecutive days. The accelerometer sensors were never removed and remounted during this period.

Table 1: Properties of the Lubricant Samples Used

Sample	**Lubricant change interval / km**	**Viscosity grade**	**Original kinematic viscosity at 40° C / cSt**	**Current kinematic viscosity at 40° C / cSt**
A	3440	10W-30	74.0	84.08
B	3760	10W-30	74.0	92.24
C	4200	15W-40	104.0	85.98
D	5000	15W-40	113.2	76.16
E	5500	10W-30	64.0	122.00

3 | Data Preprocessing

3.1 Signal Validation

3.1.1 Mounting direction validation using spectral coherence

In precedence to accelerometer sensor mounting reasoning explained in Section 2, it is important that sensors are placed close to sources of vibration that are affected by lubricant condition (e.g., journal bearings, piston skirt) for maximum transmissibility. In such cases, cross-channel coherence and multi-dimensional interpretation are used to distinguish global, fundamental excitation events that are not local structural vibrations attributed to mounting location [6], [30]. A high coherence at a specific frequency indicates that the transfer path from the excitation source to accelerometer is clean and uncompromised by local resonance. In the preliminary stages of the experiment, sensors were mounted on each axes pair successively and a 5 second vibration signal was recorded for each pair to perform a cross-spectral coherence analysis. To find the pair of axes most representative of fundamental kinematics of the engine, such as crankshaft rotation and piston movement which occur at low frequencies, those measurements were low pass filtered to isolate 0-500 Hz band before calculating the mean coherence.

Table 2 shows the mean spectral coherence between each pair of mounting axes. Welch's method [31] was used with 50% overlap and a frequency resolution of 1.5 Hz for the coherence estimation. Therefore, an estimation bias where finite data produces non-zero coherence even though two signals are linearly unrelated, is expected. Hence, statistical significant threshold is determined at 95% confidence ($p < 0.05$) level for each sensor pair based on number of segments averaged to calculate the coherence. The corresponding thresholds and the fraction of signal above the 95% confidence level noise floor are also presented in Table 2. It can be seen that $y - z$ sensor pair measures the most coherent fundamental engine dynamics as anticipated by the theoretical formulation in Section 2. Fig. 3 shows the cross spectral coherence between $y - z$ sensor pair across the 0-500 Hz band. Highest coherence was observed at order-based frequencies as indicated in Fig. 3, which means this sensor pair effectively captures crankshaft locked events below 500 Hz.

Table 2: Cross-spectral coherence between three mounting directions

Sensor axes	Mean coherence	Threshold at $\alpha = 95\%$	Fraction above threshold/noise floor
$x - y$	0.360	0.206	0.533
$x - z$	0.421	0.193	0.677
$y - z$	0.552	0.193	0.832

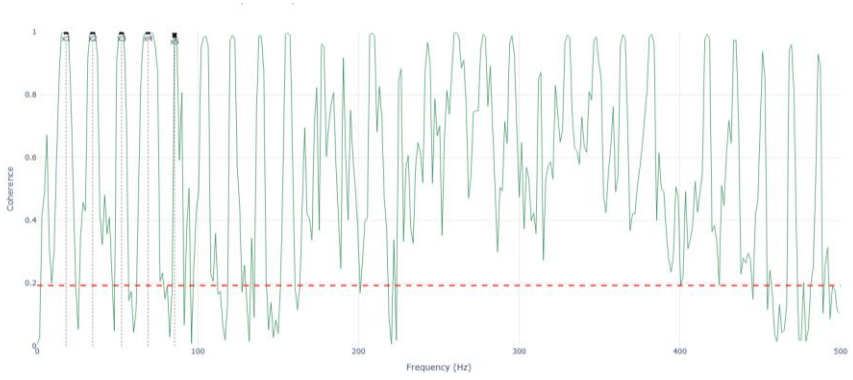


Figure 3: Cross-spectral coherence between y-z mounting directions with 95% confidence noise floor

The drawback of conventional spectral coherence estimation is that it only captures the linear relationship between two simultaneous signals, which may not capture mechanical nonlinearities such as impact, rubbing and varying contact stiffness. A complementary cross bicoherence analysis was done as a sanity check to investigate the quadratic nonlinear phase coupling between frequency components; $(f_1, f_2) \rightarrow (f_1 + f_2)$ to be specific. By visual inspection, it can be observed that $y - z$ pair exhibits strongest nonlinear coupling in frequencies below 500 Hz. Rather than a formal statistical test, only a qualitative conclusion was drawn from this analysis. Nevertheless, all observed regions of high coherence shown in Fig. 4 exceed the approximate noise floor (≈ 0.004) for all sensor pairs. Having supported that $y - z$ pair is the most representative mounting directions; they were used for all subsequent data collection and analyses performed in this study.

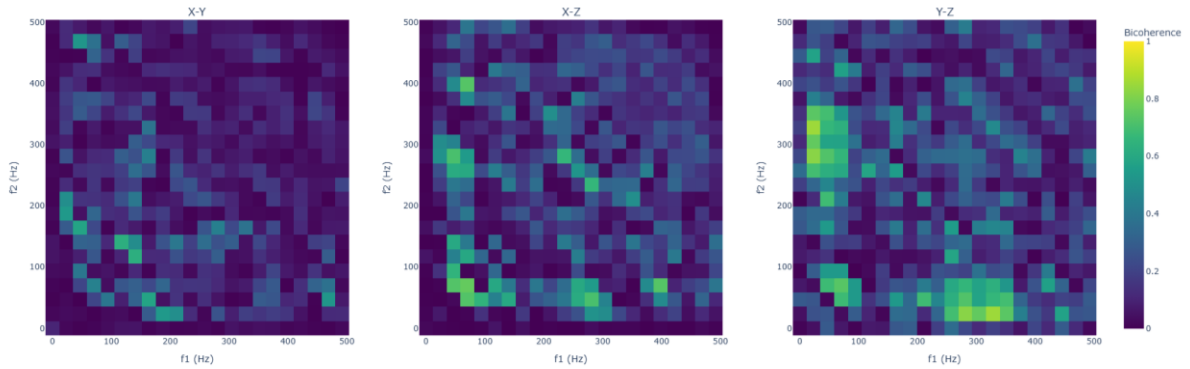


Figure 4: Quadratic nonlinear cross-spectral bicoherence map for all sensor pairs

3.1.2 Demodulation of high frequency signal

Typically, the low frequency components of ICE vibration are associated with the kinematics and periodic mechanical events of the engine, whereas stochastic or quasi-deterministic high frequency components arise from phenomena such as combustion transients, friction, metal impact ringing (resonance) and vortex shedding of intake air. Most of the principal low frequency events, including combustion pulses, valve opening and closing, and piston slap are synchronised with crankshaft position rather than a fixed frequency. Therefore, these events provide periodic mechanical excitation to the engine structure and can excite resonances in engine block [15] and engine head [32]. In particular, the engine block and head are of primary interest in the present study because they provide the structural path through which these periodic events are manifested as measurable vibration through accelerometer sensors.

While the cross-coherence analysis (0-500 Hz) corroborates the consistency of selected mounting location pair with the fundamental, crank-synchronous components of vibration, low frequency response does not represent the structural resonance occurring at high frequencies. Through two independent combustion and motor-driven tests, Ji et al. [15] have concluded that vibration components below 8 kHz are induced by combustion processes, whereas frequency components above 8 kHz (up to 25 kHz) are primarily influenced by inertia forces. Since the accelerometers used in the present study capture frequency components up to 10 kHz, it is plausible that high frequency components of the measurements may contain information associated with the underlying lubricant sample through the relationship between lubricant influenced low-order engine events and high frequency structural modes.

This possibility was investigated using envelope analysis, which provides a means of demodulating high frequency vibration to identify lower frequency information carried by a selected frequency band. In this approach, a frequency band is first isolated using a band-pass filter, after which the Hilbert

transform is applied to obtain the analytic signal. The magnitude of the analytic signal represents the amplitude modulation carried by the high frequency component. Envelope analysis was performed over the fixed frequency bands: 0-0.5, 0.5-1, 1-1.5, 1.5-2.7, 2.7-3.5, 3.5-5, 5-7 and 7-10 kHz for all samples. A fourth-order Butterworth band-pass filter was used to isolate each frequency band while minimising bandpass ripple. The resulting envelope spectra were expressed in terms of order amplitudes, which quantify the vibration components occurring at fixed multiples of crankshaft rotation frequency. The mean order amplitudes obtained across the lubricant samples were then compared between different frequency bands.

Fig. 5 shows the variation of mean order amplitudes of the envelope spectra, expressed as a ratio to the corresponding order amplitude of the 0-0.5 kHz baseband, at different engine order frequencies. The resulting trend showed a pronounced roll-off of order amplitudes up to approximately 1.5 kHz, which may represent natural spectral decay of low frequency events. This was followed by a recovery at higher frequencies above 5 kHz which is consistent with the excitation of structural resonances by discrete mechanical events. Notably, as plotted in Fig. 5, the ratios for the third- and fifth-order components in the 5-10 kHz bands are well above 1, in contrast to the decay seen at lower frequencies. This suggests that the high frequency vibration is modulated at engine orders rather than being flat broadband noise, although it is a qualitative observation, since the ratio has no fixed reference value. The observation was further supported by a simplified fixed bandwidth analogue of Antoni's kurtogram [33], which identified 5-10 kHz band to have the highest spectral kurtosis. The preceding envelope analysis provided a basis for investigating whether such information is influenced by the lubricant sample and its associated tribological behaviour in the present study.

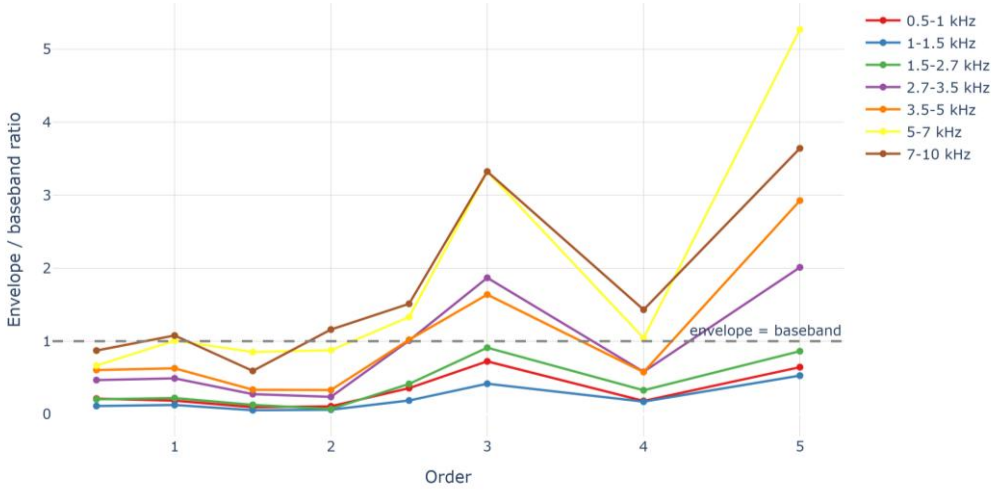


Figure 5: Normalised envelope order amplitudes across 7 fixed frequency bands (excluding baseband 0-0.5 kHz) versus rotational orders.

3.2 Extraction of Statistical Descriptors

While 30-second time series vibration recordings can be used in their raw form, they almost always require extensive pre-processing and feature extraction to be useful for data-driven modelling. Furthermore, pre-processing also improves noise attenuation and addresses data limitations such as redundancy, invariance and sparsity [34]. Recurrent Neural Networks (RNNs) and their variants (e.g., LSTM) are explicitly designed to handle sequential data such as time-series vibration signals or their fast Fourier transforms (FFTs). However, deep recurrent models typically require substantially larger number of independent training examples than what is available in this study to reliably estimate their

parameters without overfitting. For the same reason, deep learning models are rarely used in vehicle predictive maintenance [35]. Consequently, this study utilises and studies statistical descriptors extracted from the raw vibration signals in the subsequent analyses, which also offers the added benefit of low inference cost should the approach later be extended to on-board prediction [6]. Since this is a feasibility study, classical algorithms such as Support Vector Machines (SVMs) were used to test whether different lubricant samples are separable through vibration data.

Each vibration measurement captures several combustion cycles over its 30 second duration at the operational engine speeds. Despite the inherent periodicity of the system, minor combustion variations and micro-drift in operating RPM and temperature introduce subtle short-term variability to the measurements under nominally fixed operating conditions. To mitigate the constraints of a limited independent dataset, each 30-second measurement was partitioned into five non-overlapping, 5-second segments, yielding five distinct feature extractions per independent sample. For SVMs, this segmentation provides more points to shape the decision boundary, helping the model to generalise within a run's inherent variability.

Fig. 6 illustrates a summary of all the candidate statistical descriptors extracted from the raw signal. Prior to extraction of full-band statistical descriptors, a 10 kHz low-pass filter was applied to exclude the frequency range affected by nonlinear amplification due to the sensor's own resonance that distorts the structural vibration data. This introduces the limitation of this experiment not capturing information above 10 kHz. Candidate feature pool included six time-domain statistical descriptors: RMS, kurtosis, crest factor, zero crossing rate (ZCR), skewness and shape factor, which are comparatively inexpensive to compute and also correlate well with broad fault classes in automobiles [6]. In addition, two full-band frequency-domain statistical descriptors: spectral centroid and spectral entropy were incorporated into the candidate feature pool. To capture the changes in local sub-bands, band kurtosis, band spectral centroid and band energy percentages were extracted for 8 fixed frequency bands: 0-0.5, 0.5-1, 1-1.5, 1.5-2.7, 2.7-3.5, 3.5-5, 5-7 and 7-10 kHz since different frequency ranges are not equally informative as corroborated by previous research [15], [17], [32]. The preceding bands were chosen by visual inspection of the power spectral density (PSD) of the full-band signal in order to isolate visible resonance humps in both e_y and e_z spectra. Relative band energy, which is sub-band energy normalised with full-band (0-10 kHz) energy, quantifies the energy redistribution among the aforementioned fixed frequency intervals excluding the dominant effect that operating speed has on absolute band energies. Furthermore, amplitudes of integer and fractional (half-orders are informative in four-stroke cycles) orders: 0.5×, 1×, 1.5×, 2×, 2.5×, 3×, 4×, and 5×, in baseband and envelope frequency bands were included in the candidate feature pool as well. Each data sample is now represented by 208 statistical descriptors (104 from each e_y and e_z sensor) rather than a 5-second time-series (240,000 data) signal.

Because the FFT of vibration signals from a complex system such as an ICE is considerably jagged, Welch's method [31] was used to compute the PSD of the signal and extract frequency-domain statistical descriptors. The raw 5-second signal chunk was further split into 1-second segments with 50% overlap before applying Hann window and taking the average of their periodograms consecutively to mitigate spectral leakage. Certain statistical descriptors (e.g., spectral centroid, spectral entropy and relative band energy) require a smooth low variance PSD estimate as they involve integrating power over a range, and a jagged plain FFT would make that integral unreliable. In contrast, assuming the 5-second signal is near stationary, computation of baseband order amplitudes benefits from the plain single-sided FFT to identify sharp harmonic peaks (order amplitudes). However, envelope derived order amplitudes have a comparatively low signal-to-noise (SNR) due to narrowband Rayleigh distribution of envelopes, which are inherently high variance [33]. Hence, envelope spectra would benefit from the averaging process of Welch's method. It should be noted that order extraction assumes nominal operating speed due to absence of a tach channel.

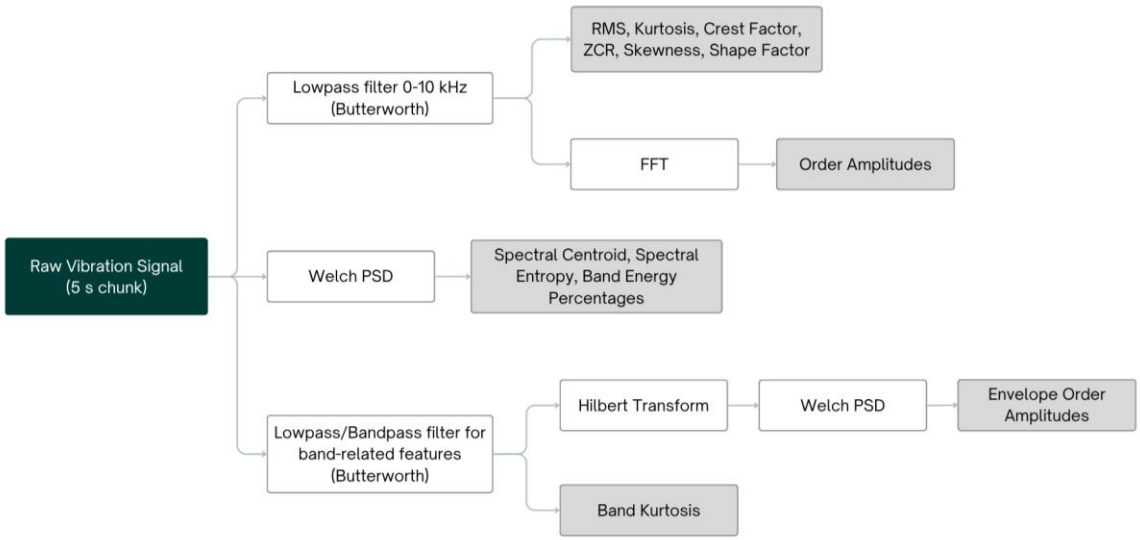


Figure 6: Extraction of statistical descriptors for the candidate feature pool

4 | Preliminary Analysis: ANOVA

4.1 Type II Factorial ANOVA

ANOVA, that is a special case of the general linear model, is used to test the effect of four categorical independent variables in the dataset: operating speed (RPM), sensor axis, lubricant (oil) sample and starting condition (hot- or cold-start) on each statistical descriptor extracted. As there is an equal number of data across the above four variable combinations, the balanced Type II factorial design will be used here [36]. Alternative options available include Type I and Type III factorial ANOVA, which better suit sequential and unbalanced designs respectively. To account for pseudo-replication arising from within-run 5-second signal chunks, all the data with same Oil × RPM × Condition × Axis combination was collapsed to its mean prior to fitting the ANOVA, in order to ensure residual degrees of freedom reflected independent experimental runs rather than correlated sub-samples.

Let $y \in \mathrm{R}$ be a vibration feature. Then β_i quantifies i parameter's contribution to that feature, $\beta_{i_1:i_2}$ quantifies the joint interaction among the parameters $i_1$and i_2, and ϵ is the residual contribution apart from main and joint interaction effects. The higher order interaction terms are not considered in the following model due to their lack of direct physical interpretability.

$$y = \beta_{\text{oil}} + \beta_{\text{RPM}} + \beta_{\text{start}} + \beta_{\text{axis}} + \beta_{\text{oil:RPM}} + \beta_{\text{oil:start}} + \beta_{\text{oil:axis}} + \beta_{\text{RPM:start}} + \beta_{\text{RPM:axis}} + \beta_{\text{start:axis}} + \epsilon$$

To explain the effect of each parameter, three metrics were used: F-statistic measures the ratio of strength of separation between different categories compared to separation within the same category; p-value denotes the probability of random variation generating the computed additional variance of that parameter; partial η^2 measures the fraction of attribution of a specific main effect or an interaction effect after controlling for all other model terms. These metrics were computed for all the statistical descriptors in the feature pool to test whether difference between lubricant samples are detectable through the vibration signal.

Table 3: Effect of the envelope-extracted statistical descriptors (features) on lubricant-related information

	Mean F-statistic (Oil)	**Mean F-statistic (Oil:RPM)**	**Mean partial η^2 (Oil)**	**Mean partial η^2 (Oil:RPM)**
Excluding envelope features (n = 80)	6.35	2.08	0.23	0.35
Including envelope features (n = 208)	3.89	1.82	0.16	0.33

Table 4: Comparison of lubricant related effects on statistical descriptors by ANOVA

Model term	**Mean F-statistic**	**Median p-value**	**Mean partial η^2**
β_{oil}	6.35	< 0.01	0.23
β_{RPM}	145.85	< 0.01	0.66
β_{start}	3.41	0.20	0.04
β_{axis}	4742.86	< 0.01	0.70
$\beta_{\text{oil:RPM}}$	2.08	0.04	0.35
$\beta_{\text{oil:start}}$	1.88	0.12	0.10
$\beta_{\text{oil:axis}}$	3.89	0.11	0.14
$\beta_{\text{RPM:start}}$	2.53	0.04	0.15
$\beta_{\text{RPM:axis}}$	40.09	< 0.01	0.54
$\beta_{\text{start:axis}}$	1.57	0.31	0.02

Recall the trend of normalised envelope amplitudes observed in Fig. 5. However, Table 3 indicates that envelope-derived order amplitudes exhibit less robust lubricant specific information as evidenced by ANOVA metrics. Consequently, statistical descriptors extracted from envelope analysis were excluded from the detailed ANOVA test and the primary study presented later. Each independent data sample is thus provisionally represented by 80 extracted features, 40 from each axis.

Table 4 shows the summary of mean F-statistic and mean partial η^2 for main effect terms and second-order interaction effects. Since operating speed and accelerometer mounting axis are the dominant physical determinants of vibration independent of lubricant sample, RPM and axis produce the highest effects (η^2 = 0.66 and η^2 = 0.70 respectively). Notably, the oil:RPM interaction is statistically significant ($p < 0.05$) in 52.5% of the features, alongside the main lubricant (oil) effect in 70%. This suggests that lubricant sample separability through vibration features is not uniform across all speeds but rather depends on the operating speed at which they are measured.

Since the temperature of lubricant was not treated as a control variable, two independent engine runs, a relative hot- and cold-start for every lubricant sample and operating speed combination were taken to capture variation associated with thermal history. According to ANOVA metrics in Table 4, the main effect due to hot- or cold-start is small (η^2 = 0.04) and explains only 0.42% of the total variance in ASCA (Table 5). However, its interaction with RPM, that is start:RPM, is statistically significant. This suggests that the difference between hot- and cold-start runs is not the same at all operating speeds. Thus, while starting condition does not constitute a dominant source of variance when averaged across the full operating range, its contribution may become important when examining the RPM-dependent behaviour of lubricant discrimination. Therefore, hot- and cold-start measurements will be retained together and treated as independent data in the primary study while their operating speed dependent behaviour will be revisited separately in Section 6.2.

4.2 ANOVA Simultaneous Component Analysis

A complementary ANOVA Simultaneous Component Analysis (ASCA) was performed to test whether this signal constitutes a coherent multivariate pattern in whole feature space in addition to the analysis in Section 4.1, which indicated a sufficiently detectable mean lubricant effect when one univariate ANOVA is fit per feature. ASCA instead decomposes the entire feature matrix X into a sum of effect matrices $\sum X_i$ per each parameter and joint interaction i and a residual E. Each feature is scaled to have zero mean and unit variance before decomposition. The higher order interaction terms were not considered in the following model due to their lack of direct physical interpretability

$$X = X_{\text{oil}} + X_{\text{RPM}} + X_{\text{start}} + X_{\text{axis}} + X_{\text{oil:RPM}} + X_{\text{oil:start}} + X_{\text{oil:axis}} + X_{\text{RPM:start}} + X_{\text{RPM:axis}} + X_{\text{start:axis}} + E$$

In this analysis, each matrix's effect is explained by three metrics: Sum of squares (SSQ) measures how much reproducible variance in the feature space is attributable to that parameter or interaction; %var measures the share of total variance in the dataset attributable to that parameter or joint interaction; permutation-based p-value denotes the probability that random chance alone would give an effect at least as large as the one actually computed, judged against the unexplained variation, based on 1000 random reshufflings of that unexplained variation.

Table 5 shows the summary of ASCA results for main effect matrices and lubricant-related interaction matrices. It is evident that ASCA on full feature space constitutes a statistically significant multivariate effect attributable to lubricant, although it explains only 2.75% of the total variance, compared with 21.21% for RPM and 44.55% for axis. The oil:RPM interaction is also significant, which supports the observation in Section 4.1 that the lubricant effect is not uniform across operating speeds. This is consistent with the proportion of features with statistically detectable lubricant effect in Section 4.1. Because the ASCA aggregates variance across all features equally, features carrying little lubricant information may dilute the share of variance attributable to lubricant. Being unsupervised, ASCA naturally surfaces the dominant RPM- and axis-driven variation rather than the comparatively weak lubricant effect. The signal is real but weak and exploiting it for discrimination requires a supervised feature representation explicitly optimised to find class-separating information, which is reported in the main analysis outlined in next section.

Table 5: Comparison of effect of parameters and joint interactions using ASCA result metrics

Model term	SSQ	%var	p-value
X_{oil}	131.07	2.75	≤ 0.001
X_{RPM}	1009.75	21.21	≤ 0.001
X_{start}	19.92	0.42	0.006
X_{axis}	2120.69	44.55	≤ 0.001
$X_{\text{oil:RPM}}$	260.47	5.47	≤ 0.001
$X_{\text{oil:start}}$	47.17	0.99	0.027
$X_{\text{oil:axis}}$	65.55	1.38	≤ 0.001
$X_{\text{RPM:start}}$	75.94	1.60	≤ 0.001
$X_{\text{RPM:axis}}$	472.79	9.93	≤ 0.001
$X_{\text{start:axis}}$	8.71	0.18	0.333

5 | Machine Learning Framework

Given the strong operating speed dependence identified in Section 4, the subsequent primary study analyses whether lubricant-related information can be extracted from a feature space dominated by operating speed and other nuisance variation. As mentioned in Section 3.2, each 30-second recording was chunked into 5 non-overlapping segments, yielding 5 replicates per run. Both hot- and cold-start recordings were retained in the analysis to expose the classifier to the uncontrolled thermal variability associated with realistic engine operation rather than restricting evaluation to a single thermal state. Since the two recordings for a given lubricant sample and operating speed correspond to the same underlying oil-RPM condition, they were treated as repeated measurements of that condition rather than independent experimental conditions. This results in 10 total replicates per oil-RPM condition. While this increases the training instances available, the presence of replicates of the same underlying physical sample introduces the possibility of a given oil-RPM condition spanning both training and testing data if not partitioned carefully. Therefore, the dataset was partitioned such that all 10 replicates representing one physical condition are grouped together, to ensure the classifiers never evaluate data drawn from a condition it was partially trained on. However, unlike run-averaged ANOVA in Section 4, subsequent classifier uses unaveraged chunks directly to maximise training density, with leakage prevented through fold-level grouping.

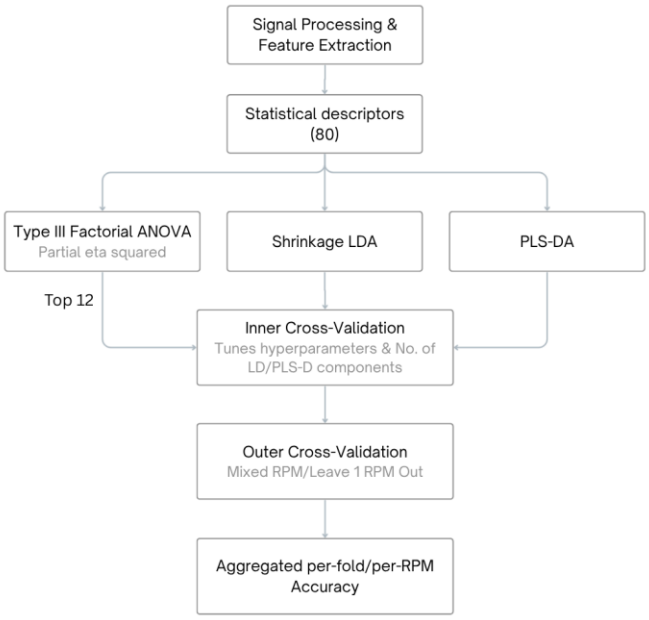


Figure 7: Machine learning pipeline used for discriminability assessment

Since there are 30 unique oil-RPM conditions (5 lubricant samples × 6 operating speeds) represented by 80 statistical descriptors, the dataset has a high feature-to-sample ratio. Such data-driven models can either suffer from high performance variance or deceptive accuracy due to overfitting [34]. To reduce the dimension of the feature vector containing statistical descriptors of the original signal while preserving the lubricant sample discriminability, three feature representation methods were used: top 12 descriptors ranked by an ANOVA score, Shrinkage Linear Discriminant Analysis (LDA) [37] projections and Partial Least Squares Discriminant Analysis (PLS-DA) [38] projections as outlined in Fig. 7. Each method assumes the following, respectively:

1. Discriminative information is concentrated in a small subset of the original univariate interpretable statistical descriptors, identifiable by their individual statistical effect size.
2. Discriminative information is distributed across many correlated descriptors and can be better represented through supervised discriminant projection combining several descriptors.
3. Discriminative information is more effectively captured through latent variables that maximise covariance between the statistical descriptors and lubricant labels.

ANOVA-based feature selection method is used as a univariate method where the machine learning model itself has interpretable vibration features in contrast to LD/PLS-D projections. A feature ranking score is given to each descriptor by its combined partial η^2 for the main oil effect and the oil:RPM interaction, normalised by the maximum across descriptors (Eq. 9,10). Within each training fold, the ANOVA used for ranking is fitted on run-averaged data using Type III sums of squares with sum-to-zero contrasts. To ensure that features with high effect sizes are genuinely statistically significant, and not simply noise, False Discovery Rate (FDR) correction (Benjamini–Hochberg procedure [39]) is applied within each training fold, separately to the oil and oil:RPM p-values; a term contributes to a descriptor's score only if its FDR-adjusted p-value (q-value) is below 0.05, and descriptors failing both receive a score of zero. Furthermore, a pairwise redundancy filter $|r| > 0.85$ was applied before the top 12 features with highest scores are selected for the data-driven model. This selects a small set of interpretable descriptors with the largest lubricant-related effect sizes, in contrast to the all-descriptor multivariate view given by ASCA

$$\eta^2_{\text{combined}}(y) = \eta^2(\text{oil} + \text{oil:RPM}, y) \tag{9}$$

$$S_{\text{ANOVA}} = \frac{\eta^2_{\text{combined}}(y)}{\max_{y_i \text{ for } i\in[1:80]}\left(\eta^2_{\text{combined}}(y)\right)} \tag{10}$$

Then, the latter two methods were therefore used to test whether multivariate supervised representation improves classification beyond individually selected interpretable descriptors. Shrinkage LDA was specifically used instead of standard sample covariance estimation (LDA) to improve stability, given the high feature-to-sample ratio. Shrinkage LDA and PLS-DA were applied to all 80 descriptors after removing descriptors correlated at $|r| > 0.95$, and all features were standardised within each training fold. Dimensionality reduction/feature selection step is undertaken exclusively within training data set given that all the aforementioned methods are supervised and prone to information leakage if done with the evaluation dataset.

Next, the classification to assess whether the lubricant samples are discriminable under all unrelated vibration noise, was conducted as a nested cross-validation (CV) using linear SVM. It is assumed that if relatively simple decision boundaries can discriminate the lubricant classes, then the information exists in the engineered feature representation without requiring a highly flexible nonlinear model. Nevertheless, a nonlinear RBF-SVM was additionally evaluated as a robustness check to determine whether the observed discrimination was dependent on the assumption of linear separability. In the nested CV method, a 3-fold inner CV was used on the outer fold's training data exclusively to tune the hyperparameters and select the optimal number of LD/PLS-D projections. Lastly, the outer CV was used to assess the accuracy of the model fit. Each classifier kernel's performance was compared against a permutation test, which shuffles the class labels for 1000 permutations and evaluates the model to generate a null distribution of accuracies. Then, the p-value from permutation test of each pipeline, which is the fraction of label-shuffled runs that matched or beat the observed accuracy. To study the effect that engine operation speed has, the outer CV was performed using two evaluation approaches: Mixed-RPM approach and Leave 1-RPM out approach.

5.1 Mixed-RPM approach

The Mixed-RPM approach uses a standard 5-fold nested CV where each training set contains data from all the operating speeds and all the lubricant samples. Each outer fold trains the model on 80% of data and tests it on the remaining 20% of data iteratively until all the data is tested. Each fold is numerically forced to contain exactly one oil at each RPM, which lubricant is randomised independently per RPM, so the test set composition is not underexposed or overexposed to one operating speed or lubricant sample. The complete 5-fold procedure is repeated 100 times (500 folds) with different random partitions to reduce the dependence of the results on any single train-test partition and to obtain a more stable estimate of model performance. Since there are 5 classification classes and 6 operating speeds, the RPM composition in each fold will be unbalanced. Hence, in this approach, the ANOVA ranking method uses Type III sum of squares with sum-to-zero contrast to account for unbalanced factorial design inside each training set. This evaluation criterion tests whether the classifier can recognise lubricant sample when all operating speeds are represented during training.

5.2 Leave 1-RPM out approach

The Leave 1-RPM out approach uses a 6-fold nested CV where the model is trained using five engine operating speeds and evaluated at the other operating speed unseen by the training set. This is done fold by fold iteratively until all six speeds are evaluated. Each fold is numerically forced to contain data from only one operating speed, which means that fold automatically spans all lubricant samples related to that speed. Since the training design is balanced, Type III with sum-to-zero contrasts is equivalent to Type II here. This evaluation criterion tests whether lubricant-related discriminative information generalises to a completely unseen operating speed.

6 | Results & Discussion

6.1 Classification Results

Table 6 shows the outer CV results for all combinations of feature representation methods, evaluation criteria and classifier kernels used. Although two kernels were investigated, subsequent interpretations employ linear SVM to explain the findings while keeping RBF SVM as a robustness check as mentioned in Section 5. Permutation tests used the fixed hyperparameters most frequently selected in nested CV, and lubricant labels were shuffled across all 30 independent conditions. In the Mixed-RPM approach this null does not preserve the design in which each lubricant appears once at every speed, so it is conservative, and the Mixed-RPM p-values should be read as upper bounds.

Table 6: Machine learning model performance and permutation test results

Machine learning pipeline			**Model accuracy**		**Permutation test (N = 1000)**		
			Mean	**Std**	**Mean**	**Std**	**p-value**
Mixed-RPM approach	ANOVA	Linear	0.2910	0.1335	0.1802	0.0628	0.064
		RBF	0.2322	0.1260	0.1839	0.0681	0.067
	LDA	Linear	0.5095	0.1615	0.1803	0.0637	≤ 0.001
		RBF	0.4674	0.1883	0.1804	0.0639	≤ 0.001
	PLS-DA	Linear	0.2190	0.1321	0.1798	0.0667	0.164
		RBF	0.1488	0.1348	0.1932	0.0773	0.483
Leave 1-RPM out approach	ANOVA	Linear	0.4700	0.1996	0.1875	0.0608	≤ 0.001
		RBF	0.4800	0.1249	0.1832	0.063	≤ 0.001
	LDA	Linear	0.5667	0.1459	0.1888	0.0616	≤ 0.001
		RBF	0.5400	0.1796	0.1886	0.0605	≤ 0.001

	PLS-DA	Linear	0.5400	0.1020	0.1860	0.0644	≤ 0.001
		RBF	0.4967	0.1104	0.1957	0.0703	0.002

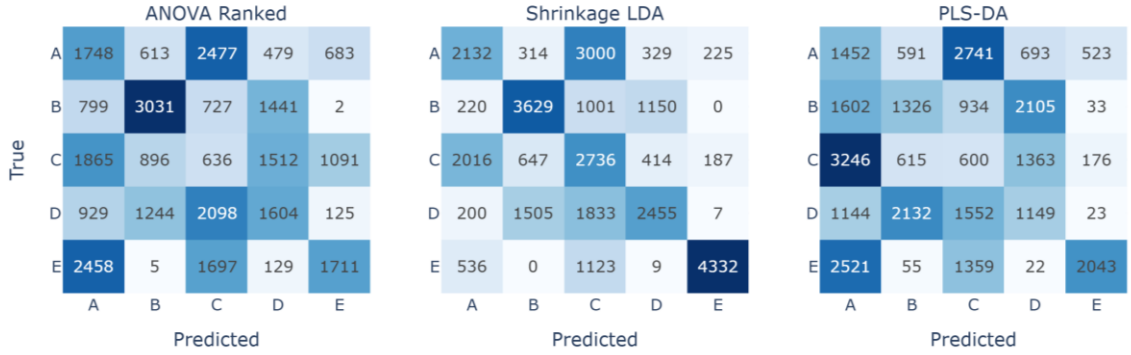


Figure 8: Combined classification matrices for Mixed-RPM approach

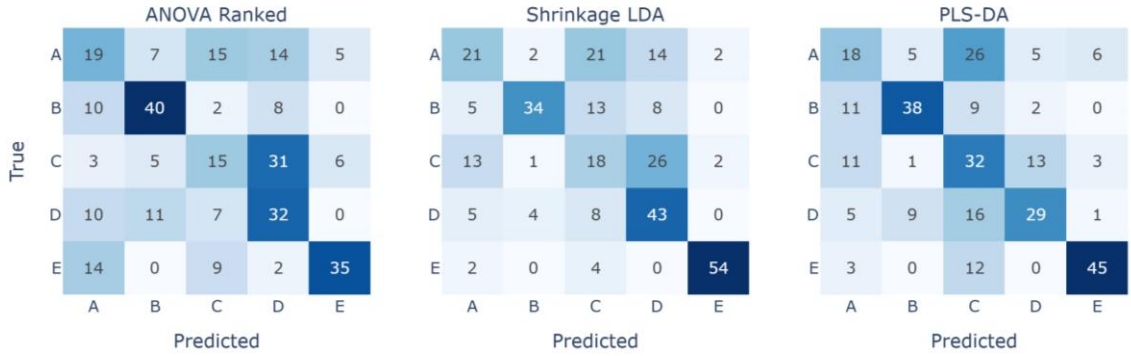


Figure 9: Combined classification matrices for Leave 1-RPM out approach

6.1.1 Mixed-RPM approach

For this evaluation criterion, from Table 6, it is observed that Shrinkage LDA pipeline performs consistently better in terms of mean accuracy and statistical significance ($p < 0.01$), than ANOVA-ranked and PLS-DA pipeline. Fig. 8 shows confusion matrices for all linear SVM fold predictions pooled together. Note that 34.1% of misclassifications of the best performing method (Shrinkage LDA) are primarily rooted in confusion of engine oil samples A and C. The same two samples are notably misclassified in the other two feature representation methods as well. This may be attributed to their relatively similar current viscosities at 40 ºC (84.08 cSt and 85.98 cSt, Table 1). No other consistent patterns are observed in PLS-DA and ANOVA-ranked feature representation methods, which may be attributed to low statistical significance (p-value, Table 6) of these two approaches.

6.1.2 Leave 1-RPM out approach

Recall that in Section 4.1, it was established that the mean additional variance on statistical descriptors by RPM is considerably high ($\eta^2 = 0.66$) compared to other effects. Since the model is tested at speeds unseen by the training set, this approach should be intuitively more rigorous given the statistical descriptors are dominated by the sole RPM effect and oil:RPM interaction. It can be seen that all feature

representations perform almost similarly with this evaluation strategy. According to Fig. 9, that is the classification matrices of all linear SVM predictions combined, the errors in this approach are attributed to samples A and C or samples D and C. Misclassification of samples A and C may be attributed to reasons mentioned in Section 6.1.1. To confirm this, pairwise confusions pooled over Leave-1-RPM out folds were compared with the difference in measured viscosity between each pair of lubricant using Spearman's rank correlation coefficient and a permutation of lubricant labels (p-value). Results in Table 7 describes that lubricant samples with closer viscosity tended to be confused more often. However, only two feature representations, namely Shrinkage LDA and PLS-DA has borderline significance ($p = 0.07$, 0.09 respectively). With only five samples, this is treated as descriptive. As for the confusion between samples D and C, it may be due to the fact that they are the only two 15W-40 lubricant samples.

Table 7: Pairwise confusion compared with difference in measured viscosity for Leave-1-RPM folds

Method	Spearman's ρ	p-value
ANOVA-ranked	-0.27	0.29
Shrinkage LDA	-0.82	0.07
PLS-DA	-0.57	0.09

From Table 6, it can be noted that Leave-1-RPM performs better than the Mixed RPM approach with +0.18, +0.06 and +0.32 increments in mean accuracy for in ANOVA, Shrinkage LDA and PLS-DA feature representations in linear SVM respectively. Furthermore, ANOVA and PLS-DA results attributed to this evaluation criterion show high statistical significance (p-value, Table 6) compared to the Mixed-RPM approach. Nevertheless, the marked difference between two evaluation strategies suggest that there is a considerable effect from this on generalisation. The same trend is observed in RBF SVM as well (Table 6).

6.2 Operating Condition Dependence

6.2.1 Effect of Engine Operating Speed

Since the statistical descriptors are dominated by operating speed compared to lubricant sample, an additional ablation analysis was carried out to test whether the presence of RPM of current test sample in the training set affected the comparative loss in accuracy in Mixed-RPM approach. In this ablation analysis, samples corresponding to the test sample RPM were excluded from the training set while retaining the remainder of Mixed-RPM structure. The results are tabulated in Table 8. The discrepancy in full training values in Table 8 with actual values in Table 6 is attributed to fixed hyperparameters used in ablation analysis compared to inner CV hyperparameter tuning performed in the actual model. Nevertheless, a substantial improvement in mean accuracy is observed, which is consistent with the suggestion that performance is conflated by the presence of competing samples from same operating speed in feature space. This also explains the drop of accuracy below permutation null mean in PLS-DA RBF (0.149 vs. 0.187, Table 6). In contrast, Leave 1-RPM out approach avoids this artefact as evaluating the sample at a completely different operating speed. However, under Leave 1-RPM out approach, accuracy at a given operating speed is a product of two phenomena: intrinsic oil separability at a given operating speed and the extrapolation quality of that RPM from other operating speeds. Hence, neither is perfect but subsequent comprehensive analyses will lead with Leave 1-RPM out approach given Mixed-RPM's feature space artefact and its weak statistical significance. Higher increment in accuracy of ANOVA-ranked and PLS-DA representations compared to shrinkage LDA in ablation test also reveals the reduced performance of those two representations in Mixed-RPM approach is attributed to this artefact in feature space. Shrinkage LDA explicitly seeks separation between lubricant sample classes relative to within-class variation (including RPM variation) by design; thus,

supervised discriminant projections of statistical descriptors seem to produce a more stable representation with a relatively low penalty when operating speed distribution is changed within dataset.

Recall that in Section 4.1, it was deduced that oil separability through vibration features depends on the engine speed at which the vibration is measured, through the statistically significant oil:RPM interaction in the type II factorial ANOVA model and in ASCA (Table 5), although its share of the total variance is small (5.5%). Therefore, a performance analysis was done as a function of operating speed to isolate the oil:RPM interaction. Fig. 10, 11 shows the variation of model performance, or in other words, the discriminability of lubricant sample at each engine operating speed for Mixed-RPM approach and Leave 1-RPM out approach respectively. As anticipated, the performance varies considerably across the RPMs at which the data samples were measured. For the Mixed-RPM approach, the accuracy peaks at 1850 RPM for ANOVA-ranked and shrinkage LDA representations and at 1600 RPM for PLS-DA representation. In contrast to that, for Leave 1-RPM out approach, the accuracy peaks shift to lower operating speeds for Shrinkage LDA (1300 RPM) and PLS-DA (1100 RPM) representations, while ANOVA-ranked feature representation retains its peak accuracy at 1850 RPM. Similar observations were made for RBF SVM as well for the most part.

Table 8: Ablation test results excluding the test sample RPM from training set for Mixed-RPM approach

Method	Full training	Excluding same RPMs	Increment	Leave 1 RPM Out
ANOVA-ranked	0.3002	0.4518	0.1516	0.4700
Shrinkage LDA	0.4423	0.5440	0.1017	0.5667
PLS-DA	0.2292	0.4409	0.2117	0.5400

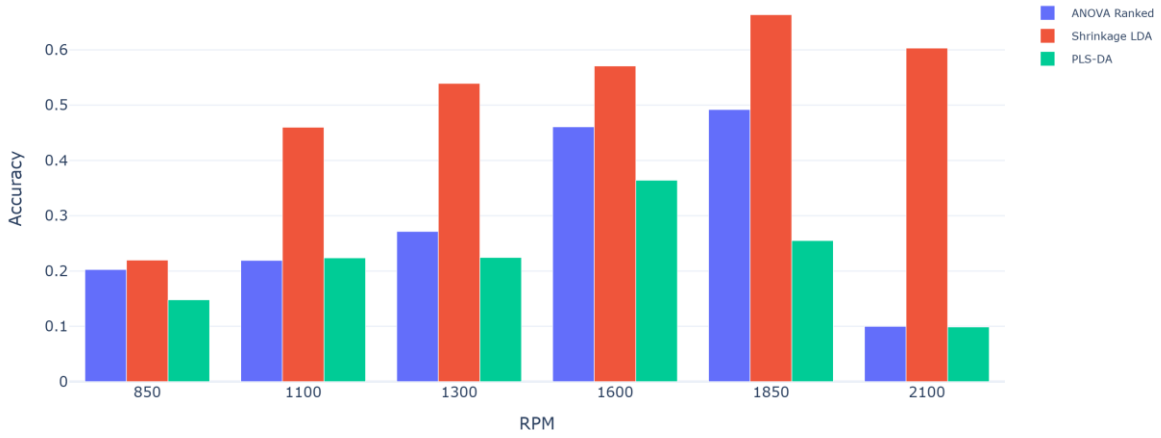


Figure 10: Breakdown of model performance by RPM for Mixed-RPM approach

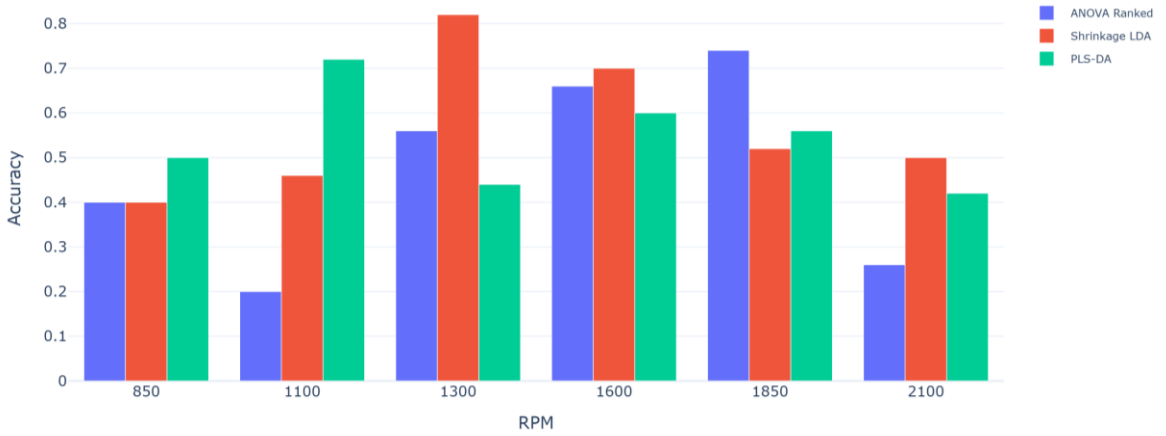


Figure 11: Breakdown of model performance by RPM for Leave 1-RPM out approach

Table 9: Comparison of lubricant-related informativeness at different engine operating speeds

Operating speed / RPM	ANOVA partial η^2	Increment of accuracy from ablation test		
		ANOVA-ranked	Shrinkage LDA	PLS-DA
850	0.3890	0.22	0.18	0.21
1100	0.3130	0.15	0.14	0.28
1300	0.2777	0.17	0.20	0.18
1600	0.3493	0.13	0.12	0.16
1850	0.4164	0.10	0.00	0.20
2100	0.3924	0.14	-0.03	0.24

The effect of the former phenomenon was analysed independently through a one-way ANOVA model and showed that the oil separability is maximum at 1850 RPM as detailed in Table 9. Furthermore, the accuracy changes per RPM breakdown for earlier ablation test revealed that the squared ablation error is also minimum at 1850 RPM for ANOVA-ranked method (Table 9), consistent with its peak accuracy occurring at the same speed in both evaluation strategies. This suggests that ANOVA-ranked, that is a subset of interpretable features is more robust towards identifying the localised operating speed range at which the lubricant sample is discriminable compared to dimensionally reduced feature representations.

6.2.2 Effect of Thermal History

The preceding analysis evaluated discriminability of lubricant samples after combining two starting conditions, therefore neglecting the thermal effect it has on the lubricant. To examine the effect of thermal history on discriminability of data, the classification accuracy was recomputed again for hot- and cold-start conditions separately. The results are tabulated in Table 10. From its results, it is evident that the accuracy of the data-driven models has increased as the engine heated up. This observation holds directionally in all ML pipelines, but the size of the effect varies considerably by the feature representation. The average increase in mean accuracy of two evaluation strategies is quite significant for ANOVA-ranked (+0.15) representation; modest for shrinkage LDA (+0.09) and PLS-DA (+0.05) representations. This means ANOVA-ranked feature representation is also more sensitive to the thermal effect of the lubricant on vibration characteristics.

This behaviour is consistent with the temperature-viscosity relationship of automobile lubricants. As engine operating speed increases, greater frictional and combustion heat is transferred to the lubricant via the cylinder wall and piston skirt, raising lubricant temperature. Consequently, this suggests that in hydrodynamic lubrication, the viscosity of the lubricant is reduces and therefore alters the damping

effect on vibration (Eq. 4,7). According to Stribeck curve, a sufficient reduction in Hersey number $H = \mu N/P$ (Eq. 8) can shift the lubrication regime toward boundary lubrication, where damping instead depends on the surface film properties of the lubricant rather than its bulk viscosity. In boundary lubrication, thermal changes are mainly affected through the chemical film formation and depletion on solid metal surfaces [2], [19]. To establish the ground-truth regarding the contribution of surface chemistry, TBN measurement is needed to reach a conclusion, which was not measured in this study. Moreover, as the Hersey number indicates, viscosity, and hence thermal history alone does not dictate this transition: for a given pressure (load) P, the lubrication regime also depends on rotational speed N, which varies substantially across the six operating speeds tested in this study. Operating speed therefore influences discriminability both directly, through its established effect on lubrication regime, and indirectly, by governing the rate of heat transfer to the lubricant and hence its thermal history.

This motivates an examination of the combined effect of speed and starting condition. Recall that factorial ANOVA in Section 4 indicated that the effect of starting condition was dependent on operating speed (significant start:RPM interaction in ASCA, $p \leq 0.001$), despite its relatively small overall main effect ($\eta^2 = 0.04$). To investigate this joint interaction effect, variation of lubricant discriminability, that is the model performance, for different engine operating speeds for separate hot- and cold-starts is broken down in Fig. 12-15.

Table 10: Machine learning model performance for hot-start and cold-start separately for linear SVM

Machine learning pipeline		**Cold-start only**		**Hot-start only**	
		Mean	**Std**	**Mean**	**Std**
Mixed-RPM approach	ANOVA	0.2032	0.1393	0.4058	0.1769
	LDA	0.4407	0.1658	0.5006	0.1578
	PLS-DA	0.1978	0.1350	0.2483	0.1631
Leave 1-RPM out approach	ANOVA	0.3667	0.1335	0.4733	0.1413
	LDA	0.4867	0.1590	0.6267	0.1692
	PLS-DA	0.4200	0.1301	0.4733	0.1896

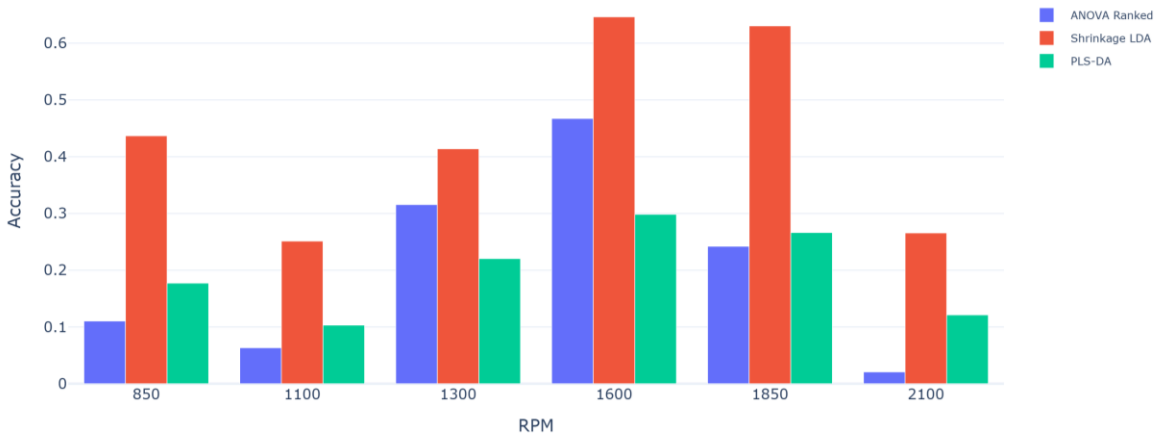


Figure 12: Breakdown of model performance by RPM for Mixed-RPM approach for cold-start condition

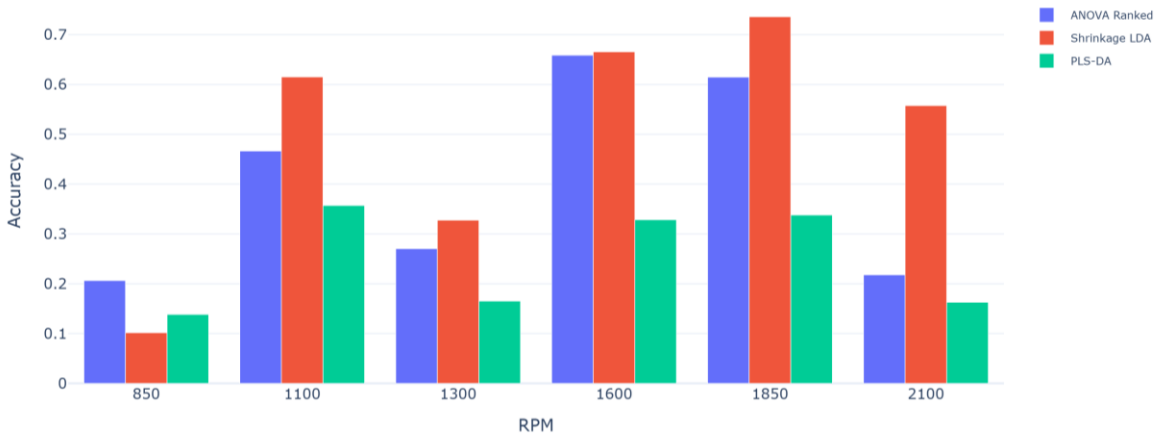


Figure 13: Breakdown of model performance by RPM for Mixed-RPM approach for hot-start condition

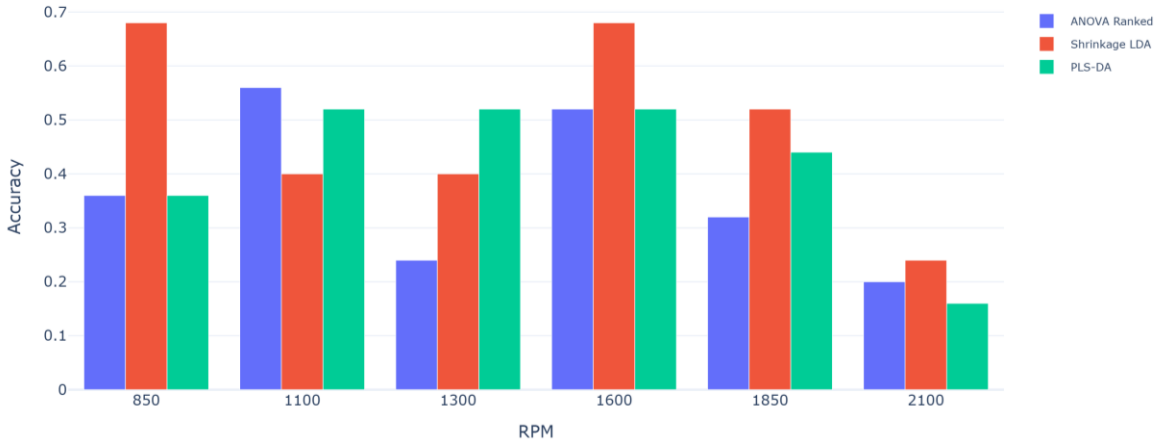


Figure 14: Breakdown of model performance by RPM for Leave 1-RPM out approach for cold-start condition

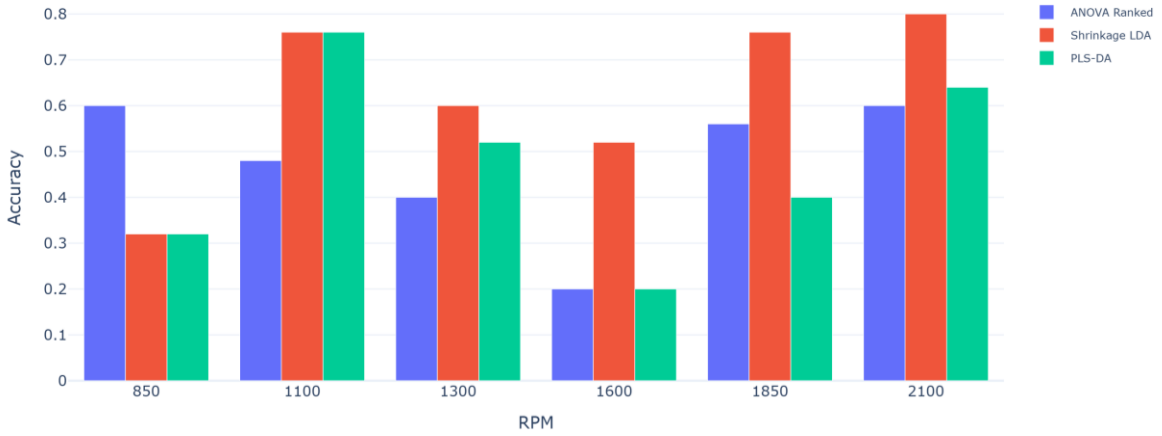


Figure 15: Breakdown of model performance by RPM for Leave 1-RPM out approach for hot-start condition

It can be established descriptively that both Mixed-RPM pair's (Fig. 12, 13) and Leave 1-RPM out pair's (Fig. 14,15) model performance by RPM notably performs better towards higher engine operating speeds at hot-start condition whereas for cold-start condition, the models perform better towards low engine operating speeds. This description is well defined at the operating speed extremes (850 RPM and 2100 RPM, Fig. 12-15). Given the low number of operating speeds, no statistical claims can be drawn with regard to this. This also means that the performance variation observed in Fig. 10, 11 is an average of the temperature effect of data that belongs to two thermal histories being pooled together.

This behaviour points toward the possibility that there exists some optimal Hersey number window in Stribeck curve where discriminability peaks, and hot- and cold-start conditions needs different operating speeds to reach that window. For example, at high viscosity (cold-start), a lower operating speed is required to achieve maximum discriminability whereas for low viscosity (hot-start), a higher operating speed is required to reach this same H window for no load. Furthermore, comparatively low model performance of the cold-start implies that the cold start's optimal region sits at or below the idle speed of the engine. In the same sense, hot-start's higher accuracy suggest that it is already within or approaching the most of this optimal window's span. Whether this optimal region on Stribeck curve corresponds to hydrodynamic, boundary or mixed lubrication cannot be determined from present vibration data. It should be noted that these are mechanistic hypotheses entirely suggestive of ML performance patterns observed.

The same concept emphasises that lubricated parts transition from hydrodynamic regime to boundary regime at higher engine loads. Therefore, given that experiment was only conducted for no load condition, this specific RPM behaviour can only be attributed to no load condition for this engine. Then, if the Hersey number hypothesis holds, at higher loads, the engine speeds at which this peak accuracy occurs would be expected to move to higher speed ranges in order to overcome the load acting on the lubricated parts (e.g., side thrust on piston skirts) to reach the optimal Hersey number window. To quantify the exact dependence of the engine load, a separate experiment has to be designed which is beyond the scope of this paper. Another plausible hypothesis that contributes to this operating speed dependence behaviour would be the excitation of certain vibration modes that underline lubricant related dynamics at these specific engine operating speeds, which warrants a more rigorous vibrational analysis.

6.3 Feature Interpretation

Relying on statistical descriptors to discriminate between lubricant samples without understanding their physical origin risks undermining the generalisability and transferability of the findings of this data-driven analysis. Therefore, the highest contributing descriptors identified independently by each feature representation method were examined through an auxiliary qualitative analysis on the full dataset. Fig. 16 summarises the most significant descriptors selected by ANOVA ranking and by the top three components of the shrinkage LDA and PLS-DA projections. To corroborate the findings inferred from feature loadings in Fig 16, a feature-group ablation test was performed excluding the classifiers from each descriptor category (Table 11). It should be noted that fixed hyperparameters were used in this test rather than nested CV tuning, thus increments are not directly comparable to results in Table 6. Also, loadings are from a single full-data fit and are interpreted qualitatively.

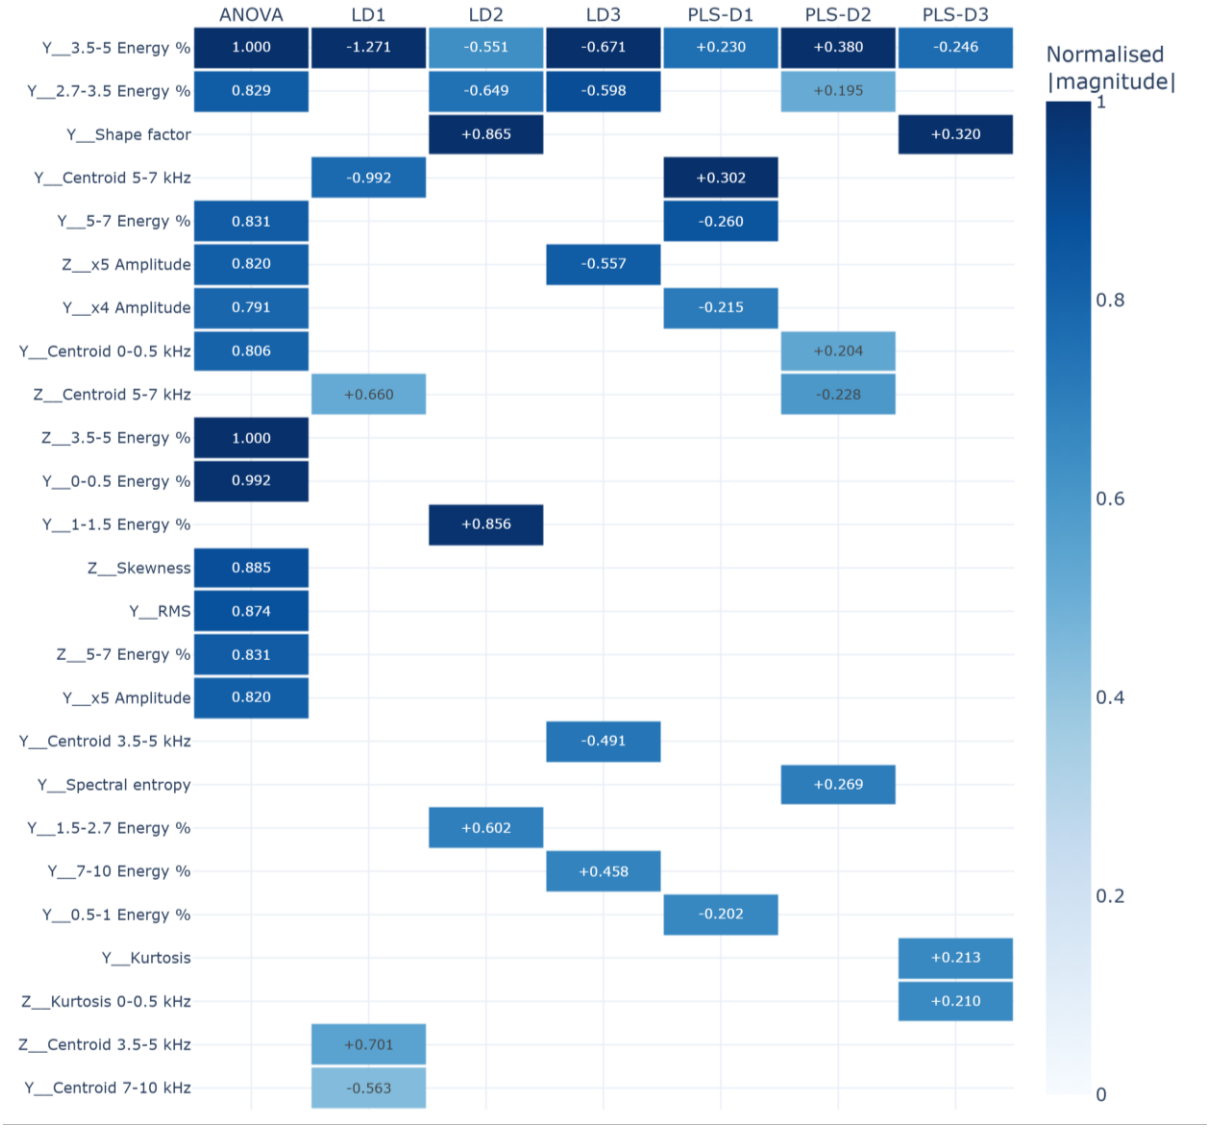


Figure 16: Column 1: Top 12 features ranked based on the ANOVA score of full dataset; Column 2-4: Top 3 components of the Shrinkage LD projections of full dataset; Column 5-7: Top 3 components of the PLS-D projections of full dataset.

Table 11: Increment in model performance for ablation tests excluding feature groups. Results are averaged for three feature representations and separate results for each evaluation strategy is tabulated.

Feature group	**Mixed-RPM accuracy increment**	**Leave 1-RPM out accuracy increment**
Time-domain features	-0.0098	-0.0311
Band-based frequency domain features	-0.2310	-0.2345
Broadband frequency domain features	-0.0146	+0.0144
Order amplitudes	+0.0136	+0.0322
Features from e_y sensor	-0.2076	-0.2089
Features from e_z sensor	-0.0056	-0.0089

Two patterns are apparent from this comparison. First, sub-band statistics are well represented: 17 of the 25 descriptors shown in Fig. 16 are extracted from fixed frequency bands, 5 are full-band time- or frequency-domain descriptors and 3 are order amplitudes. Since band-based descriptors make up 60% of the candidate pool (full-band 20%, order amplitudes 20%), this is close to what the pool composition

alone would give. The accuracy also drops (−0.2345) when band-based features are excluded, but this is the largest group, so the drop reflects group size as well as information content. Second, within the sub-band descriptors, band energy percentage is over-represented: it accounts for 10 of the 17 band entries, compared with one third of the band-based pool. Band spectral centroids account for 6 of 17, about their pool share, and band kurtosis appears once.

Notably, band energy percentages, which are normalised with respect to the total energy (that scales with operating speed), are comparatively more discriminative. This suggests that lubricant sample identity is better reflected in how vibration energy is redistributed across frequency bands than in the absolute magnitude of that energy (e.g., RMS), which is heavily confounded by RPM. Shape factor, a similar ratio-type time-domain statistic, shows the same behaviour, appearing among the top features in both the LD and PLS-D projections unlike other time-domain descriptors. These ratio descriptors still depend on RPM, however (partial η^2 for RPM of 0.73 for the 3.5–5 kHz band energy percentage and 0.68 for shape factor, compared with 0.99 for RMS), so normalisation reduces but does not remove the speed effect.

Band spectral centroids are also ratio-type descriptors and recur consistently across methods. A fair hypothesis for this take is the shift in resonant behaviour: as lubricant condition changes, certain resonant modes may appear to strengthen relative to others at a given operating speed. This reading is consistent with the significant oil:RPM interaction (Sections 4). Of all descriptors examined, band energy percentage in the 3.5-5 kHz range contributes most consistently.

The sensor axis and order-domain results point to a common mechanical origin. The majority of high-contributing descriptors originate from the e_y axis rather than e_z, implying that block vibration near the crankshaft is more sensitive to lubricant sample than vibration measured at the cylinder head. This is evident from ablation test as well; the accuracy drops significantly (-0.2089) when measurements from e_y sensor is excluded. In contrast, excluding measurements from e_z sensor has little to no effect on accuracy. Consistent with this, x4 order amplitude in e_y, the order equivalent to the number of pistons and journal bearings in the engine, appears among the top features in two of the three methods. The ablation test suggest that; however, inclusion of all order amplitudes reduce the accuracy by 0.0322 (Table 11). Together, these results suggest that piston side thrust, piston slap, crankshaft forces, bearing forces, and other lubricant-induced damping changes in e_y direction are more strongly reflected than camshaft and valvetrain motion, which are more directly associated with e_z. It should be, however, noted that preceding analysis is only significant in a statistical or mathematical sense, and not a definitely conclusive sense.

6.4 Practical Considerations

Although the physics behind the damping behaviour was explained in Section 1 under the assumption that the lubricant behaves as a Newtonian fluid (Eq. 1-4), this assumption does not hold for multigrade lubricants such as the 15W-40 and 10W-30 samples used in this study [12], [40]. Their polymer viscosity-index improvers make them shear-thinning, meaning that effective viscosity decreases as shear rate increases, particularly at the high shear rates present in bearing and piston skirt films [9]. Eq. 4 and 7, which assume a constant viscosity μ, therefore describe only a qualitative trend. Furthermore, the viscosities in Table 1 were measured at 40 °C under a low shear rate, and do not represent the effective viscosity of each lubricant in the engine at operating temperature and shear rate. The viscosity-based interpretation of sample similarity in Section 6.1 should be read with this in mind.

From a computational perspective, the processing cost of the proposed approach was not benchmarked in this study. The descriptors require operations such as Welch spectral estimation and band-pass filtering. The ANOVA-ranked pipeline would only need its 12 selected descriptors at inference, whereas

Shrinkage LDA and PLS-DA project all of the retained descriptors and therefore require the full feature set to be computed. In either case, the approach is likely to be considerably simpler than processing the raw vibration signal with a large neural network, which is favourable for low-resource on-board implementation should the method be extended beyond this feasibility study. Under severe memory and power constraints, approximate computing techniques have also been shown to reduce the latency and power consumption of feature-extraction hardware with limited loss in accuracy [41].

Several limitations of the experimental design should also be noted. First, each lubricant sample was recorded in a separate session consisting of drain, refill, warm-up and run, and the samples were run in the order A to E over two consecutive days. Sample identity therefore can be confounded with session factors such as ambient temperature, engine thermal state or residual lubricant, although measures were taken to minimise their effect. The confusions and top descriptors showed no monotonic relation with run order, but with only five samples this is descriptive and session effects cannot be excluded. Second, each sample was drawn from a different vehicle, so sample identity is also tied to viscosity grade, additive package, contamination and degradation history. As three samples are 10W-30 and two are 15W-40, the classifiers may be learning sample-specific signatures rather than viscosity or degradation state, which is why this study is limited to discriminating sample identity. Third, the cold-start sweep always preceded the hot-start sweep, so thermal history is confounded with sweep order, and lubricant temperature was not recorded. Temperature is also uncontrolled in practical scenarios, so logging lubricant temperature and including it as a model input may help compensate for such effects, which remains to be tested. Randomised run order and additional lubricant samples, including controlled blends or degradation levels, would be needed to move from discriminating identity towards predicting lubricant properties and degradation.

7 | Conclusion

This study investigated the feasibility of discriminating five field-used lubricant samples from engine recorded at six operating speeds under no-load conditions, with particular focus on the relative influence of lubricant sample, the robustness of different feature representations and the identification of interpretable vibration descriptors.

ASCA showed that lubricant sample identity was significantly associated with the vibration features, accounting for 2.75% of the total variance in the standardised feature matrix. Operating speed and sensor axis accounted for much larger shares, at 21.21% and 44.55%, respectively. The significant oil × RPM interaction accounted for another 5.47%, indicating that differences between samples varied with operating speed. However, each lubricant was tested in a separate recording session, so lubricant differences could be confounded with session effects.

Among the three feature representations investigated, shrinkage LDA followed by a linear SVM gave the highest mean classification accuracies when hot and cold-start data were analysed together. The results show mean accuracies of 0.51 and 0.57 under Mixed-RPM and Leave-1-RPM-out evaluation, respectively, against a normal chance level of 0.20. Both results were statistically significant ($p \leq 0.001$). These results support the second assumption in Section 5, where lubricant-related information is distributed across correlated descriptors and can be better represented through supervised discriminant projections. Although PLS-DA achieved a comparable accuracy of 0.54 under Leave-1-RPM-out evaluation, its performance was considerably weaker under Mixed-RPM evaluation, indicating less consistent discrimination across the two strategies. In contrast, ANOVA-ranked features provide an interpretable representation and were more informative in identifying the operating speed range at which lubricant discrimination was pronounced. Therefore, Shrinkage LDA is the preferred

representation for overall classification performance, whereas ANOVA-ranked features remain useful for interpreting the physical origin of the discriminative information.

Feature analysis identified band energy percentages, particularly in the 3.5–5 kHz range, as consistently useful for sample discrimination. Relative energy distribution therefore appears informative, although normalisation did not remove the influence of operating speed. Furthermore, descriptors extracted from the e_y sensor, mounted on the engine block near the crankshaft, were substantially more informative than those from the e_z sensor on the engine head. Excluding e_y features reduced the mean Leave-1-RPM-out accuracy by 0.21, whereas excluding e_z features produced a reduction of only 0.01. These results suggest that vibration associated with crankshaft dynamics and piston motion may be more sensitive to lubricant-related changes than vibration associated with the valvetrain and camshaft. The observed discrimination also varied with thermal history, with higher classification accuracy generally obtained from hot-start measurements, indicating that lubricant discrimination depends on both operating speed and the thermal condition of the engine.

Overall, this study demonstrates that the engine vibration contains lubricant-related across multiple operating speeds, despite the dominant influence of operating conditions. However, the maximum mean classification accuracy of 0.63 (with hot-start only) remains modest, and the findings demonstrate discrimination between five field-used lubricant samples rather than direct prediction of lubricant properties or degradation. Since sample identity is confounded with factors such as viscosity grade, additive package and degradation history, further experiments involving a larger number of lubricant samples and controlled degradation conditions are required to establish generalisability. Nevertheless, the results establish the feasibility of vibration-based lubricant discrimination and provide a basis for further investigating engine vibration as a complementary sensing method for lubricant condition monitoring.

Data availability:

Data will be made available on request.

Declaration of Competing Interest:

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References:

[1] J. B. Heywood, *Internal Combustion Engine Fundamentals*. in McGraw-Hill Series in Mechanical Engineering. New York: McGraw-Hill, 1988.

[2] A. Wolak and K. Fijorek, ‘Laboratory Diagnostics of Engine Oils as a Tool for Identifying Mechanical Faults and Supporting Sustainable Vehicle Maintenance’, *Sustainability*, vol. 18, no. 7, p. 3355, Jan. 2026, doi: 10.3390/su18073355.

[3] T. Omiya, K. Hanyuda, and E. Nagatomi, ‘Predicting engine oil degradation across diverse vehicles and identifying key factors’, *Mech. Syst. Signal Process.*, vol. 229, p. 112524, Apr. 2025, doi: 10.1016/j.ymssp.2025.112524.

[4] A. Wolak and W. Krasodomski, ‘Reducing Oil Waste Through Condition-Based Maintenance: A Diagnostic Study Using FTIR and Viscosity Monitoring’, *Sustainability*, vol. 17, no. 18, p. 8214, Jan. 2025, doi: 10.3390/su17188214.

[5] M. Habyarimana and A. A. Adebiyi, ‘Real-Time Engine Oil Quality Monitoring: A Review and Future Perspectives on Microcontroller-Based Sensor Fusion and AI’, *Appl. Sci.*, vol. 16, no. 6, p. 2919, Jan. 2026, doi: 10.3390/app16062919.

[6] R. Shah, V. Mittal, and M. Lotwin, ‘Recent Advances in Vibration Analysis for Predictive Maintenance of Modern Automotive Powertrains’, *Vibration*, vol. 8, no. 4, p. 68, Dec. 2025, doi: 10.3390/vibration8040068.
[7] A. C. Leite and M. L. de O. e Souza, ‘Fault Detection and Diagnosis (FDD) on a Knock Sensor’, SAE International, Oct. 2008. doi: 10.4271/2008-36-0369.
[8] S. S. Rao and Y. F. Fah, *Mechanical Vibrations*, 5. ed. in SI units. in Always learning. Singapore: Prentice Hall/Pearson, 2011.
[9] B. J. Hamrock, S. R. Schmid, and B. O. Jacobson, *Fundamentals of Fluid Film Lubrication*, 2nd edn. Boca Raton: CRC Press, 2004. doi: 10.1201/9780203021187.
[10] R. Stribeck, ‘Die wesentlichen Eigenschaften der Gleit-und Rollenlager [The key qualities of sliding and roller bearings]’, *Z Vereines Seutscher Ing*, vol. 46, pp. 1432–1437, 1902, Accessed: Sep. 13, 2026. [Online]. Available: https://cir.nii.ac.jp/crid/1572824500551324032
[11] M. D. Hersey, ‘The laws of lubrication of horizontal journal bearings’, *J. Wash. Acad. Sci.*, vol. 4, no. 19, pp. 542–552, 1914, Accessed: Sep. 13, 2026. [Online]. Available: https://www.jstor.org/stable/24520857
[12] S. Ardah, F. J. Profito, and D. Dini, ‘A comprehensive review and trends in lubrication modelling’, *Adv. Colloid Interface Sci.*, vol. 342, p. 103492, Aug. 2025, doi: 10.1016/j.cis.2025.103492.
[13] C. M. Santana, J. Mautone, and H. Almeida, ‘Analysis of the Increase Level of Vibration in an Internal Combustion Engine due to the Degradation of the Lubricating Oil’, pp. 2019-01–0780, Apr. 2019, doi: 10.4271/2019-01-0780.
[14] A. Moosavian, M. MehrabiVaghar, and M. Ghanbari, ‘Experimental Investigation on Engine Vibration Behavior Caused by Crankshaft Mechanism Under Different Lubricant Type and Temperatures’, *Automot. Sci. Eng.*, vol. 15, no. 1, pp. 4634–4645, Mar. 2025, doi: 10.22068/ase.2025.691.
[15] S. Ji *et al.*, ‘Study on the correlation between combustion status and block vibration induced by piston slap’, *Appl. Acoust.*, vol. 199, p. 109006, Oct. 2022, doi: 10.1016/j.apacoust.2022.109006.
[16] C. M. Santana, ‘Correlation Between Vibration Level, Lubricating Oil Viscosity and Total Number Base of an Internal Combustion Engine Operated with Gasoline and Ethanol’, in *SAE Technical Paper*, Mar. 2022. Accessed: Jul. 03, 2026. [Online]. Available: https://trid.trb.org/View/1953214
[17] G. Wu, G. Jiang, C. Chen, G. Jiang, X. Pu, and B. Chen, ‘An Experimental Study of the Effects of Cylinder Lubricating Oils on the Vibration Characteristics of a Two-Stroke Low-Speed Marine Diesel Engine’, *Pol. Marit. Res.*, vol. 30, pp. 92–101, Dec. 2023, doi: 10.2478/pomr-2023-0062.
[18] A. Nagy, J. Knaup, and I. Zsoldos, ‘Investigation of Used Engine Oil Lubricating Performance Through Oil Analysis and Friction and Wear Measurements’, *Acta Tech. Jaurinensis*, Jul. 2019, doi: 10.14513/actatechjaur.v12.n3.495.
[19] M. G. A. Nassef *et al.*, ‘Physics-Informed Transfer Learning for Predicting Engine Oil Degradation and RUL Across Heterogeneous Heavy-Duty Equipment Fleets’, *Lubricants*, vol. 13, no. 12, p. 545, Dec. 2025, doi: 10.3390/lubricants13120545.
[20] P. Lukášik, ‘Monitoring of engine oil degradation and possiblities of life predictions in combustion engine’, *Diagnostyka*, vol. 24, pp. 1–10, Jul. 2023, doi: 10.29354/diag/169032.
[21] Y. Chen, J. Shen, F. Zhou, H. Li, K. Yang, and L. Wang, ‘Multi-Sensor Fusion and Deep Learning for Predictive Lubricant Health Assessment’, *Lubricants*, vol. 13, no. 8, p. 364, Aug. 2025, doi: 10.3390/lubricants13080364.
[22] Y. Chen, H. Li, D. Li, K. Yang, and J. Zhou, ‘Diesel engine lubricating oil fault prognosis: A hybrid Bayesian LSTM and deep generative model architecture for multilayer anomaly detection’, *Tribol. Int.*, vol. 215, p. 111434, Mar. 2026, doi: 10.1016/j.triboint.2025.111434.
[23] J. Antoni, N. Ducleaux, G. NGhiem, and S. Wang, ‘Separation of combustion noise in IC engines under cyclo-non-stationary regime’, *Mech. Syst. Signal Process.*, vol. 38, no. 1, pp. 223–236, Jul. 2013, doi: 10.1016/j.ymssp.2013.02.015.
[24] D. Abboud, J. Antoni, S. Sieg-Zieba, and M. Eltabach, ‘Deterministic-random separation in nonstationary regime’, *J. Sound Vib.*, vol. 362, pp. 305–326, Feb. 2016, doi: 10.1016/j.jsv.2015.09.029.

[25] M. Yasuda and H. Maruyama, ‘A New 1.6-Liter Twin-Cam 16-Valve Nissan Engine’, *SAE Trans.*, vol. 100, pp. 1120–1131, 1991, Accessed: Jul. 03, 2026. [Online]. Available: https://www.jstor.org/stable/44554468
[26] H. Mahdisoozani *et al.*, ‘Performance Enhancement of Internal Combustion Engines through Vibration Control: State of the Art and Challenges’, *Appl. Sci.*, vol. 9, no. 3, p. 406, Jan. 2019, doi: 10.3390/app9030406.
[27] S. Shokrollahi, F. Adel, and H. Ahmadian, ‘An investigation into the accelerometer mounting effects on signal transmissibility in modal measurements’, *Sci. Iran.*, vol. 0, no. 0, pp. 0–0, Aug. 2017, doi: 10.24200/sci.2017.4244.
[28] H.-B. Jun, F. Conte, D. Kiritsis, and P. Xirouchakis, ‘A Predictive Algorithm for Estimating the Quality of Vehicle Engine Oil’, *Int. J. Ind. Eng. Theory Appl. Pract.*, vol. 15, Dec. 2008.
[29] M. O. Duarte, L. M. Margalho, W. Gołębiowski, M. Mendes, J. M. T. Farinha, and B. Šarkan, ‘Monitoring the condition of city bus engines by analysing used oil using PCA and K-Means clustering’, *Sci. Rep.*, vol. 16, no. 1, p. 9392, Feb. 2026, doi: 10.1038/s41598-026-39045-x.
[30] G. Wojnar, R. Burdzik, A. N. Wieczorek, and Ł. Konieczny, ‘Multidimensional Data Interpretation of Vibration Signals Registered in Different Locations for System Condition Monitoring of a Three-Stage Gear Transmission Operating under Difficult Conditions’, *Sensors*, vol. 21, no. 23, p. 7808, Jan. 2021, doi: 10.3390/s21237808.
[31] P. Welch, ‘The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms’, *IEEE Trans. Audio Electroacoustics*, vol. 15, no. 2, pp. 70–73, Jun. 1967, doi: 10.1109/TAU.1967.1161901.
[32] B. Łazarz, H. Madej, G. Peruń, and Z. Stanik, ‘Vibration based diagnosis of internal combustion engine valve faults’, *Diagnostyka*, no. 2(50).
[33] J. Antoni and R. B. Randall, ‘The spectral kurtosis: application to the vibratory surveillance and diagnostics of rotating machines’, *Mech. Syst. Signal Process.*, vol. 20, no. 2, pp. 308–331, Feb. 2006, doi: 10.1016/j.ymssp.2004.09.002.
[34] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems and Control*. Cambridge University Press, 2021.
[35] A. Theissler, J. Pérez-Velázquez, M. Kettelgerdes, and G. Elger, ‘Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry’, *Reliab. Eng. Syst. Saf.*, vol. 215, p. 107864, Nov. 2021, doi: 10.1016/j.ress.2021.107864.
[36] M. S. Paolella, *Linear models and time-series analysis: regression, ANOVA, ARMA and GARCH*. Hoboken: John Wiley & Sons, 2019.
[37] J. H. Friedman, ‘Regularized Discriminant Analysis’, *J. Am. Stat. Assoc.*, vol. 84, no. 405, pp. 165–175, Mar. 1989, doi: 10.1080/01621459.1989.10478752.
[38] R. G. Brereton and G. R. Lloyd, ‘Partial least squares discriminant analysis: taking the magic away’, *J. Chemom.*, vol. 28, no. 4, pp. 213–225, 2014, doi: 10.1002/cem.2609.
[39] Y. Benjamini and Y. Hochberg, ‘Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing’, *J. R. Stat. Soc. Ser. B Methodol.*, vol. 57, no. 1, pp. 289–300, Jan. 1995, doi: 10.1111/j.2517-6161.1995.tb02031.x.
[40] B. A. Gecim, ‘Non-Newtonian Effects of Multigrade Oils on Journal Bearing Performance’, *Tribol. Trans.*, vol. 33, no. 3, pp. 384–394, Jan. 1990, doi: 10.1080/10402009008981968.
[41] R. K. Amartey and D. Zhao, ‘Low-Latency Edge Computing Architectures for Real-Time Vehicle Warning Systems: A Review’, *Future Internet*, vol. 18, no. 8, p. 387, Aug. 2026, doi: 10.3390/fi18080387.