15th IFK
Proc. of the 15th International Fluid Power Conference <Symposium/Conference>
<DOI>

Predicting leakage and operating conditions in axial piston pumps using vibration-based machine learning

Simon Knoll, Svenja Horn and Jürgen Weber

Institute of Mechatronic Engineering, Technische Universität Dresden,

E-mail: simon.knoll@mailbox.tu-dresden.de, fluidtronik@mailbox.tu-dresden.de

Abstract.

Condition monitoring for axial piston pumps is often hindered by high sensor and infrastructure costs. This paper presents the groundwork for a low-cost, edge-computing "virtual sensor" that uses vibration-based machine learning (ML) to predict pump operating conditions (OCs) and leakage. Using data from a test rig with multiple artificially generated fault types, ML models (XGBoost and Multilayer Perceptron) were evaluated.

Results show that direct leakage prediction from vibration data is accurate for known faults ($R^2 \geq 0.97$) but fails to generalize to unknown faults ($R^2 > 0.50$). This prediction for unknown faults improves significantly (e.g., $R^2 \approx 0.80$) when OCs are provided as features. It is shown that these OCs (pressure, speed, and stroke) can be accurately predicted ($R^2 > 0.96$) from the same vibration signal. This high accuracy, however, is critically dependent on a robust training dataset that includes various fault types.

Finally, edge-computing feasibility is confirmed by deploying a Multilayer Perceptron model to an ESP32-S3 microcontroller starting at less than 10 €. This work validates the key data requirements for developing a cascaded, low-cost smart sensor.

Keywords. Economic fault detection, condition monitoring, machine learning, Edge-Computing.

1. INTRODUCTION

Condition Monitoring (CM) for fluid power systems, particularly axial piston pumps, is a subject of growing importance. These pumps are critical components in many industrial settings, such as steel mills, where they can be difficult to access for maintenance. Traditional maintenance relies on fixed replacement intervals, which is often costly and does not prevent spontaneous, expensive failures. Predictive Maintenance (PM) offers a needs-based alternative, but its implementation in existing systems faces significant hurdles.

Integrating new sensors into established industrial plants is often complicated and cost intensive. Furthermore, many applications involve sensitive process data, making cloud-based monitoring solutions undesirable due to data security concerns. These practical and economic constraints point toward the need for low-cost, stand-alone retrofit solutions that perform calculations locally. This "Edge-Computing" approach, using low-power microcontrollers, has become increasingly viable. While much recent research focuses on high-performance Deep Learning algorithms, this paper targets a robust, low-hardware solution by focusing on predicting leakage flow, a highly interpretable health index for axial piston pumps.

To achieve this without an expensive flow meter, a "virtual sensor" concept is proposed. This paper demonstrates the ground work for a cascaded regression method where vibration data from a single, low-cost accelerometer—identified in previous work as a highly promising sensor for ML-based monitoring [1]—is first used to predict the pump's operating conditions (pressure, speed, and stroke). This predicted operating condition information is then, in a second step, used to improve prediction of the pump's leakage flow.

This paper is structured as follows: Chapter 2 provides an overview of the existing work in fault detection for Axial Piston Units (APU). Chapter 3 introduces the artificially generated faults, the test rig, the installed sensors, and the machine learning (ML) methods. Chapter 4 presents the results of the ML algorithms including the proof-of-concept implementation and performance of this approach on a low-cost microcontroller for edge deployment.

2. State of the Art

Condition Monitoring (CM) and Predictive Maintenance (PM) are established fields for hydraulic systems, motivated by the high economic cost of unexpected failures and the growth of the Industrial Internet of Things (IIoT). A significant amount of research focuses specifically on axial piston units (APUs), developing CM systems for both mobile and stationary applications. While model-based fault detection exists [2, 3], data-driven methodologies using sensor data (e.g., pressure, acceleration) are gaining interest.

Most research employs Machine Learning (ML) for fault classification, utilizing algorithms like K-Nearest Neighbors [4], SVMs [5], or Deep Learning (DL) [6, 7] to identify the type of fault or a severity measure defined by multiple classes. Notably, Duensing [22] predicted wear progress in APU using data from actual degradation rather than artificial faults, identifying Multilayer Perceptrons and Random Forests as the most effective models. Keller et al. [4] identified specific acceleration and operating parameters as essential for classifying valve plate damage. Aligning with this, previous work [1] validated the triaxial acceleration sensor for ML tasks and demonstrated that higher sampling rates improve prediction quality, motivating the 16.4 kHz rate used here. Some approaches omit manual feature extraction using DL pipelines. For instance, Tang et al. [6] utilize wavelet transforms to convert signals into images for DL classification. While effective, this lacks an interpretable, continuous degradation measure necessary for maintenance planning. Additionally, data-driven models often struggle with small datasets, requiring transfer learning or simulated data [8, 9].

Addressing this, regression models enable continuous condition prediction. Schimke et al. [10] applied Random Forests and CNNs to predict damage severity. To simulate cavitation erosion holes are drilled in specific areas of the valve plate. The area of the drilled-out portion then functions as the regression target. Their use of 96 operating conditions highlights the importance of dataset diversity. They emphasize that inclusion or exclusion of damage types in the training dataset significantly influences prediction accuracy.

Other studies quantify degradation via abstract "Health Indices" (HIs) [11, 12]. However, unlike physical metrics such as leakage flow, HIs often lack the direct meaning required to set concrete maintenance thresholds. While Liu et al. [9] proposed a "virtual flow sensor" to predict leakage directly, their reliance on simulated data limits real-world relevance. Furthermore, high-performance architectures like PINNs [13], CNNs [10], and LSTMs [14, 15] demand computational resources often unavailable in low-cost, retrofit scenarios.

Walani et al. [16] argue for lightweight models on edge platforms to ensure real-time processing and low latency, though hardware limitations restrict model complexity.

Consequently, this review identifies a gap for an APU monitoring solution that is simultaneously: 1) Low-Cost & Edge-Deployable, avoiding heavy DL models; 2) Interpretable, by predicting physical metrics like leakage; and 3) Robust against variable conditions and unknown faults, a challenge for static classification models.

3. Experimental and Methodical Set-up

This section provides an overview of the artificially worn pump parts, the test rig and sensors, followed by a description of the measurement procedure and a brief introduction of the methodical approach for the implementation of ML algorithms. More details as well as the pump efficiencies for the different faults are presented in [17].

3.1. Artificially worn pump parts

Figure 3.1 illustrates three representative types of axial piston pump damage, each reproduced as realistically as possible. To identify the location and severity of a typical valve plate cavitation damage, simulations were conducted at elevated speeds, complemented by measurements at very low suction pressures. Real cavitation erosion at three different intensities (I, II, III) was induced locally in the transition zone between low and high pressure using a sonotrode, compare Figure 3.1 a).

Axial clearance between slipper and piston, a failure that can result in total pump breakdown, was implemented with a tensile testing machine. The required axial force was estimated via finite element analysis (FEA), as shown in Figure 3.1 b). All nine slipper-piston assemblies were then elongated until reaching a 179 µm clearance using a tensile testing machine, with a custom fixture to clamp both parts without damaging the sliding interfaces.

a) Valve plate cavitation **b) Axial clearance slipper-piston** **c) Slipper surface wear**

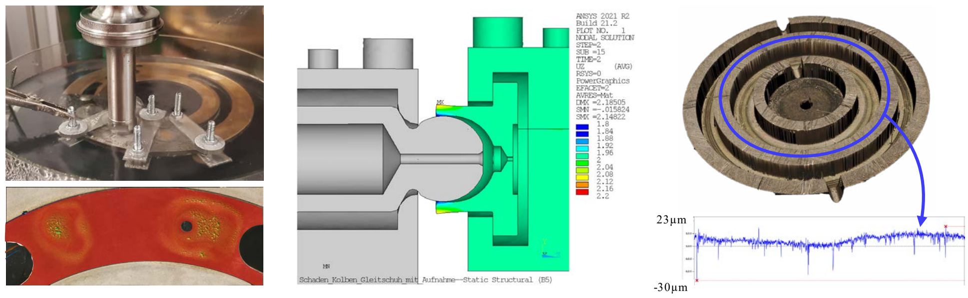

Figure 3.1 Artificially damaged pump parts, manufactured as realistically as possible (more details in [17])

Finally, slipper sealing lands were intentionally damaged using a file, and the achieved damage was quantified by profilometer measurements, as depicted in Figure 3.1 c).

3.2. Pump Test-Rig and Measurement Procedure

A standard test rig setup, designed for pump efficiency measurements, was equipped with high-frequency pressure and acceleration sensors. These were sampled at 16.4 kHz and located at the pump's front plate and housing, compare Figure 3.2. The triaxial acceleration sensor at the front plate provided the most informative signals for machine learning applications. In this study, all three acceleration components were included in the predictive pipeline. Suction pressure was maintained at 1.8 bar to prevent cavitation throughout all operating conditions. Low-frequency pressure and temperature data were acquired to monitor and control the test environment. Measurements were conducted across the full operating range of the variable displacement APU, covering six pressure levels, four stroke values and four rotational speeds, resulting in 96 distinct operating conditions. Most faults resulted in minor efficiency losses, with the exception of severe valve plate cavitation III, as clearly shown by the leakage results (see Figure 4.1).

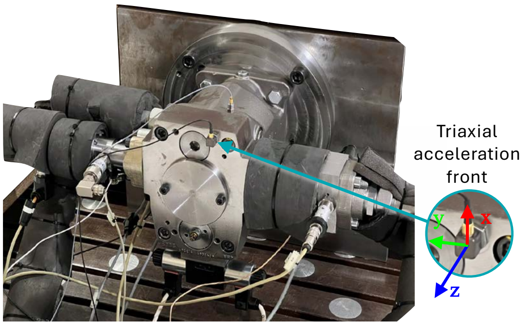

Property	Value
Sensor	MMF KS903.10
Sensor type	triaxial piezoelectric accelerometer
Mounting method	threaded stud
Upper frequency limit (3 dB)	11 kHz (X), 22 kHz (Y/Z)
Sampling Frequency	16.4 kHz

Figure 3.2 Mounting location and technical information for the triaxial acceleration sensor used in this work.

The study's findings are based on a representative subset of 28 operating conditions that are presented in Table 1. Each measurement lasted 10 seconds, sufficient to capture multiple pump rotations at each speed. Since temperature significantly affects leakage and vibration damping, the inlet temperature was controlled at 40°C ± 3°C. In practical applications, temperature variations can be measured cost-effectively and incorporated into the evaluation process.

Table 1: Overview of representative set of 28 operating conditions (OCs) grouped by stroke value relative to the maximum

100 % stroke		**75 % stroke**		**50 % stroke**		**25 % stroke**	
n/rpm	Δp/bar	n/rpm	Δp/bar	n/rpm	Δp/bar	n/rpm	Δp/bar
500	200	500	150	500	50	500	100
1000	100	500	250	1000	150	500	250
1000	250	1000	50	1000	300	500	300
1500	150	1000	200	1500	50	1000	150
1500	200	1500	100	1500	250	1500	50
2000	50	1500	300	2000	100	2000	200
2000	300	2000	150	2000	200	2000	250

Initially, a reference measurement was conducted under healthy conditions without any damaged parts ("new" condition). This data serves as a baseline for comparison with subsequent measurements involving damaged parts. Each defect was measured separately, without combining multiple defects.

3.3. Methodical approach for implementing machine learning

This section details the standardized methodology applied to all supervised machine learning tasks performed in this paper using Python (v3.12.3). As illustrated in Figure 3.3 the approach encompasses data acquisition, data preprocessing, feature engineering, model selection, hyperparameter optimization and final model evaluation. Additionally, a proof-of-concept edge deployment on an ESP32 micro controller is presented for one selected model.

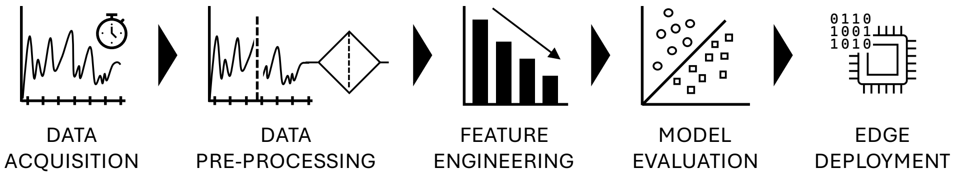

Figure 3.3 Methodical approach to implement machine learning algorithms.

3.3.1. Data segmentation and feature engineering

The raw acceleration signals exhibit cyclostationary behaviour, with patterns repeating for each pump revolution. To create a high volume of training samples, the data is split into smaller segments, each containing two full pump revolutions. Each segment underwent normalization (zero mean, unit variance) to emphasize signal shape over absolute amplitude across operating conditions. Feature extraction was performed using TSFresh (v0.20.2) [18], a Python library validated in prior studies [1, 17].

Table 2: Selected feature types ordered by count in the final top 700 (names matching the TSFresh implementation)

Feature Type	Amount in final selection	Feature Type	Amount in final selection
fft_coefficient	384	linear_trend	9
change_quantiles	126	number_cwt_peaks	6
agg_linear_trend	62	mean_abs_change	3
ar_coefficient	26	longest_strike_below_mean	3
autocorrelation	22	longest_strike_above_mean	3
partial_autocorrelation	21	count_above_mean	3
number_peaks	15	count_below_mean	3
fft_aggregated	11	abs_energy	3

From an initial candidate set of approximately 2,400 features, a preliminary reduction using univariate p-value analysis selected the top 16 feature types based on median ranking. This yielded 1,600 spectral and time-domain features, which essentially describe the shape and patterns of the signal. A second p-value analysis further reduced this selection to a final set of 700 features, detailed in Table 2. The most significant predictors included time-series autocorrelation, longest strikes above and below mean as well as spectral attributes (FFT coefficients, skewness, variance, kurtosis) up to 1.7 kHz equivalent frequency. Finally, the feature matrix was standard-scaled to ensure equal contribution during training. The regression label for each sample is represented by the unscaled mean leakage flow for the corresponding time segment.

3.3.2. Model selection

To balance performance with deployment feasibility, two regression models are employed. First, **XGBoost** (v2.0.3) [19], a decision-tree-based ensemble, serves as a robust benchmark due to its proven success in prior studies. Second, a **Multilayer Perceptron** (MLP) using Scikit-Learn (v1.5.0) [20] is selected specifically for its computational efficiency, making it ideal for resource-constrained edge deployment.

3.3.3. Hyperparameter optimization and training

Before the models are optimized and trained, 20 % of the entire dataset is separated as a final, untouched hold-out set for the calculation of performance metrics. The remaining 80 % of the data is used for hyperparameter optimization. The optimization for both models is conducted using the Optuna framework (v3.6.1), an automated, trial-based hyperparameter optimization tool [21]. This process employs a GroupKFold cross-validation strategy with 5 folds per trial. For this strategy, the data is grouped by the specific "damage kind" (e.g. valve plate cavitation II, slipper wear). This ensures that for each fold of the optimization process, the validation set contains data from at least one damage type that the model has not been exposed to during training in that fold. The objective of this approach is to develop final deployment models that are as robust as possible and generalize well to novel or unseen fault conditions. The hyperparameter search spaces for both models are represented in Table 3 in the appendix.

3.3.4. Final model evaluation

Performance results in the following sections are generated using the 20 % hold-out set that was segregated prior to hyperparameter optimization. The evaluation on this hold-out set is conducted using a KFold cross-validation strategy with 10 folds and random sampling aiming to represent the real-world deployment of pretrained models. In each iteration, a model is instantiated with the optimized hyperparameters, trained on 90% of the development set, and evaluated against the fixed hold-out set. This process is repeated 10 times, ensuring that varying samples in the original set are used for training. The aggregated results from these 10 folds provide a robust and comprehensive assessment of the model's final performance in a deployment setting.

4. Predicting pump leakage

4.1. Critical faults rise leakage flow

Leakage flow measurement is a widely used indicator of hydraulic pump condition, as increased leakage often signals wear and the need for replacement. However, clamp-on flow meters are expensive (beginning from 1000 €) and typically require a straight section of pipe or specific housing to ensure accurate readings, making them less practical for many installations. In-line flow meters, on the other hand, are difficult to implement in existing systems and can be prone to failures. For these reasons, the vibration signal should be used to predict the pump's leakage. In comparison to traditional fault analysis as for example the classification in "good" and "bad" or the use of a "health index" [12], operators that have a long-standing experience with their pumps and systems know the critical leakage values to decide whether the pump needs to be exchanged. As a further advantage, leakage has the potential to also work for different pump sizes if normalized. Not all pump faults - such as those involving the shaft or bearings - directly affect leakage, so relying solely on leakage measurements may miss some mechanical issues. Therefore, vibration analysis provides a more comprehensive assessment, allowing for better detection of both hydraulic and mechanical faults. In addition to predicting leakage, other models can monitor the vibration signal to detect fault types that do not measurably increase leakage, a capability demonstrated in previous fault classification work [1, 17].

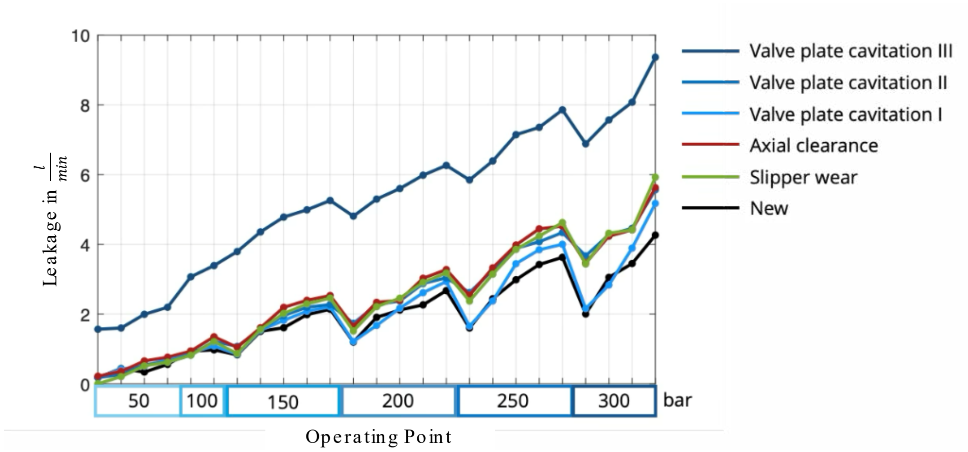

Figure 4.1 Leakage of the different faults over many operating conditions at maximum stroke (100 %) with rising speed for each pressure level

Figure 4.1 illustrates the measured leakage for the maximum stroke (100 %) for the different faults across a wide range of operating conditions, clustered by pressure levels and rising speeds for each pressure. Measurements were conducted using a VSE VS1 gear type flow rate sensor recording the total external leakage flow of the APU at a sampling frequency of 330 Hz. All faults are situated within the minor wear range, causing only slight increases in pump leakage, except for valve plate cavitation III, which shows a more pronounced effect. The figure demonstrates that higher pressures and speeds particularly accentuate leakage differences between fault types. Additionally, leakage is shown to be significantly dependent on operating conditions and temperature, supporting the hypothesis that machine learning models could improve fault detection by incorporating additional variables such as speed and pressure.

4.2. Correlation between pump vibration and fault types

While leakage indicates severe faults (e.g., valve plate cavitation III), it shows only minor increases for defects like slipper wear (Figure 4.1). To capture a holistic view of pump health, a vibration-based ML approach is employed. The correlation between vibration and condition is analysed using heatmaps of the y-acceleration amplitude against the "pump order". Due to cyclostationary behaviour, spectral peaks occur at harmonics of the piston frequency f_p, where Z is the piston count and the n rotational speed. Unlike raw frequency, the pump order normalizes the signal frequency f against f_P, enabling comparisons across varying speeds:

$$pump\ order = \frac{f}{f_p} = \frac{f}{Z * n}$$

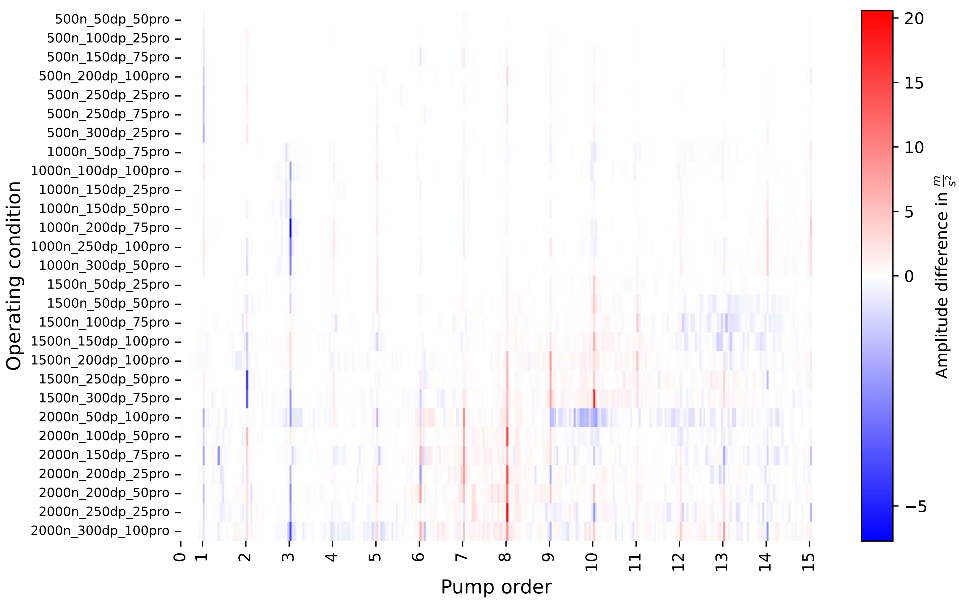

Figure 4.2 Differential amplitude spectrum of the y-acceleration signal highlighting changes from new pump condition to valve plate cavitation damage (stage III) across varying operating conditions.

Figure 4.2 illustrates the spectral difference between a healthy pump (new condition) and a pump exhibiting valve plate cavitation (stage III) as a function of pump order and operating condition. The difference spectrum directly reveals how the fault alters the acceleration amplitude at each harmonic order: blue regions indicate amplitude reductions relative to the healthy baseline, while red regions denote amplitude increases induced by the cavitation damage. Notably, certain orders, such as the 3rd harmonic, exhibit particularly pronounced differences, and the magnitude of these deviations varies significantly across operating conditions. This order-dependent and operating-point-sensitive behaviour confirms that spectral analysis of the acceleration signal provides a robust means of detecting and characterizing condition changes in the pump.

4.3. Leakage flow prediction with supervised learning

The dataset of the 28 OCs (see Table 1) is used, including the data from the new condition, axial clearance, slipper wear, as well as valve plate cavitation I and III. The fault of valve plate cavitation II is not included, as it will be used to judge the model's ability to perform on unknown damages. When training deployment models, data may only be available for low and high fault intensity (e.g. valve plate cavitation I and III). During the deployment of the final model, actual faults are expected to lie in between these extremes as represented by valve plate cavitation II. The triaxial sensor is evaluated, whereby for this investigation all three axes x, y, and z are used at the full sampling rate of 16.4 kHz.

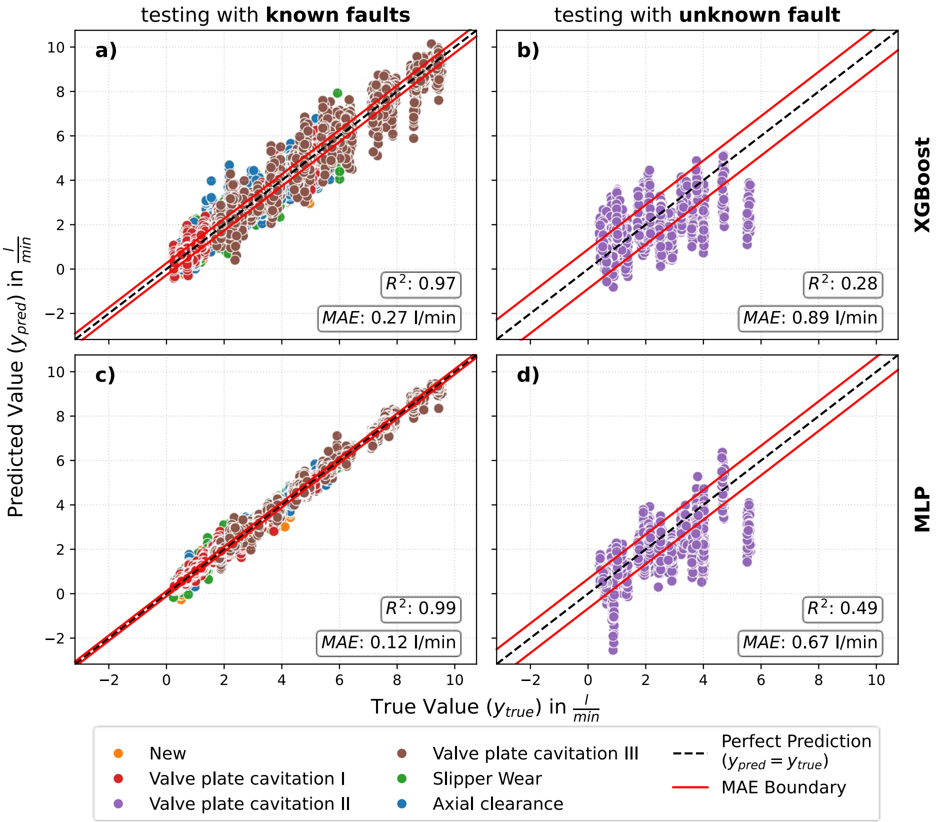

Figure 4.3 Comparison of leakage prediction with known pump conditions and unknown valve plate cavitation II damage not seen during training using XGBoost and MLP regression models.

Figure 4.3 a) shows the prediction result of the XGBoost regression model on known fault types. When predicting the leakage, it achieves a mean absolute error (MAE) of 0.27 l/min and a coefficient of determination (R^2) of 0.97. This high regression quality can be attributed to the fact that the same operating points and fault types were included in both the training and test data. As a next step, the dataset valve plate cavitation II is used for prediction, as it is still completely unknown for the model. The prediction results are depicted in Figure 4.3 b) and show a larger variation and a MAE of 0.89 l/min with R^2 of 0.28. This lower performance is expected when predicting samples with an unknown fault type.

Figure 4.3 c) shows the same kind of scatter plot for predictions made using the Multilayer Perceptron model. Prediction quality improves across the board when compared to the previous XGBoost predictions reaching a MAE of 0.12 l/min when predicting for known fault types. Performance on the unknown valve plate cavitation II damage also improves to 0.67 l/min of mean absolute error in Figure 4.3 d). Judging by this first analysis the Multilayer Perceptron model shows a higher adaptability to the unknown fault type combined with better performance metrics on known damages than the XGBoost model.

For the valve plate cavitation II predictions (Figure 4.3 b) and d)), both models clearly reveal distinct operating conditions through the dispersion of the vertical lines. Since leakage flow strongly depends on these parameters (speed, pressure and stroke/displacement), the objective is to enhance the prediction of leakages under unknown faults by incorporating operating condition information as model features.

4.4. Leakage flow prediction with known operating condition

Up to this point all features used by the models are derived from acceleration time series data. For the following analysis, the operating parameters speed (revolutions per minute), pressure (bar) and stroke (% of maximum displacement) are added to the feature space. Using these features, models can now deduce that, for example high speed and stroke values lead to higher overall leakage across all fault types. As can be seen in Figure 4.4 the addition of operating parameters improves prediction quality for the unknown valve plate cavitation II fault and both models. The XGBoost Regression model show the biggest improvement for predictions on this fault type, increasing R^2 values from below 0.4 to above 0.8 and almost halving MAE to 0.47 l/min. The Multi-Layer-Perceptron shows smaller improvements for the unknown fault type with a higher variance of R^2 values between folds.

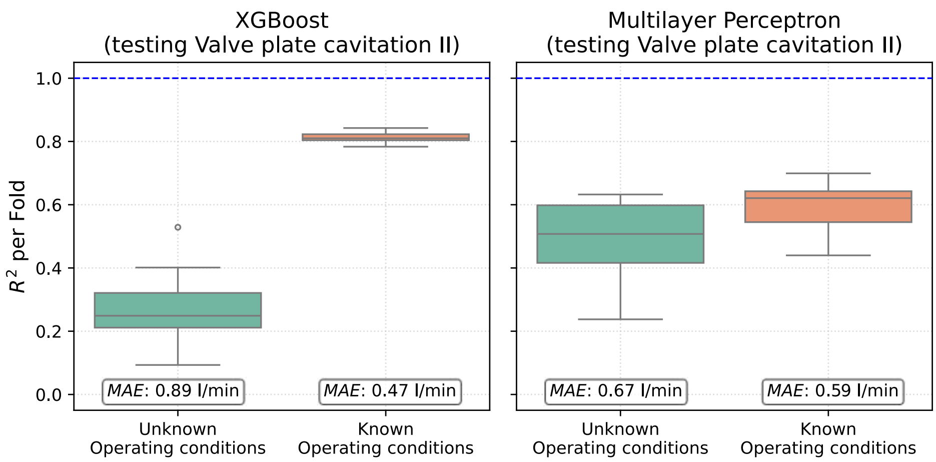

Figure 4.4 Improvement of leakage flow prediction when testing on unknown valve plate cavitation II fault and adding operating conditions to the feature space

4.5. Operating condition determination using vibration data

As demonstrated in the previous section, incorporating operating parameters (speed, pressure, and stroke) as features enhances the model's ability to predict leakage, especially for unknown fault types (Figure 4.4). However, this crucial information is often unavailable in retrofit scenarios, as dedicated sensors for pressure or stroke (i.e. swashplate angle sensors) may not exist or are unavailable to the CM system. The following analysis, therefore, investigates the feasibility of using the same high-frequency acceleration data to predict the operating conditions themselves.

Determining the operating condition from the vibration signal enables three key functions for a 'smart sensor'. First, it allows the recording of load history. By logging these conditions, load patterns become visible, which are important for assessing the system state. Second, it enables indirect monitoring. In applications without dedicated sensors, parameters like flow rate can be derived directly from the vibration data. Finally, it improves predictive maintenance. Analysing the connection between wear and load allows for a better prediction of potential pump failures. To achieve this, the following sections will test the ability of machine learning models to determine the operating condition (speed, pressure, and stroke value) using *only* the acceleration data, with the goal of realizing this functionality within a single sensor.

While the rotational speed can be analytically calculated from the frequency spectrum using spectral and order analysis, a similar analytical approach for pressure (e.g., correlating the amplitude of the first harmonic) was found to be unreliable. While a clear trend exists for the healthy pump, this correlation breaks down completely in the presence of faults, as the fault signatures interfere with the pressure-based amplitudes. This high complexity and variance, especially under fault conditions, lends itself to the application of machine learning models to predict speed and the remaining stroke parameter.

4.6. Operating condition prediction with ML-models

To predict the operating point from vibration data, both XGBoost and MLP regression models are employed. The feature space and the representative set of 28 operating conditions (Table 1) are identical to those used in the leakage prediction tasks.

A key challenge is that a deployed model will initially only be trained on data from a healthy pump. To simulate this, models are first trained exclusively on data from the undamaged pump (new) and then tested on data from the faulty components. As shown in Figure 4.5, this approach fails.

While the models achieve near-perfect accuracy on undamaged test data and XGBoost can still predict speed on faulty samples with $R^2 > 0.8$, the R^2 scores for pressure and stroke drop to negative values. This is attributed to the substantial impact of damage on the vibration spectra (see Section 4.2), which indicates the models have overfitted to the specific patterns of the healthy pump and cannot generalize to unknown fault conditions.

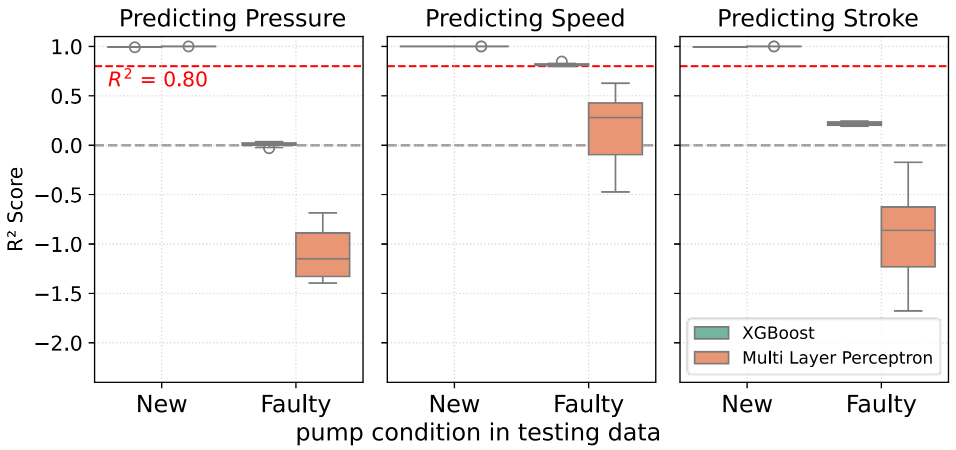

Figure 4.5 Prediction of pressure, speed and stroke with XGBoost and MLP on new and faulty condition training samples, with only new condition samples used for training

While training solely on healthy pump data yields limited generalization capability, extending the training set to include fault conditions significantly improves model performance. Figure 4.6 demonstrates this effect: when trained on a comprehensive dataset comprising all fault types (new condition, valve plate cavitation stages, axial clearance, and slipper wear), both XGBoost and MLP models achieve R^2 values exceeding 0.98 for all operating parameters.

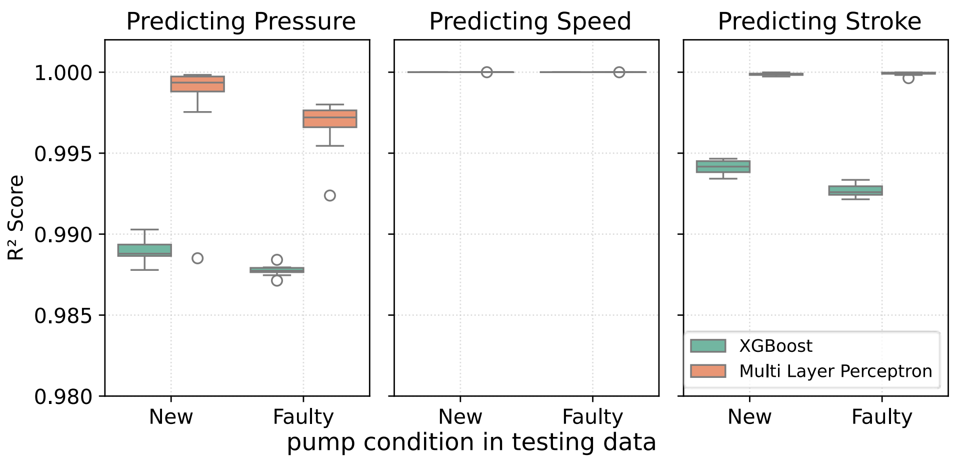

Figure 4.6 Prediction of pressure, speed and stroke with XGBoost and MLP on new and faulty condition training samples, with samples from all fault types used for training

This analysis confirms that accurate operating condition prediction solely from acceleration data is feasible, provided the training dataset captures the full range of expected pump conditions. Since industrial applications initially provide only healthy pump data, an adaptive training method is proposed. As the pump ages, the model is continuously retrained using its own predictions as 'quasi-true' labels, as long as internal performance metrics remain stable. This process could allow the model to adapt emerging degradation patterns and maintain high OC prediction quality over the pump's service life.

4.7. Proof of concept for Edge Deployment

Implementing the model directly on an embedded system enables a secure, self-contained "smart sensor". This study validates the deployment of the MLP model on a low-cost ESP32-S3 microcontroller (< 10 €). Although XGBoost showed slightly higher accuracy in certain analyses, the MLP was selected as it lends itself more readily to efficient conversion into C++ code. The final model architecture accepts 703 input features (700 base features and three operating parameters). It consists of 3 hidden layers with 238, 41, and 187 neurons respectively, and produces a single output using the tanh activation function.

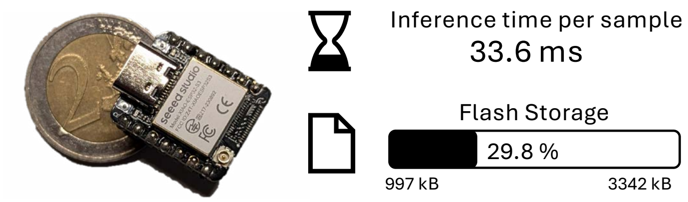

Figure 4.7 Validation of the 703-feature MLP model on a 10 €-ESP32 microcontroller: despite its complexity, the model occupies only 29.8 % of the storage and enables real-time predictions

For deployment, the optimized MLP model used in Section 4.4 (with knowledge of operating conditions) was converted to a C++ implementation via the ONNX format. Validation was performed by sending test samples to the microcontroller to measure real-world inference performance. As shown in Figure 4.7, the compiled code requires 997 kB, occupying only 29.8 % of the available flash storage. The average inference time is minimal at 33.6 ms per sample. This confirms that complex leakage prediction is viable on low-cost hardware.

5. SUMMARY AND OUTLOOK

This paper successfully demonstrated the foundational work for a low-cost, vibration-based "virtual sensor" for axial piston pumps and highly diverse operating conditions. The key findings are: 1) Predicting leakage flow for unknown fault types is only possible when the pump's operating conditions (OCs) are included as features. 2) These OCs (pressure, speed, and stroke) can be accurately predicted ($R^2 > 0.98$) using ML models from the same vibration signal. 3) This high accuracy is critically dependent on a robust training dataset that includes examples of various fault conditions; models trained only on healthy data fail to generalize. 4) The entire leakage flow prediction can be deployed on a low-cost (<10€) ESP32-S3 microcontroller with a minimal footprint (997 kB) and fast inference time (33.6 ms per sample), confirming its viability for edge computing.

The outlook for this work is focused on developing a fully autonomous, cascaded prediction system. This involves first predicting the OCs from the vibration signal (e.g., analytically calculating speed and using ML for pressure and stroke), and then feeding these predicted OCs into a second model to accurately estimate leakage. Future work will also explore more sophisticated feature engineering, such as using the heatmap analysis to define custom frequency bands and combining the leakage prediction with anomaly detection for event recognition. Also feature types fit for leakage flow prediction may not represent the best features for OC prediction further motivating more in-depth feature engineering approaches. For a fully standalone system, calculating features directly on the microcontroller is essential. Future work will explore the ESP32-S3's dual-core design to run feature extraction and model inference in parallel. On a systemic level, the connectivity of modern microcontrollers opens the possibility of using Federated Learning (FL) a method where only model parameters, not raw data, are shared—to train robust models across a fleet of pumps without compromising data security.

6. ACKNOWLEDGEMENT

The presented research activities are part of the project "Causal Condition Monitoring von Verdrängereinheiten durch Verschleiß- und Schadensdetektion in Korrelation zur Betriebs- und Lastsituation" (Ref. No. 22012 BR/1). The authors would like to thank the Fluid Power Research Fund of the VDMA for the funding and support.

7. APPENDIX

Table 3: Hyperparameter optimization search space (unspecified parameters use defaults). Parameter names match variable names in both models.

MLP Parameter	Distribution	XGBoost Parameter	Distribution
learning_rate	[1e-4, 1e-1]	learning_rate	[0.01, 1]
alpha (L2 Regulization)	[1e-5, 1e-1]	n_estimators	[1, 1000]
max_iter	[200, 1000]	max_depth	[1, 32]
batch_size	[16, 256]	min_child_weight	[1, 10]
activation	ReLU, Tanh	subsample	[0.5, 1.0]
Number of hidden layers	[1, 3]	colsample_bytree	[0.5, 1.0]
Layer Dimensionality	[32, 256]		

AUTHOR CONTRIBUTIONS

As the main author, S.K. was responsible for designing and refining the condition monitoring procedure, handling the edge implementation, and creating the written content and figures for the ML and CM sections. S.H. was responsible for the figures and written content regarding artificial fault creation, the experimental setup, and the leakage flow analysis. J.W. acted as the overall scientific supervisor.

REFERENCES

[1] S. Horn, S. Knoll, A. Shorbagy, M. Lenz, and J. Weber, "Economic Fault Detection for Axial Piston Pumps Under Variable Operating Conditions," International Journal of Fluid Power, vol. 26, no. 2, pp. 319–342, 2025, doi: 10.13052/ijfp1439-9776.2628.

[2] C. Dong, J. Tao, H. Sun, Q. Chao, and C. Liu, “Inverse transient analysis based calibration of surrogate pipeline model for fault simulation of axial piston pumps,” Mechanical Systems and Signal Processing, vol. 205, Dec. 2023, doi: 10.1016/j.ymssp.2023.110829.

[3] N. Kumar, R. Kumar, B. K. Sarkar, and S. Maity, “Condition monitoring of hydraulic transmission system with variable displacement axial piston pump and fixed displacement motor,” in Materials Today: Proceedings, Elsevier Ltd, 2019, pp. 9758–9765. doi: 10.1016/j.matpr.2020.09.327.

[4] N. Keller, A. Sciancalepore, and A. Vacca, “Demonstrating a Condition Monitoring Process for Axial Piston Pumps with Damaged Valve Plates,” International Journal of Fluid Power, vol. 23, no. 2, pp. 205–236, 2022, doi: 10.13052/ijfp1439-9776.2324.

[5] Y. Zhao, A. Jiang, W. Jiang, Y. Zhang, and S. Zhang, “Hydraulic pump fault diagnosis by modified slime mold algorithm optimized support vector machine,” Measurement, vol. 254, p. 117878, Oct. 2025, doi: 10.1016/J.MEASUREMENT.2025.117878.

[6] S. Tang, Y. Jiang, T. Zhou, Y. Zhu, and K. M. Lim, “Bayesian algorithm optimized deep model for multi-signal fault identification of hydraulic pumps,” Alexandria Engineering Journal, vol. 121, pp. 465–483, May 2025, doi: 10.1016/j.aej.2025.02.084.

[7] S. Tang, Z. Yan, Y. Xue, K. M. Lim, and Y. Zhu, “Deep transferable model with multi-head attention for small sample fault diagnosis of hydraulic pumps,” Engineering Applications of Artificial Intelligence, vol. 160, Nov. 2025, doi: 10.1016/j.engappai.2025.111808.

[8] S. Liu, Y. Zhang, C. Ai, and P. Liang, “Semi-supervised multi-adversarial domain-adaptive fault diagnosis for hydraulic pumps from pressure simulation data to experimental data,” Advanced Engineering Informatics, vol. 68, Nov. 2025, doi: 10.1016/j.aei.2025.103658.

[9] M. Liu, G. Kim, K. Bauckhage, and M. Geimer, “Data-Driven Virtual Flow Rate Sensor Development for Leakage Monitoring at the Cradle Bearing in an Axial Piston Pump,” Energies, vol. 15, no. 17, Sep. 2022, doi: 10.3390/en15176115.

[10] A. Schimke, O. Gnepper, and E. Boos, “Data-Driven Damage Assessment of Axial Piston Pumps,” in Proceedings of the 30th ICE IEEE/ITMC Conference on Engineering, Technology, and Innovation: Digital Transformation on Engineering, Technology and Innovation, ICE 2024, Institute of Electrical and Electronics Engineers Inc., 2024. doi: 10.1109/ICE/ITMC61926.2024.10794220.

[11] T. Wang, S. Liu, Y. Fu, and J. Ma, “Health state evaluation index system for aviation piston pumps based on asynchronous multi-source information fusion,” Results in Engineering, vol. 27, Sep. 2025, doi: 10.1016/j.rineng.2025.107076.

[12] Q. Chao, Y. Shao, C. Liu, and X. Yang, "Health evaluation of axial piston pumps based on density weighted support vector data description," Reliability Engineering and System Safety, vol. 237, Sep. 2023, doi: 10.1016/j.ress.2023.109354.

[13] Q. Chao, Y. Hu, and C. Liu, "Physics informed neural networks for detecting the wear of friction pairs in axial piston pumps," Reliability Engineering and System Safety, vol. 261, Sep. 2025, doi: 10.1016/j.ress.2025.111144.

[14] Ö. Özmen, C. Sinanoğlu, A. Caliskan, and H. Badem, "Prediction of Leakage from an Axial Piston Pump Slipper with Circular Dimples Using Deep Neural Networks," Chinese Journal of Mechanical Engineering (English Edition), vol. 33, no. 1, Dec. 2020, doi: 10.1186/s10033-020-00443-5.

[15] Y. Zhang, S. Wang, C. Zhang, H. Dui, and R. Chen, "Application of physics-informed machine learning in performance degradation and RUL prediction of hydraulic piston pumps," Reliability Engineering and System Safety, vol. 261, Sep. 2025, doi: 10.1016/j.ress.2025.111108.

[16] C. C. Walani and W. Doorsamy, "Edge vs. Cloud: Empirical Insights into Data-Driven Condition Monitoring," Big Data and Cognitive Computing, vol. 9, no. 5, May 2025, doi: 10.3390/bdcc9050121.

[17] S. Horn, A. Shorbagy, M. Lenz, S. Knoll, M. Winter, and Prof. J. Weber, "CAUSAL CONDITION MONITORING FOR AXIAL PISTON PUMPS: PREDICTION OF DIFFERENT FAULTS UNDER VARIABLE OPERATING CONDITIONS," pp. 181–204, Oct. 2024, doi: 10.13052/RP-9788770047456.014.

[18] M. Christ, N. Braun, J. Neuffer, and A. W. Kempa-Liehr, "Time Series FeatuRe Extraction on basis of Scalable Hypothesis tests (tsfresh – A Python package)," Neurocomputing, vol. 307, pp. 72–77, Sep. 2018, doi: 10.1016/J.NEUCOM.2018.03.067.

[19] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, vol. 13-17-August-2016, pp. 785–794, Mar. 2016, doi: 10.1145/2939672.2939785.

[20] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[21] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2623–2631, Jul. 2019, doi: 10.1145/3292500.3330701.

[22] Y. Duensing, "Zustandsüberwachung hochdynamischer bidirektionaler Axialkolbenpumpen mittels maschinellen Lernens," 2025. doi: 10.18154/RWTH-2025-10412.

Biographies

Simon Knoll

Simon Knoll graduated from the Technical University of Dresden in 2025 with a diploma degree in Mechanical Engineering, specializing in fluid power. During his studies, he focused on research related to the slipper–swashplate interface in axial piston pumps and using machine learning with sensor data to monitor pump conditions. He is currently working as a research assistant at the Chair of Fluid-Mechatronic Systems.

Svenja Horn

Svenja Horn received her diploma degree in Aircraft Construction from the Technical University of Dresden in 2018. Since then, she has been working as a research assistant at the Chair of Fluid-Mechatronic Systems, focusing on improving the efficiency and service life of axial piston pumps. Her research concentrates on tribological phenomena at the slipper–swashplate interface, which she investigates using simulations, a hydrostatic tribometer and pump measurements. Furthermore, she applies machine learning algorithms for pump condition monitoring and lifetime prediction.