

A Concept for Accelerating Long-Term Prototype Testing Using Anomaly Detection and Digital Twins

Vincent Nebel¹^a, Pia Goßrau-Lenau², Harshvardhan Agarwal¹^b and Dirk Werth¹

¹August-Wilhelm Scheer Institute for Digital Products and Processes gGmbH, Saarbrücken, Germany

²Technology Center Drives, Miele & Cie. KG, Euskirchen, Germany

Keywords: Anomaly Detection, Digital Twin, Prototyping, Machine Learning.

Abstract: Developing mechanical components, especially complex assemblies like pumps, is a resource and time intensive process. Testing pump prototypes for long-term durability is critical to ensure error-free operation of the final product. Prototypes undergo material and operational tests to determine their expected lifespan, focusing on defects caused by material degradation and water contamination. Long-term tests, lasting months, are necessary to simulate real-world conditions, but limited test bench capacities create bottlenecks, restricting material experimentation. Moreover, monitoring the internal state of pumps during tests is challenging. Undetected defects can worsen or trigger secondary issues, complicating the root cause analysis, which provides valuable information for further product improvements. To address these challenges, a digital twin that integrates geometry and material data, simulations, and sensor measurements was developed. This twin is used as data source for machine learning based anomaly detection, allowing tests to stop sooner and preventing further damage when first signs of a defect are detected. A modular serverless architecture is used to host the model inference on the cloud, improving resource usage and scalability as well as reducing operational costs.

1 INTRODUCTION


The development of mechanical components is a resource-intensive undertaking. It requires a lot of time and effort to design different parts, select materials, and construct working prototypes. To ensure error free operation and longevity of the finished product, the prototypes then have to be put through a variety of tests. Many of these tests require expensive equipment and specific training which limits the possibilities for parallelization. Especially for long-term tests, this lack of parallelization can constitute a significant bottleneck in the product development pipeline.


Complex assemblies, like pumps, consist of many individual parts and hence are exposed to multiple different kinds of faults. To test such assemblies, it is common to use a combination of material and operating tests. Since pumps for household appliances need to run maintenance and error free for many years, the central goal during prototype testing is to determine the expected lifetime of a pump. In our case, most of the defects, that we are interested in, are caused

by material degradation or water contamination over time. The speed and severity of the degradation is influenced by the forces that act on the materials during operation as well as by the contamination of the pumped water. In order to replicate these influences when testing a new prototype, long-term tests, which can last multiple months, are necessary. Since for every iteration of a prototype a significant amount of pumps needs to be tested and only a limited number of test benches is available at a given time, this represents a severe bottleneck for the development of new product revisions and limits the variety of materials that can be tested.

Since it is not possible to monitor the inside of a pump during operation testing, a defect might become much more severe or even cause other issues before it is noticed by the test operator. This can make the identification of the root cause of the defect very challenging or even impossible in some cases. Therefore, it is very important to stop the testing procedure as soon as possible once a fault occurs and examine the pump to gain insights on possible improvements to the pump construction or choice of materials.

To address these issues we propose using a digital twin to aggregate geometric models, material data,

^a <https://orcid.org/0009-0009-2713-802X>

^b <https://orcid.org/0000-0001-6527-7658>

simulation results, and sensor measurements during testing. The digital twin then acts as the single source of data for detecting anomalies and classifying different kinds of defects in the sensor measurements using machine learning (ML). This enables the operator to stop the test earlier if an issue is detected, hence, decreasing the probability of additional defects and freeing up valuable test bench capacity. Furthermore, knowledge about the degradation behavior of the used materials is integrated to improve fault detection and lifespan estimation. The aggregation of all relevant data in a central digital twin also simplifies data access for manual inspection of the test results, streamlining the product development and testing processes.

To ensure cost-effective operation and scalability, the digital twin is hosted on the cloud and a serverless architecture for model inference was developed. The presented approach reduces the maintenance overhead and simplifies the integration of additional test benches, sensors, and analysis components by prioritizing modularity in the architecture design.

The remainder of this paper is structured as follows. Section 2 gives an overview of the general construction of the pumps and the different defects that can occur. Then, Section 3 discusses related work on ML anomaly detection and serverless cloud computing. After that, the proposed concept is described in detail in Section 4. Finally, Section 5 concludes this work by summarizing and providing an outlook on the next steps.

2 PUMP PROTOTYPE TESTING

In this work, the focus is on dry-running centrifugal dishwasher pumps driven by a brushless permanent synchronous motor (PMSM).

As shown in Figure 1, the pump complex of these specific pumps can be divided into three main assemblies:

1. pump cover with tubular heating element and pressure switch
2. pump housing with rotary vane mechanism and water gate
3. main pump drive with pump impeller

The drive is controlled by an electronic frequency converter which allows the pump to run speed controlled.

Through rotor turning, the pump impeller, mounted on the rotor shaft, is moved. It has different blade geometries in the intake and outlet areas. Rotation of the impeller creates a pressure gradient from the center of the impeller to the outside, whereby the




Figure 1: Assemblies of the pump complex.

pressure increases in radial direction (see Figure 2). The higher the speed, the higher in general the centrifugal force and the associated delivery pressure of the pump. In this way, the water enters the connected pump housing and is pumped from there into an open outlet nozzle and reaches different levels of the dishwasher.




Figure 2: Main pump drive with impeller and approximate flow paths.

The pump complex is subject to certain service life requirements. For this reason, there are long-term test benches for the pumps in which they are tested (in this case separately from the device) for their service life. The primary aims of the testing are particularly based on the development stage of the pump and range from the qualification of prototypes to pilot series releases and series releases.

There are various reasons for carrying out such service life tests of pumps specially on test benches developed for this purpose:

- fewer interactions than expected in the device (individual pump system can be examined more specifically)
- measuring points are more accessible
- operation of individual components outside the specification is possible (limit or overload tests)
- control of the test environment

Various faults are signs of wear that could occur during the service life of a pump. Some typical ones are:

- defects in the hydraulic (e.g. concerning the impeller)
- deposits (of foreign materials)
- signs of ageing on seals
- damage to bearings

These can have a number of undesirable effects, e.g.:

- loss of function of the pump (worst case)
- increased power consumption (e.g. due to difficulty of movement)
- reduced water delivery
- acoustic abnormalities
- leakages

The aim is to minimize errors and (material) wear as far as possible in advance. The physical long-term test is a way of testing different prototypes and providing direct feedback to the R&D department. Of course, the number of test benches is limited and only a certain number of variants can be tested in a certain amount of time. This makes it quite important to become aware of anomalies as early as possible using suitable data analysis or machine learning and, ideally, to be able to predict certain correlations or parameters.

2.1 Relevant Measured Variables

Suitable data is required for ML based anomaly detection. Relevant (measured) variables that allow the condition of a pump to be assessed must be identified. The following measured variables are considered to be particularly relevant for this pump type:

- motor phase current (because of the PMSM drive proportional to the torque when the pump is in control mode)
- pressure (difference between intake and outlet of the pump)
- vibration / structure-borne noise (usually early observation of abnormalities possible)
- mechanical static friction torque of the sealing system (measure for ease of movement of the sealing system)

All these values are recorded by the test benches at different, for the specific measurement appropriate, frequencies and uploaded to the digital twin either immediately or in batches.

3 RELATED WORK

This section outlines various state-of-the-art techniques for anomaly detection and serverless machine learning inference. A large variety of model architectures and methods has been proposed to detect anomalies in different kinds of data. In the following widely used methods and applications related to pump testing are outlined.

The Local Outlier Factor (LOF) method (Breunig et al., 2000) developed by Breunig et al. and extensions like the Cluster-Based Local Outlier Factor (CBLOF) (He et al., 2003) work by using a measure for the local density of a dataset and classifying outliers as points with substantially lower density than their neighbors. While LOF uses the distance to the k nearest neighbors to determine the local density, CBLOF extends this idea by clustering the points (e.g. using k -means) and calculating the density based on cluster size and distance to other clusters.

Other approaches are to estimate Minimum Covariance Determinant (MCD) which provides a robust clustering of the data that can be used for outlier detection (Hardin and Rocke, 2004) or to use a Histogram-Based Outlier Score (Goldstein and Dengel, 2012) which estimates the density based on the frequency of samples in a bin of the univariate histogram of each feature.

Isolation Forests (Liu et al., 2008) are also used for anomaly detection. They are particularly useful for high-dimensional datasets. In this case anomalies are determined via the path length needed to isolate a specific sample averaged over multiple trees. Another popular method is Bayesian Online Change Point Detection (BOCPD) (Adams and MacKay, 2007) which is used to detect change points in time series data. BOCPD works by maintaining a probability distribution over the run length, which is the time elapsed since the most recent change point. When a new observation comes in, this distribution is updated using Bayes' theorem for inverting conditional probabilities. This method can be adapted to manage time series with a non-constant baseline (e.g. different modes of operation) and has been used to monitor the vibration of centrifugal pumps in HVAC systems (Lu et al., 2020).

Generative Adversarial Networks (GAN) can also be used for unsupervised anomaly detection as, e.g., shown by Liu et al. (Liu et al., 2020). A GAN consists of two competing neural networks: a generator and a discriminator. The generator tries to generate realistic artificial data, while the discriminator distinguishes between real and artificial data. Using this method, anomalies are identified by the discriminator

as non-real values.

Furthermore, different encoder-based architectures have been proposed for anomaly detection (Zhou and Paffenroth, 2017; Kieu et al., 2019; Siegel, 2020; Abhaya and Patra, 2023). In general, these work by producing a latent representation of the input using an encoder and reconstructing the inputs through decoding. The difference between original inputs and decoder outputs is then used as an anomaly score.

Anomaly detection has also been used for the specific variables discussed in Section 2.1, which are especially relevant for pump testing.

For detecting pressure anomalies in axial piston pumps Jiang et al. used a composite method that combines Isolation Forest with a random convolutional kernel for feature extraction and dynamic time warping (DTW) to effectively handling time series of varying lengths (Jiang et al., 2023). They validated their method using a specialized test bench for pump failure simulations and showed that it outperformed traditional methods. In (Dong et al., 2023) the authors tackle a similar challenge using a subsequence time series (STS) clustering-based approach. Their method is a two step process consisting of a step for identifying a "norm cluster", that represents the normal behavior of the time series, by performing multiple STS clustering operations and an anomaly subsequence clustering step which clusters the remaining subsequences to detect anomalies. Additionally, DTW was utilized to enable the detection of variable-length sub-sequences and enhance the robustness against variations in different operational parameters. Their method compared favorably to other methods, such as Isolation Forest and LOF, particularly in scenarios with recurrent anomalies and varying loads.

For working with vibration data to detect gear and bearing faults in helicopters a semi-supervised approach, that relies on training models using only healthy signals due to the scarcity of faulty data in real-world applications, was proposed in (Vos et al., 2022). Long-strong-term memory regression was utilized to remove deterministic components from the signal and the residual signal classified using a one class support vector machine. This method was shown to be suitable for early fault detection. In (Hu et al., 2022) vibration vectors, which consist of amplitude and phase information of the measured vibration, are used as the primary indicator for detecting anomalies in the vibrations caused by a steam turbine and a steam feed pump. The vibration vectors are extracted using Fast Fourier Transform-Based Order Analysis and Support Vector Data Description (Tax and Duin,

2004) is used to learn an acceptance region that can self-evolve to accommodate changes in machine conditions. This approach outperformed other methods like MCD and Isolation Forest, especially when the data distribution was non-Gaussian.

Ribeiro et al. conducted research on detecting anomalies in the screw-tightening process using angle-torque pairs. In (Ribeiro et al., 2021) they compared three different unsupervised models (LOF, Isolation Forest, and an Autoencoder based approach) and used a supervised Random Forest as a benchmark. A realistic rolling window approach was employed to evaluate models over time, simulating real-world use-cases with continuous data flow. They found that the Isolation Forest Approach could compete with the supervised approach in terms of accuracy while the Autoencoder performed slightly and LOF noticeably worse. Based on those findings, the authors concluded that in a real world application the Autoencoder is to be preferred due to its better computational efficiency. They confirmed this conclusion in a follow-up study (Ribeiro et al., 2022) comparing Isolation Forest and Autoencoder on a larger dataset. In this scenario both models achieved similar accuracy, but the Autoencoder required 2.7 times less time for training and 3.0 times less time for inference.

Initial investigations into the presented use-case will focus on simple cluster-based methods, like CBLOF, to establish a performance baseline. However, the existing work on related use-cases suggest that significant improvements over that baseline might be possible using more sophisticated approaches like the the method by Jiang et al. (Jiang et al., 2023). Hence, the adaption of different methods to the specific use-case presented in this work and an in-depth performance comparison is planned.

Serverless computing has gained popularity over the last decade as it provides a cost-effective option for hosting cloud-based services with minimal management overhead and built-in scalability (Hassan et al., 2021; Shafiei et al., 2022). While the solutions from large companies like Amazon, Microsoft, and Google are most commonly used, open source solutions like Apache OpenWhisk and custom implementations are actively worked on and researched (cf. (McGrath and Brenner, 2017; Djemame et al., 2020)). Furthermore, usage of serverless compute platforms specifically for machine learning has grown in recent years (Barrak et al., 2022). Even though cold starts and GPU access still pose issues for serverless ML deployments (Barrak et al., 2022; Kojs, 2023), it has become a popular option, especially for model inference. Moreover, frameworks specialized for machine learning inference like BATCH by Ali et al. (Ali et al.,

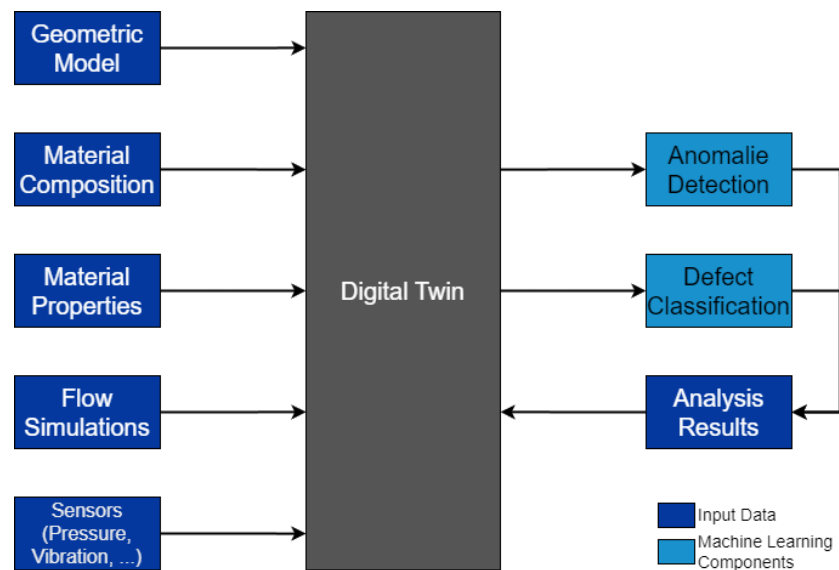


Figure 3: General Overview of the Digital Twin showing the different Input Data and Machine Learning Components.

2020) which utilizes an optimizer to obtain inference tail latency guarantees and enable adaptive batch processing, or INFless, a serverless platform purpose built for the ML domain (Yang et al., 2022) promise performance improvements over the available general purpose solutions from the large cloud providers. For the purpose of this work it was determined that a custom serverless architecture that uses Azure Functions and is discussed in more detail in Section 4 is the best choice because it allowed to use existing expertise and reduces the complexity future maintenance.

4 CONCEPT

In this chapter the details of our solution concept are described, the motivation and reasoning behind the choices made are outlined, and an overview of the developed digital twin and cloud architecture is provided.

In order to alleviate the impact of the testing bottleneck on the product development and qualification pipeline, accelerating the tests is essential. Since accelerating the testing cycles themselves can place additional strain on the materials, this would make it much more difficult to correlate the testing results to real world use. Shortening the long-term tests also is not an option, because reaching a certain minimum runtime is mandatory for product qualification. Consequently, stopping test runs in which the pump did fault or is showing signs of a defect that will lead to a fault is the best available option for accelerating the testing. This approach has the additional benefit of simplifying the identification of the root cause of

a defect by preventing the development of follow-up defects.

To achieve this, machine learning based anomaly detection is used. For applying such an approach in a product development setting where the used materials, and therefore certain characteristics of the pumps, change regularly, aggregating a broad spectrum of information on the pump prototypes and ensuring high data quality are extremely important. For addressing these challenges, we propose using a digital twin of the pumps that combines the information about the pump itself with detailed data on the materials used and measurements from the long-term tests. The digital twin is stored in a cloud-based data platform that allows fast and simple access to the data for manual inspection or the creation of test reports and also provides various endpoints for automatic upload of data from the test benches and connecting the ML components of the system. Figure 3 shows an overview of the data that is aggregated in the digital twin. Besides the geometric model of the pump and corresponding material compositions of the individual parts, the digital twin also contains in depth information on the properties and aging characteristics of the different materials and flow simulations that provide expected values for the pressure measurements under different conditions. These data are then combined with the real world sensor measurements from the test bench. The ML components receive this data from the digital twin and store their analysis results in it.

For achieving good results using machine learning, the quality of the data plays a major role. The data for the use-case presented in this work is acquired using purpose built test benches which

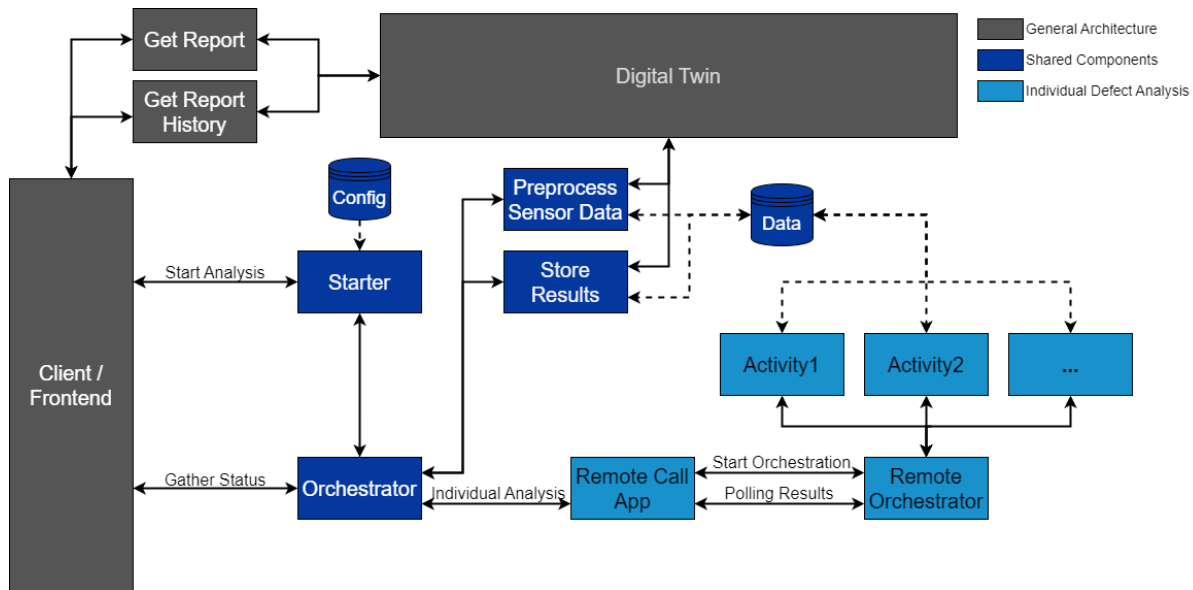


Figure 4: Structure of the Cloud Architecture used for Model Inference.

record the sensor data at different frequencies depending on the specific measurement. Low frequency data is immediately uploaded to digital twin, while high frequency measurements, e.g. the phase current, are stored locally using CSV files and uploaded as batches in predefined time intervals. To mitigate sensor drift, a calibration step is performed before testing. Additionally, general information on the testing conditions like water contamination, environment temperature, and total runtime of the pump are integrated to better represent expected changes over the entire test time and, thus, improve the detection accuracy of abnormal pump behavior. Where necessary, denoising and data transformations (e.g. normalization or fast Fourier transform for phase current measurements) will be performed by the cloud architecture during the preprocessing step.

To run the ML inference a serverless approach was chosen to reduce operational cost and maintenance effort. The models and data preprocessing pipelines are deployed using Azure Functions, a serverless compute solution developed and provided by Microsoft. To efficiently work with the resource and runtime limitations that follow from using Azure Functions, the custom architecture depicted in Figure 4 was developed. The architecture enables the modularization of the analysis pipeline across multiple function apps to avoid resource bottlenecks. Furthermore, the chosen orchestration approach allows the architecture to effectively handle function timeouts, which are especially likely when one of the functions has to perform a cold start.

When an analysis is requested by the client, which

can either be done manually or automatically at predefined points during a long-term test, a starter function gathers the configuration data for the specific analysis and starts the orchestration. Via an endpoint returned by the starter function, the client can then request status updates from the orchestrator. The orchestrator triggers the data preprocessing pipeline, which obtains all required data from the respective endpoints of the digital twin platform, and the remote function apps that handle further data processing and inference for the different models. A separate model is used for anomaly detection for each kind of sensor data (cf. Section 2.1) as well as for the defect classification, allowing for even more modularity. The preprocessed data and intermediate results are stored in a data cache until all steps of the analysis are complete and the final results are stored in the digital twin.

5 CONCLUSION AND OUTLOOK

In this work the need for accelerating long-term tests during product development was demonstrated in the use-case of household appliance pumps. It was argued that the best option for achieving this is to employ anomaly detection to detect defects during testing early, allowing to shorten the test time for defect pumps and aid the classification of a defect and the reconstruction of its cause. Furthermore, a cloud architecture that handles model inference using a serverless approach and relies on a cloud based digital twin as data source for the machine learning components was presented. The developed architecture can

handle timeouts through the chosen orchestration approach and alleviate the impact of memory limitations through its modularity. Furthermore, the modularity ensures the expandability of the system, e.g. for integrating new sensors or classifiers for additional kinds of defects.

While the development and deployment of the digital twin and other general architecture components is already completed, the data acquisition and model training is still in an early stage. Initial manual investigation of the vibration and phase current data show correlations between the measurements and the state of the pump, but further experiments regarding the pressure measurements are necessary as it proves to be difficult to identify small scale leakages based on the acquired data. At the time of writing, baseline measurements using correctly functioning pumps have been completed and long-term test for the first defect category have started.

Once the data acquisition phase is completed, a variety of different models and methods has to be tested to determine the best approach for the presented use-case and additional long-term tests have to be executed to validate the obtained results. Future investigations should also focus on testing the transferability of this approach to other product categories, e.g. turbines or gearboxes. As a first step towards generalization, it is planned to apply the approach for other models of pumps once satisfactory results have been achieved for the pumps addressed in this work. Since the characteristics of the vibration and other measurements change based on the design of the pump, some amount of retraining will be necessary to transfer the results to different pumps or other product categories.

ACKNOWLEDGEMENTS

The authors thank the German Federal Ministry for Economic Affairs and Climate Action (BMWK) for financial support of the project ProDiNA through project funding FKZ 01MN23016A. The project ProDiNA is a joint effort of the August-Wilhelm Scheer Institute for Digital Products and Processes gGmbH, the Miele & Cie. KG, the adesso SE, the dive solutions GmbH, and the Leibniz-Institute for New Materials gGmbH.

REFERENCES

- Abhaya, A. and Patra, B. K. (2023). An efficient method for autoencoder based outlier detection. *Expert Systems with Applications*, 213:118904.
- Adams, R. P. and MacKay, D. J. C. (2007). Bayesian online changepoint detection.
- Ali, A., Pincioli, R., Yan, F., and Smirni, E. (2020). Batch: Machine learning inference serving on serverless platforms with adaptive batching. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15.
- Barrak, A., Petrillo, F., and Jaafar, F. (2022). Serverless on machine learning: A systematic mapping study. *IEEE Access*, 10:99337–99352.
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. (2000). Lof: identifying density-based local outliers. *SIGMOD Rec.*, 29(2):93–104.
- Djemame, K., Parker, M., and Datsev, D. (2020). Open-source serverless architectures: an evaluation of apache openwhisk. In *2020 IEEE/ACM 13th International Conference on Utility and Cloud Computing (UCC)*, pages 329–335.
- Dong, C., Tao, J., Chao, Q., Yu, H., and Liu, C. (2023). Subsequence time series clustering-based unsupervised approach for anomaly detection of axial piston pumps. *IEEE Transactions on Instrumentation and Measurement*, 72:1–12.
- Goldstein, M. and Dengel, A. (2012). Histogram-based outlier score (hbos): A fast unsupervised anomaly detection algorithm. *KI-2012: poster and demo track*, 1:59–63.
- Hardin, J. and Rocke, D. M. (2004). Outlier detection in the multiple cluster setting using the minimum covariance determinant estimator. *Computational Statistics & Data Analysis*, 44(4):625–638.
- Hassan, H. B., Barakat, S. A., and Sarhan, Q. I. (2021). Survey on serverless computing. *Journal of Cloud Computing*, 10(1):39.
- He, Z., Xu, X., and Deng, S. (2003). Discovering cluster-based local outliers. *Pattern recognition letters*, 24(9–10):1641–1650.
- Hu, D., Zhang, C., Yang, T., and Chen, G. (2022). An intelligent anomaly detection method for rotating machinery based on vibration vectors. *IEEE Sensors Journal*, 22(14):14294–14305.
- Jiang, W., Ma, L., Zhang, P., Zheng, Y., and Zhang, S. (2023). Anomaly detection of axial piston pump based on the dtw-rck-if composite method using pressure signals. *Applied Sciences*, 13(24):13133.
- Kieu, T., Yang, B., Guo, C., and Jensen, C. S. (2019). Outlier detection for time series with recurrent autoencoder ensembles. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2725–2732. International Joint Conferences on Artificial Intelligence Organization.
- Kojs, K. (2023). A survey of serverless machine learning model inference.
- Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2008). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422.
- Liu, Y., Li, Z., Zhou, C., Jiang, Y., Sun, J., Wang, M., and He, X. (2020). Generative adversarial active learning

- for unsupervised outlier detection. *IEEE Transactions on Knowledge and Data Engineering*, 32(8):1517–1528.
- Lu, Q., Xie, X., Parlikad, A. K., and Schooling, J. M. (2020). Digital twin-enabled anomaly detection for built asset monitoring in operation and maintenance. *Automation in Construction*, 118:103277.
- McGrath, G. and Brenner, P. R. (2017). Serverless computing: Design, implementation, and performance. In *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pages 405–410.
- Ribeiro, D., Matos, L. M., Cortez, P., Moreira, G., and Pilastrri, A. (2021). A comparison of anomaly detection methods for industrial screw tightening. In *Computational Science and Its Applications–ICCSA 2021: 21st International Conference, Cagliari, Italy, September 13–16, 2021, Proceedings, Part II 21*, pages 485–500. Springer.
- Ribeiro, D., Matos, L. M., Moreira, G., Pilastrri, A., and Cortez, P. (2022). Isolation forests and deep autoencoders for industrial screw tightening anomaly detection. *Computers*, 11(4):54.
- Shafiei, H., Khonsari, A., and Mousavi, P. (2022). Serverless computing: A survey of opportunities, challenges, and applications. *ACM Comput. Surv.*, 54(11s).
- Siegel, B. (2020). Industrial anomaly detection: A comparison of unsupervised neural network architectures. *IEEE Sensors Letters*, 4(8):1–4.
- Tax, D. M. and Duin, R. P. (2004). Support vector data description. *Machine Learning*, 54(1):45–66.
- Vos, K., Peng, Z., Jenkins, C., Shahriar, M. R., Borghesani, P., and Wang, W. (2022). Vibration-based anomaly detection using lstm/svm approaches. *Mechanical Systems and Signal Processing*, 169:108752.
- Yang, Y., Zhao, L., Li, Y., Zhang, H., Li, J., Zhao, M., Chen, X., and Li, K. (2022). Infless: a native serverless system for low-latency, high-throughput inference. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS ’22*, page 768–781, New York, NY, USA. Association for Computing Machinery.
- Zhou, C. and Paffenroth, R. C. (2017). Anomaly detection with robust deep autoencoders. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 665–674.