

Nr. 81 16. June 2021

Preprint-Series: Department of Mathematics - Applied Mathematics

Machine Learning for Nondestructive Wear Assessment in Large Internal Combustion Engines

C. Angermann, S. Jonsson, M. Haltmeier, A. Moravova, C.Laubichler, C. Kiesling, M. Kober, W. Fimml



Technikerstraße 13 - 6020 Innsbruck - Austria Tel.: +43 512 507 53803 Fax: +43 512 507 53898 https://applied-math.uibk.ac.at

Machine Learning for Nondestructive Wear Assessment in Large Internal Combustion Engines

Christoph Angermann¹, Steinbjörn Jónsson², Markus Haltmeier¹, Adéla Moravová¹, Christian Laubichler³, Constantin Kiesling³, Martin Kober³, and Wolfgang Fimml²

¹Department of Mathematics, University of Innsbruck, Technikerstraße 13, 6020 Innsbruck, Austria applied-math.uibk.ac.at

²INNIO Jenbacher GmbH & Co OG, Achenseestrasse 1-3, 6200 Jenbach,

Austria

www.innio.com/en ³LEC GmbH, Inffeldgasse 19, 8010 Graz, Austria www.lec.at

June 16, 2021

Abstract

Digitalization offers a large number of promising tools for large internal combustion engines such as condition monitoring or condition-based maintenance. This includes the status evaluation of key engine components such as cylinder liners, whose inner surfaces are subject to constant wear due to their movement relative to the pistons. Existing state-ofthe-art methods for quantifying wear require disassembly and cutting of the examined liner followed by a high-resolution microscopic surface depth measurement that quantitatively evaluates wear based on bearing load curves (also known as Abbott-Firestone curves). Such reference methods are destructive, time-consuming and costly. The goal of the research presented here is to develop nondestructive yet reliable methods for quantifying the surface condition. A deep-learning framework is proposed that allows computation of the bearing load curves from reflection RGB images of the liner surface that can be collected with a wide variety of simple imaging devices, without the need to remove and destroy the investigated liner. For this purpose, a convolutional neural network is trained to predict the bearing load curve of the corresponding depth profile from the collected RGB images, which in turn can be used for further wear evaluation. Training of the network is performed using a custom-built database containing depth profiles and reflection images of liner surfaces of large gas engines. The results of the proposed method are visually examined and quantified considering several probabilistic distance metrics and comparison of roughness indicators between ground truth and model predictions. The observed success of the proposed method suggests its great potential for quantitative wear assessment on engines during service directly on site.

Keywords— large gas engine, cylinder liner wear, condition monitoring, bearing load curve, modality transfer learning, convolutional neural network

1 Introduction

Digitalization offers a wide range of promising tools for large internal combustion engines. While sensors are needed for recording data, networks and other IT infrastructure provide access and connect the information. From electronic control units of specific engine components via the entire engine control system to application software for monitoring, the data is repeatedly processed and analyzed. Today, original equipment manufacturers or third-party suppliers are already providing solutions that instantaneously obtain insight into the overall performance of an engine [1, 2, 3, 4, 5, 6, 7] or the condition of individual key engine components such as bearings [8], injectors [9] or turbochargers [10]. In addition, with modern data analysis methods and advanced simulation techniques such as digital twins, there is great potential to improve condition monitoring systems that were previously based on expert systems only [11]. All these factors, their interrelations as well as the knowledge gain from modern analysis or simulation techniques give rise to a concept which can be referred to as a "digital engine".

When online condition monitoring systems are given a predictive character, e.g. by using advanced data analysis methods such as machine learning, there is the potential to increase the lifetime of an engine, to avoid unwanted downtime due to maintenance and therefore to reduce the engine's carbon footprint and its total cost of ownership [12]. Carvalho et al. [13] and Zhang et al. [14] provide recent reviews about data-driven approaches for predictive maintenance in general. Both include studies applying traditional machine learning techniques such as support vector machines or random forests and neural network-based deep learning approaches, respectively. Lei et al. [15] consider transfer learning techniques in their review of methods for intelligent fault diagnosis. Although there are many valid machine learning approaches for implementing a predictive maintenance system in a digital engine, its success still heavily depends on the quality of the available data and the structure of the learning system.

1.1 Wear Monitoring of Engine Components

While information on operating parameters, performance and ambient conditions can be evaluated permanently, position-based measurement of the wear of most components during engine operation is currently not possible. Some online wear monitoring systems focus on indirect effects of wear, for example by detecting wear debris in the lubricant oil system [16]. In many situations, however, wear assessment requires disassembly and destruction of the engine component to be examined. Therefore, the goal of the research presented here is to use fast yet reliable and informative methods in combination with advanced machine learning algorithms to evaluate the wear condition of engine components without the need to remove and destroy them. Machine learning-based systems acquire their knowledge from high-dimensional raw data [17]. Recent research shows that in combination with data-driven machine learning techniques, monitoring systems can build an effective wear prediction model [16]. In addition, machine learning also provides the opportunity to permanently approximate computationally expensive wear predicting simulation techniques and to conduct online wear predictions of machinery parts [18]. The aim of this research is to demonstrate the plausibility and advantages of combining simple measurement methods with advanced machine learning techniques for nondestructive and fast wear assessment. The main focus of this study lies on the inner surface/wear inspection of cylinder liners in large gas engines. Nevertheless, it should be emphasized that the proposed transfer learning framework can be similarly applied to advance the assessment of various other components in large internal combustion engines.

1.2 Cylinder Liner Wear Assessment

Due to the movement of the piston during the combustion cycle, cylinder liners of internal combustion engines are subject to constant wear. Wear decreases the volumetric efficiency of the engine and increases blow-by; oil consumption; power loss and HC, CO, CO₂ and NO_x-emissions [19, 20, 21]. In addition, wear impairs the hydrodynamic support of the piston rings and thus increases the risk of a fatal engine failure [22]. There are multiple on-site and external methodologies that characterize the surface structure including stylus instruments, white light interferometry [23], scanning electron microscopes, bore measurements, thin layer activation [24] and 3D optical confocal measurements [23, 25]. Various methods have been proposed for cylinder liner wear assessment including component mass losses [26], precise coordinate measurements [27], volume losses [28, 29] and repeated measurements of component positions [30].

A common approach in the large engine industry is to evaluate wear by taking spatial measurements of an unworn and a worn area of the cylinder liner and then to compare the corresponding bearing load curves (BLC) after height compensation [29, 31, 32, 33]. BLC is also known as bearing surface curve or Abbott-Firestone curve. These spatial measurements are performed with an advanced optical device that generates a high-resolution microscopic surface depth profile that allows for calculation of the corresponding BLC. Since the surfaces of interest are located on the inside of the cylinder liner and the optical device is mounted at a fixed measurement station due to its cumbersome size, it is necessary to dismount the cylinder liner and to cut it into segments containing the respective surface areas. A specific issue with this method is that the components are destroyed and cannot be re-used after measurement. Thus a single component cannot be measured again at a later time when it has accumulated more wear. In addition, handling the sophisticated measurement equipment requires a high level of technical understanding of roughness assessment and software-specific training on the measurement device. In summary, the current measurement pipeline is a time- and resourceintensive process chain, which makes it unsuitable for constant evaluation of production quality and long-term studies on the wear of cylinder liner components. The proposed liner surface evaluation method presented in this research does not require removal and destruction of the examined cylinder liner.

1.3 Modality Transfer Learning

Assessment methods that rely only on a simple handheld optical device have the potential to enable on-site condition assessment of liners on series engines. The current state-of-the-art technology integrates optical sensors and especially digital image sensors into comparatively small and inexpensive devices such as cell phones, compact cameras and handheld microscopes. These devices are capable of producing a large amount of image data in a time-saving, simple and non-invasive manner [34, 35]. The challenge, however, is to make a reliable condition analysis based on this data, since an RGB modality does not directly describe a depth profile. This study demonstrates that RGB image data in combination with sophisticated machine learning methods is sufficient to provide an objective and quantitative assessment of inner surface wear.

The core of the proposed framework for nondestructive and fast on-site wear assessment is a data-driven modality transfer model. In modality transfer, observed data of a simpler modality in terms of acquisition effort is mapped to desired properties of a more complex modality. In the context of cylinder liner wear assessment, this means that the BLC, which is usually derived from a surface depth profile, is predicted from a simpler RGB modality. Although the RGB modality contains no per-pixel depth information, it offers contextual information about the surface. A neural network is trained from noisy and low-resolved RGB reflection images of the inner surfaces of gas engine cylinder liners to predict the BLC of the depth profile. To combine optical measurements with a data-driven prediction methodology, a database is built from scratch. Reflectance RGB images are acquired via a handheld digital microscope, which exemplifies the wide range of possible simple handheld optical methods. For corresponding ground truth data, engine liners are removed and 3D depth profiles of worn and unworn areas are generated using a confocal microscope. The quantile functions of the depth profile pixel values represent the desired BLCs.

Based on the database, a convolutional neural network (CNN) is developed and trained to enable prediction of the BLC of a completely different modality, namely the depth profile, from an RGB modality. Estimation errors are evaluated via probabilistic metrics (Wasserstein-1 and Wasserstein-2 distance [36, 37, 38]) and surface roughness indicators (Sk, SMr1, SMr2 [39, 40]). Once the model is trained, the BLC of a new liner can be predicted from a simple RGB image of the surface. The proposed methodology for targeted modality transfer promises a tremendous improvement in terms of time resources, replacement liner costs and measurement costs by eliminating the need to remove and destroy cylinder liners for measurement with expensive and sophisticated microscopy equipment.

1.4 Outline

The rest of this article is organized as follows. Section 2 provides background information on cylinder liners, wear assessment and replacement modalities used in this work. Section 3 presents the proposed transfer learning framework for nondestructive and accelerated wear assessment. Section 4 contains results and discussions on the proposed BLC prediction procedure. Section 5 summarizes the main findings of this study and describes important future research directions. Finally, Appendix A presents data analysis results that highlight correlations between the reflection image and the corresponding BLC.

2 Fundamentals of Cylinder Liner Surface Assessment

This section discusses the importance of the surface properties of the liner and honing structure. In addition, optical methods for measuring surface texture and surface depth are presented. The information quality of bearing load curves regarding roughness analysis and wear assessment is described in detail.

2.1 Liner Surface Machining

Surface finish is considered to be an important factor for oil consumption, piston-ring friction, wear propensity and blow-by of cylinders [41, 42, 43]. In addition, the cylinder liner surface topography also affects the engine emissions [44, 45]. The desired surface properties are achieved by a manufacturing process known as honing. In this process, the surface roughness is machined to a production standard that is usually defined by the design limits of the manufacturer. The surface is formed into a fine plateau structure that simultaneously minimizes friction and provides oil retention.



Figure 2.1: Illustration of the honing process. By adjusting the rotational velocity v_{rot} and the vertical velocity v_{vert} as well as the pressure p on the cylinder wall, the honing structure can be varied to achieve the required honing angle α and surface properties.

Honing of a cylinder surface is performed by rotating a device with abrasive stone inserts and simultaneously moving the device vertically (cf. Figure 2.1). Plateau honing is a conventional honing method used for cylinders in the gas engine industry [46]. It is usually implemented in two or three consecutive steps with either ceramic or diamond cutting stones, where the first step is pre-roughening with a removal of up to 50 μ m and the last step is a finishing process that removes about 5 μ m to obtain a finer surface [47]. Other honing techniques such as spiral-glide honing, variable-roughness honing, laser structuring, coating and form honing are also considered surface finish techniques for liner production [48].

2.2 Surface Roughness Characterization

In the field of large engine operation, evaluation of surface roughness and lubrication characteristics of engine cylinder liners can help to determine signs of failure or premature wear. The bearing load curve is an important tool for assessing surface structure conditions. As illustrated in Figure 2.2, left, the BLC plots the reduced height value against the percentage of measurement points above that height. In terms of cylinder liner functionality, the BLC describes two important parts of the honing structure: the valley part, which specifies the oil retention capacity, and the peak part, which specifies debris or asperities of the honing structure [49]. Depending on the measurement method, either profile or surface parameters are visualized by the BLC.



Figure 2.2: Left: BLC of a surface depth measurement. The corresponding surface roughness indicators Sk, Spk, Svk, SMr1 and SMr2 are derived directly from the BLC depth representation [39, 40]. Right: Example of the typical observed circular and vertical wear distribution of a cylinder liner in operation.

Profile values (denoted with an R prefix) are determined using conventional stylus profilometers [23]. These devices measure the surfaces with a needle that traverses a defined distance and outputs the corresponding BLC. Surface values (denoted with an S prefix) are determined by using optical methods. Unlike stylus instruments, optical methods are able to record the entire surface and are also less sensitive to measurement errors. Depending on whether profile or surface measurements are used, commonly extracted parameters from the BLC are:

- reduced peak heights Rpk and Spk
- heights of core roughness areas Rk and Sk
- reduced valley heights Rvk and Svk
- material ratios of peaks RMr1 and SMr1
- material ratios of valleys RMr2 and SMr2.

Additional parameters that are often considered are the arithmetical mean deviation and the maximum height of profile or surface. A definition and detailed information on parameter calculation are given in ISO 13565-2 (profiles) and ISO 25178-2 (surfaces) and further discussed in [39, 40].

2.3 Optical Metrology

In this study, optical surface imaging methods based on confocal microscopy are employed as the reference method. The basic idea is to acquire multiple 2D images at different height levels of a sample in order to reconstruct the 3D structure of the object. This is achieved by using point illuminations that pass through an optical path and pinhole in front of the detector to eliminate out of focus signals. Therefore, only light in the immediate vicinity of the focal plane is detected, and as the height level of the focal plane is varied, high depth resolution is achieved [23]. Figure 3.1 shows in the left column examples for relative depth profiles generated by a confocal microscope, where darker pixels denote valleys up to a relative depth of 8 µm and brighter values denote peaks up to a relative height of 4 µm. Other than confocal microscopy but not used for this study, white light interferometry as well as interferometric microscopy represent other potential optical reference measurement methods with high spatial and depth resolution [23].

Although surface imaging methods enable the recording of high-resolved depth maps, they are quite limited for industrial on-site use. For example, they must be isolated from external influences to avoid any degradation in their reliability caused by vibrations or thermal effects. In addition, liner components must be removed from the engine and treated with great effort (cutting, cleaning) to enable measurement with the external stationary devices. A comprehensive literature study indicates that no handheld optical alternative exists that permits use of these microscopy techniques, which essentially limits their practical utility for on-site surface measurement.

2.4 Liner Wear Assessment

As described in Section 1.2, the unavoidable wear of cylinder liners alters important surface properties. As operating time increases, wear can become extremely hazardous at certain local points and thus the honing no longer provides sufficient lubrication and fatal engine failure is highly likely to occur [50]. Besides mechanical friction, the honing structure can also wear out due to intrusion of foreign particles from either oil coking or fuel contamination [51, 52, 53]. These potential occurrences indicate the importance and benefits of condition assessment of the surface structure at regular intervals.

Comprehensive simulation work at INNIO Jenbacher GmbH & Co OG on gas engines has established that the greatest wear occurs near the top dead center (TDC). In this region, the cylinder liner wall temperature is comparatively high due to the thermal impact from combustion. In turn this leads to a low oil viscosity, a thin oil film, and thus to asperity contacts between the compressed piston ring and the cylinder wall [54]. The continuous movement of the piston causes a change in the thrust force in the cylinder and also a continuous change both in the circumferential and in the vertical directions, namely from TDC to the bottom dead center (BDC). The greatest wear occurs in the liner surface area parallel to the piston pin axis on the thrust/anti-thrust side of the cylinder, while less wear is observed in the area perpendicular to the pin axis (cf. Figure 2.2, right). Due to the peak firing pressure in the cylinder, an additional contact pressure arises between the piston ring and the liner, which further increases wear at the liner areas near the TDC.

In general, standard roughness indicators (arithmetic mean deviation and maximum height of profile) are considered insufficient for accurate liner surface wear analysis, and therefore surface height characterization of the BLC according to ISO 25178-2 is preferred. The ISO defines the exact procedure for determining BLC parameters that characterize height (peaks, core area, valleys and material ratios), which can be used for the subsequent wear analysis. This again emphasizes the benefit of an efficient BLC prediction method.



Figure 2.3: Visualization of the wear calculation: In the left plot, the worn/top BLC (orange) is vertically compensated to map the valley region of the unworn/bottom BLC (blue), where alignment of corresponding valley regions is shown in the middle plot. The right plot indicates that the wear is calculated as the difference between the aligned curves at 50% of the bearing load area.

As an integral part of the liner wear assessment, cylinder liners are cut to obtain segments with measurement regions at touched (thrust and anti-thrust) and untouched areas regarding contact with the piston and piston rings. Each segment is then measured with an optical device to obtain corresponding depth images from which the BLCs are calculated based on the empirical distribution of the depth values. Subsequently, the BLCs of worn (top, touched) and unworn (bottom, untouched) areas are adjusted in respect of height at a certain significant point or region of the curve. There are several common criteria for the height compensation of the two curves [29]. The vertical shift of the worn BLC to match the valley region of the unworn BLC as shown in Figure 2.3 is considered a reliable method for assessing wear. The prevailing assumption is that the valleys of the surfaces are not worn by the movement of the piston, thus allowing an estimation of the wear between the worn and unworn regions. An appropriate value for the vertical shift is found automatically by minimizing a distance function evaluated on both curves in the valley region; in some cases, this can also be determined visually.

After the valleys of the unworn and worn BLCs are matched, wear is quantified by taking the difference in height at the core roughness of both surfaces (cf. Figure 2.3, right). The 50% mark of the bearing load area is often considered as quantification position. In order to obtain a more reliable assessment of wear, several measurements are taken of both the worn and unworn areas of the targeted cylinder liner. Since measurement artifacts caused by material chipping or pitting in the honing structure could distort the BLC and indicate excessive wear, a more robust estimation of the wear is achieved by averaging the wear calculations on these multiple measurements.

3 Methodology

This section presents the proposed method for nondestructive and fast wear assessment of cylinder liners based on surrogate image acquisition methods and transfer learning. A modality transfer approach is provided to derive BLC information from measurement data of a surrogate modality that can be applied without full disassembly and destruction of engine components and without use of highly expensive and unportable scanners at sites of existing series engines.

3.1 Optical Surrogate Modalities

Existing wear analysis techniques require a complex workflow to generate high-resolution depth measurements. Optical reflection images are used in the modality transfer approach proposed here. The specific surrogate image modalities in this work are described below.

Fixed confocal device outputs: As visualized in Figure 3.1, the Alicona InfiniteFocus confocal microscope [55] used in this work provides the user with a depth image as well as with a high-resolved RGB reflection image of the same measurement area. This gives access to a reflection RGB image of exactly the same regions that were used to calculate the corresponding BLCs. Because the same invasive liner acquisition as the depth modality is required to determine this reflection modality, its practicality is limited. However, it is ideal as a verification concept for transferring reflection images into a BLC of the corresponding depth image.

Handheld device outputs: Of greater practical relevance is the second optical approach, in which a Mic-Fi Wi-Fi microscope [56] with a length of 135 mm captures RGB reflection images in SXGA format (1280×1024 pixels) at up to 220 times the magnification. The use of such a handheld device has the potential to enable inspection of cylinder liners without full disassembly of the component, which facilitates inner surface imaging of a permanently mounted liner. Figure 3.2 visualizes images taken with the handheld device. It highlights additional practical challenges of using this handheld device such as non-uniform light intensities, blurred areas due to cylinder curvature and the lower resolution compared to the depth profiles.

The later presented transfer learning model is not limited to this specific handheld device, which is mainly used here because this microscopic camera fits well with the cylinder dimensions. All kinds of handheld devices for RGB image acquisition such as compact cameras, smartphones, or endoscopes could be used instead.

3.2 BLC Estimation by Modality Transfer

The imaging modality underlying the BLC is a high-resolution nanoscale cylinder liner surface depth image measured in a spatial domain of approximately $1.9 \times 1.9 \text{ mm}^2$ (Figure 3.1, left). The goal is to replace these depth profile measurements with the simpler modality described above, which does not contain any information about pixel-wise depth, and to predict the BLC of the depth profile measurement from this indirect contextual RGB information. In the following paragraphs, the task is described in detail from a mathematical point of view.

Let $A \in \mathbb{R}^{d_1 \times d_2}$ denote the recorded depth image where $d_1 \times d_2$ denotes the number of sampling points (pixels). The corresponding BLC is defined as the reversed empirical quantile function $B: (0, 1) \to \mathbb{R}$ of pixel values of the depth image,

$$B(x) \triangleq \inf \left\{ y \in \mathbb{R} \mid 1 - x \le \frac{1}{d_1 \cdot d_2} \sum_A \mathscr{W}\{a_i \le y\} \right\}.$$
(3.1)



Figure 3.1: Left: Examples for optical depth measurements of worn and unworn $1.9 \times 1.9 \text{ mm}^2$ surfaces (top and bottom, respectively), obtained with an Alicona InfiniteFocus confocal microscope [55]. Right: Corresponding optical reflection measurements (RGB images), simultaneously generated by the same microscope.

Figure 3.2: Examples of RGB measurements with a handheld optical device, namely a Mic-Fi Wi-Fi microscope [56], recorded with $80 \times$ magnification, of worn positions of distinct liners. Every measurement covers a region of $5.2 \times 4.2 \text{ mm}^2$.

Here the sum runs over all pixels a_i in depth image A and $\mathbb{V}\{z \leq y\}$ denotes the indicator function, returning 1 if $z \leq y$ and 0 otherwise. In fact, the discretized BLC $b \in \mathbb{R}^K$ is considered, obtained after sampling the left-continuous BLC B in (3.1) at K equidistant sampling positions. This allows the entire curve to be represented in a one-dimensional vector of size K in subsequent implementations.

The task consists of predicting BLC *b* derived from the depth image by processing an RGB reflection image $I \in [0, 255]^{M_1 \times M_2 \times 3}$ of the same object. Here $M_1 \times M_2$ is the number of sampling points and 3 the number of color channels. In mathematical terms, the task is to find a transfer function

$$\mathcal{N} \colon [0, 255]^{M_1 \times M_2 \times 3} \to \mathbb{R}^K, \ I \mapsto b \,, \tag{3.2}$$

that maps the RGB reflection image to the corresponding BLC. The analysis in Appendix A suggests that RGB reflection images of a cylinder liner surface do indeed provide sufficient information to distinguish between light and severe wear. This supports the assumption that RGB reflection images permit the determination of the corresponding BLC of the depth image. It is worth noting that the prediction of the empirical quantile function of a different modality can be seen as a novel variant of modality transfer [13, 57, 58, 59, 60, 61], in which low-dimensional information of a desired measurement modality is synthesized from a simpler data acquisition technique.

3.3 Machine Learning Approach

A supervised machine learning approach is proposed to construct a transfer function (3.2) for predicting BLCs from RGB images. It is composed of the following components. First, a database is created that contains training data pairs (I_n, b_n) of RGB input images $I_n \in [0, 255]^{M_1 \times M_2 \times 3}$ and corresponding BLC outputs $b_n \in \mathbb{R}^K$ for $n = 1, \ldots, N$, where $N \in \mathbb{N}$ denotes the number of training data pairs. Second, a class of networks $\mathcal{N}_{\theta} \colon [0, 255]^{M_1 \times M_2 \times 3} \to \mathbb{R}^K$ is designed, where θ denotes the model parameters and $\theta \mapsto \mathcal{N}_{\theta}$ is referred to as the network architecture. Third, a loss function $\ell \colon \mathbb{R}^K \times \mathbb{R}^K \to [0, \infty]$ is constructed that measures the deviations between the model output $\mathcal{N}_{\theta}(I)$ and the ground truth BLC b. Finally, mathematical optimization algorithms are used to minimize the empirical risk function

$$\mathcal{R}(\theta) \triangleq \arg\min_{\theta} \frac{1}{N} \sum_{n=1}^{N} \ell \left(\mathcal{N}_{\theta}(I_n), b_n \right).$$
(3.3)

The final trained neural model is given by $\mathcal{N} = \mathcal{N}_{\hat{\theta}}$, where $\hat{\theta}$ is determined to be close to a minimizer of the empirical risk function (3.3). During inference, prediction of the BLC from an RGB reflection image will simply involve applying the trained neural network.

A variety of efficient optimization algorithms exist for minimizing the risk functional (3.3). For example, it can be solved with a wide range of stochastic gradient descent implementations [62, 63, 64]. In this work, (3.3) is minimized using the Adam algorithm (stochastic gradient descent with adaptive learning rates). Adam is an iterative minimization algorithm that updates the model parameters by reusing all the training data pairs in a cyclic manner. In machine learning, the number of training cycles is commonly referred to as the number of epochs. Important task-specific components of the machine learning framework are the generation of the underlying database as well as the design of the network and the loss function. These elements are described in detail in the following subsections.

3.4 Database Generation

In the scope of this research project, a database was created from scratch. The examined objects are ten type 6 gas engines from INNIO Jenbacher GmbH & Co OG. Each engine consists of 12, 16, 20 or 24 cylinders with a displacement of approximately 6 dm³. The number of operating hours of the corresponding cylinder liners varies between 2550 h and 60 000 h, where this number may vary within an engine due to past replacements of individual liners. The inspection of each liner goes through a time-consuming and resource-intensive logistics chain, which initially consists of disassembling the liner, removing it from the engine and marking it. A segment of 45° is cut out of each liner which contains positions parallel to the piston pin axis (06-h-segment) as described in Figure 3.3. Since these areas are permanently in contact with the piston and the piston rings, this is where the greatest wear is expected. These segments are superficially cleaned and the confocal microscope measures three to five different areas within the TDC (upper reversal point; see Figure 3.3). The results for each measured area are an RGB image and a depth profile with an associated BLC.

After the depth images are generated, the segments are forwarded to a second distinct measurement station, where approximately the same areas of the TDC are recorded with the handheld device using a self-constructed segment holder (Figure 3.4, left) placed in a darkened room. Due to very fine scaling on micrometer range, it is not possible to measure the exact same



Figure 3.3: Left: A segment of 45° is cut out of each liner which contains the area parallel to the piston pin axis where the greatest wear is expected (in this case the 6 h position). Another segment is cut out at the 3 h position, which is the liner area perpendicular to the piston pin axis and used for later comparisons to unworn areas. Right: Comparison of worn and unworn areas. The grayscale on the topographic scans represents the relative depth, i.e. the distance of the surface from the normalized surface height. The scans of the worn areas show less contrast due to the abrasions of the peaks. The greatest wear occurs on the wall of the 6 h position near the top reversal point. The bottom dead center shows marginal surface abrasions on the wall at the 3 h position.

spatial range of the segment surface as the depth profile with the handheld device. The segment holding device ensures that the areas of the confocal microscope measurements are embedded in the corresponding handheld device records (Figure 3.4, right). Therefore, the whole area from which the empirical quantile function is calculated is also visible in the handheld device observation. A total amount of 73 liner segments and 3 to 5 measurement areas per segment yield a total of 271 modality pairs.



Figure 3.4: Left: Measurement station with a self-constructed segment holder. The holding device enables measurement of comparable positions relative to the confocal microscope. The camera transfers the recorded measurement via a Wi-Fi connection to mobile work stations. Right: Depth profiles from the fixed confocal device for BLC calculation are fully contained in the corresponding handheld microscope measurements.

3.5 Network Architecture

In this work, deployment of convolutional neural networks is proposed for BLC prediction. CNNs represent the state of the art for several image processing tasks such as segmentation [65, 66, 67], image reconstruction [68, 69] or modality transfer [13, 57, 58, 59, 60, 61]. The main building block of these architectures consists of convolutional layers, where each channel of the input is convolved with several data adapted kernels (Figure 3.5). For each kernel, convolution with the layered input followed by application of a nonlinear activation function results in hidden feature maps representing properties of the input channels on different scales. To train a CNN, the most informative feature maps are found by fitting convolution kernels to the given data. The overwhelming success of CNNs has been the hallmark of the last decade of machine learning in image processing. The parameters θ used in a CNN are mainly composed of millions of convolution kernel values, which in this work will be adjusted to obtain BLCs from RGB input.

A variation of a well-known CNN architecture called VGG Net [70] is used in this study. In the last decade of machine learning, VGG Net architectures have redefined the state of the art in image classification. The VGG Net enables deep multi-scale feature learning via convolutional blocks with 3×3 kernels, which has been successfully employed in several public classification challenges to date. Its high generalization properties and superior ability to capture image distribution also suggested the use of the VGG network architecture beyond classification, for example in segmentation [66, 67] or tomographic image reconstruction [68, 71]. The key idea is to prune the network in the latent space (encoder output) and to transform the low-scale feature representations into a vector of length K describing the discretized BLC (cf. Figure 3.6). The VGG13 network architecture [70] is used, where the univariate output layer is replaced by a dense layer of size K = 256 representing the discretized BLC prediction. The RGB images are subsampled by a factor of 2 and propagated through the network, which consists of a total of 19.4×10^6 adjustable parameters.



Figure 3.5: Convolution of a three-layered input with two filters. Both filters slide vertically and horizontally over the given input layers and convolve each input channel with individual 3×3 kernels. Summation of the resulting feature maps of filter 1 and filter 2 yields the first and second output layer, respectively.



Figure 3.6: Visualization of the proposed variation of the VGG13 Net [70] for BLC prediction. Every convolution and dense layer is followed by a ReLU activation function [72].

3.6 Loss Function Design

In addition to the model architecture and the quality of training data, the choice of the loss function also has a major impact on the model's performance. In general, it is desirable to find a loss function that matches the targeted task best while minimizing a task-dependent distance measure. Mathematically, a BLC is an empirical quantile function, which in turn is the inverse function of the empirical cumulative distribution function. Therefore, comparing a predicted BLC with the corresponding ground truth is equivalent to comparing two probability distributions.

A commonly considered metric for measuring the closeness of two probability distributions is the Wasserstein-1 distance [36, 37, 38], which is also called the earth-mover distance. The Wasserstein-1 distance between two probability distributions P_1 and P_2 is defined by

$$\mathcal{W}_1(P_1, P_2) \triangleq \inf_{J \in \mathcal{J}(P_1, P_2)} \mathbb{E}_{(x, y) \sim J} \|x - y\|,$$

where the infimum is taken over the set of all joint probability distributions that have marginal distributions P_1 and P_2 . The Wasserstein-1 distance can be interpreted in the setting of optimal mass transport. In this setting, one aims to find an optimal transfer plan to transport one mass distribution into another as cheaply as possible in reference to a given cost function [73]. Taking a metric as cost function results into the Wasserstein-1 distance. It is important to note that, contrary to the standard L^p -norms, the Wasserstein distance not only compares distribution values point-wise but also quantifies how far the distributions have to be moved. Wasserstein distances have several useful properties and dual representations, which makes iterative solution of the transport problem computationally feasible [36, 37]. Moreover, given two one-dimensional probability distributions P_1 and P_2 , the Wasserstein-1 distance simplifies to $W_1(P_1, P_2) = \int_0^1 |Q_1(z) - Q_2(z)| dz$, where Q_1, Q_2 denote the corresponding quantile functions.

Because BLCs are the empirical quantile functions associated with depth profiles, the Wasserstein-1 distance between ground truth depth distribution and predicted depth distribution corresponds to the component-wise L^1 -distance between the corresponding discretized BLCs. A significant property of a BLC, which is also important for subsequent depth valley alignment (Section 2.4), is its curvature at specific positions. Therefore, to ensure that the model prediction \hat{b} and the corresponding ground truth b have nearly the same curvature, the Wasserstein-1 distance after application of the second order central difference quotient is also considered during loss minimization. This results in the proposed penalized loss function

$$\ell_{1}^{\lambda}(\hat{b}, b) \triangleq \frac{1}{K} \sum_{k=1}^{K} \left| \hat{b}_{k} - b_{k} \right| + \lambda \frac{1}{K} \sum_{k=1}^{K} \left| \partial^{2} \hat{b}_{k} - \partial^{2} b_{k} \right|,$$
(3.4)

where ∂^2 denotes the second order central difference quotient and $\lambda \geq 0$ is a regularization parameter indicating the penalty during training when predicting curves with dissimilar curvature. The final transfer learning model minimizes the empirical risk (3.3) using the penalized Wasserstein loss (3.4), training data (I_n, b_n) described in Section 3.4 and the architecture $(\mathcal{N}_{\theta})_{\theta \in \Theta}$ described in Section 3.5.

4 Results and Discussion

This section presents the results for supervised BLC prediction from RGB reflection images using data pairs from the stationary confocal device as well as from the handheld device. To assess the similarity between BLC prediction \hat{b} and ground truth b, the Wasserstein-p distance $W_p(\hat{b}, b) \triangleq (1/K \sum_{k=1}^K |\hat{b}_k - b_k|^p)^{1/p}$ for $p \in \{1, 2\}$ and roughness metrics Sk, SMr1 and SMr2 are evaluated for both devices. Sk core roughness is measured on µm scale, and peak and valley material ratios SMr1 and SMr2 are given as a percentage. Computation of considered roughness metrics is quite demanding and requires separate optimization of a linear regression model for each BLC. More detailed information can be found in [39, 40].

4.1 Results with the Fixed Device

The first reported results in Table 4.1 use as inputs the fixed device reflection RGB images taken by the confocal microscope simultaneously with the depth profiles. As stated in Section 3.4, the database consists of 271 reflection images of the fixed device and the corresponding depth images. To increase the number of training data, nine non-overlapping patches of size 512×512 pixels are extracted from each reflection image and the BLC of the corresponding depth patch is generated, where the length of the discretized BLC equals K = 256. This yields $271 \cdot 9 = 2439$ available training pairs.

The evaluation is carried out by means of a three-fold cross-validation. More specifically, the dataset is split into three separate parts of nearly identical size. During training, one part (evaluation data) is omitted and then used for testing on unseen data while the model is optimized on the other two parts (training data). The process is repeated three times, each time omitting another third. Note that there is always a strict separation between the liners. This means that no liners contribute to the training data and the evaluation data at the same time. This is employed to mimic use in a practical scenario where the network would also be applied to liners that the underlying model has never seen during training. While splitting it is ensured that all folds have a similar liner operating hour distribution.

The risk functional (3.3) is minimized with the Adam algorithm using 40 epochs and an initial learning rate of 5×10^{-5} , as this optimizer setting provided a quite stable training progress in most of the experiments. Training stability is enforced by a learning rate scheduler that constantly decreases the learning rate. The results of three-fold cross-validation with $(\lambda > 0)$ and without $(\lambda = 0)$ curvature penalty are shown in Table 4.1. Due to considered three-fold cross-validation, it is possible to predict a BLC for every sample using a model that has never seen this sample during training. Running this experiment several times with different values for the curvature penalty parameter $\lambda \in \{k \times 0.05 \mid k = 0, ..., 10\}$ has shown that $\lambda = 0.4$ achieves the best improvement in terms of the evaluation metrics. It can be observed that if the penalty term is employed, there is a significant performance gain in terms of the unseen liner evaluation metrics for the first two folds. Considering that the quantile function of an unknown depth profile was predicted from an RGB reflection image, the results in Table 4.1 are quite satisfactory.

Table 4.1: Evaluation results for the fixed device RGB data observed during three-fold cross-validation (3 training runs). For all metrics, the mean absolute errors between the predictions on unseen data (of the excluded fold) and the ground truths are reported (smaller is better).

Run	\mathcal{W}_1	$\mid \mathcal{W}_2$	Sk μm	SMr1	SMr2				
Without curvature penalization $(\lambda = 0.0)$									
1	0.103	0.222	0.182	1.229	2.236				
2	0.106	0.223	0.192	1.501	3.003				
3	0.131	0.270	0.224	1.360	2.682				
Avg.	0.113	0.238	0.199	1.363	2.640				
With curvature penalization $(\lambda = 0.4)$									
1	0.095	0.187	0.157	1.537	2.557				
2	0.094	0.212	0.147	1.494	2.880				
3	0.144	0.308	0.242	1.405	2.958				
Avg.	0.111	0.237	0.182	1.478	2.798				

Table 4.2: Evaluation results for the handheld device RGB data observed during three-fold cross-validation (3 training runs). For all metrics, the mean absolute errors between the predictions on unseen data (of the excluded fold) and the ground truths are reported (smaller is better).

Run	\mathcal{W}_1	\mathcal{W}_2	Sk μm	$SMr1 \mid SMr2$;			
Without curvature penalization $(\lambda = 0.0)$								
1	0.104	0.219	0.164	$1.070 \mid 2.374$				
2	0.102	0.231	0.156	$1.256 \mid 2.497$				
3	0.110	0.237	0.179	1.172 2.876				
Avg.	0.105	0.229	0.166	$1.166 \mid 2.582$				
With curvature penalization $(\lambda = 0.1)$								
1	0.103	0.230	0.154	1.026 2.232				
2	0.099	0.236	0.155	1.120 2.570				
3	0.106	0.231	0.178	1.265 2.836				
Avg.	0.103	0.232	0.162	$1.137 \mid 2.546$				

4.2 Results with the Handheld Device

The input RGB reflection images cover a spatial area of $4.2 \times 4.2 \,\mathrm{mm}^2$ with a total of 1024×1024 pixels and are subsampled by factor 2 (i.e. taking every second pixel) to a total of 512×512 pixels. The corresponding depth image covers an area of about $1.9 \times 1.9 \,\mathrm{mm}^2$ and is fully contained within the range of the reflection image (Figure 3.4, right). The RGB images show a significant degradation in image quality compared to the fixed device inputs in terms of inconsistent light intensities and blurred areas due to cylinder curvature (Figure 3.2). Since the resolution of the input images is comparatively small and there is no one-to-one correspondence to the given depth images regarding the exact position, it is not feasible to use patch extractions for data augmentation as in the previous experiment. Instead, the set of training data is augmented to $271 \cdot 2 = 542$ by vertically flipping the reflection image while the corresponding BLC remains unchanged. Training is performed within only 25 epochs using Dropout regularization to prevent overfitting [74]. By numerous experiments, the value 3×10^{-4} turned out to be the best learning rate in terms of training stability. Furthermore, the ReLU activation function initially proposed in Figure 3.6 is replaced by the hyperbolic tangent activation function, which yields better results for the handheld data.

The results of three-fold cross-validation without $(\lambda = 0)$ and with $(\lambda > 0)$ curvature penalization are shown in Table 4.2. The best performing CNN is obtained for $\lambda = 0.1$ (cf. Figures 4.1 and 4.2), where again the set $\{k \times 0.05 \mid k = 0, ..., 10\}$ is used for regularization hyperparameter search. For this model, the mean absolute difference between the ground truth Sk values and the model predictions is 0.162 µm, which is only 17.5% of the average ground truth Sk value 0.925 µm. For the SMr1 values, which denote the material ratios of the peak regions, the observed mean absolute difference is 1.14 percent points, and for the SMr2 values, which denote the material ratios of the valley regions, the observed mean absolute difference is 2.55 percent points. Note that the evaluation results of the handheld device are even better than the results in Section 4.1. This is remarkable, since images acquired with the handheld device create additional practical challenges such as different light intensities and artifacts (cf. Figures 3.1 and 3.2). However, this observation could be explained by the fact that the handheld device data covers a broader measurement range than the fixed device data.



Figure 4.1: The subplots show exemplary BLC predictions using penalized Wasserstein loss for samples that have not been used during model training. Each subplot contains the ground truth (blue dashed line), BLC prediction of the model (orange solid line) and Wasserstein-1 and Wasserstein-2 metrics as well as the Sk value of the model output and the ground truth, respectively.



Figure 4.2: Quantitative evaluation for the proposed model is conducted on a fine discretization of the bearing load area. Each interval on the horizontal axis covers 5% of the bearing load area and the corresponding boxplot diagrams represent the mean absolute error between model predictions and ground truths in this area on µm scale. The predictions are always generated on unseen samples via three-fold cross-validation, i.e. for each prediction a model is used where corresponding fold has been excluded during training.

While Figure 4.1 visualizes exemplary BLC predictions on unseen data of the proposed model, Figure 4.2 provides for all test folds a quantitative evaluation based on the bearing load position. Prediction errors are quite satisfying in the core roughness area and slightly increase near the peak area (interval [0%, 5%)). Large prediction errors can be found in the valley region (interval [85%, 100%]) as these errors are mainly caused by high variations of the training BLC values in this load area. In addition, dependency between model performance and the amount of accumulated operating hours of the tested liner is investigated in Figure 4.4. Although no significant correlation between the boxplot diagrams of the targeted metrics (Wasserstein-1 distance and Sk errors) and the lifetime of the examined liner can be observed, it is obvious that for both metrics the distance of the outliers (small circles) to the box itself (and the whiskers) increases with the amount of accumulated operating hours.

The BLC predictions shown in Figures 4.1 and 4.2 demonstrate a reasonable approximation of the true BLC, which represents the underlying depth profile of the surface. Since the collection of paired data for supervised training is a very time-consuming and laborious process in this application, the size of the proposed data set is rather modest. Nevertheless, it was possible to generate modality pairs of a representative quantity of cylinders with a wide variety of operating hours, and to build a novel machine learning model on this database for reliable depth representation learning. Unwanted consequences caused by small data sets in machine learning like overfitting or unrepresentative training data are almost entirely avoided by use of regularization techniques, which successfully led to visual satisfying BLC predictions on unseen evaluation data. The predicted BLCs provide quite satisfactory core roughness approximations but unreliable material fraction indicators, which are very sensitive to curve values in the peak and valley regions (Figure 2.2). This is further analyzed in Figure 4.3, where roughness indicators extracted from BLC predictions are plotted against their corresponding ground truths. However, this does not come as a surprising result, as all models have been optimized for accurate BLCs via minimization of the (penalized) Wasserstein-1 distance but are not optimized regarding quantification of roughness indicators. Future work will aim to tailor predictions to roughness indicators and actual wear values.



Figure 4.3: Roughness indicators Sk, SMr1 and SMr2 are extracted from the BLC predictions of each fold and plotted against their ground truths. Perfect prediction of these indicators corresponds to points lying on the diagonal in each case.



Figure 4.4: The dependency between the accumulated operating hours of a liner and the performance of the model on the unseen folds is visualized in a boxplot diagram. The performance of the model is assessed via the Wasserstein-1 distance (left) and the Sk errors (right) between model predictions and corresponding ground truths.

5 Conclusion

Current engine liner evaluation protocols use sophisticated depth measurements that require disassembly and destruction of the engine components under investigation. This research proposes the use of simpler acquired reflection images in combination with CNNs to predict the BLC of a high-resolution depth measurement. For this purpose, a multiscale architecture is proposed together with a penalized version of the Wasserstein-1 distance as a suitable loss function for BLC prediction. The visual quality of the predicted BLC confirms the plausibility of using surrogate modalities in combination with advanced data-driven models. Furthermore, the current models are able to give coarse approximations of the core roughness parameters (Sk) but are not yet able to provide reliable estimates for the material ratio indicators (SMr1, SMr2). In addition, several regularization experiments were required to stabilize model training and thus compensate for the rather small amount of available training data. There are several directions in which to improve model performance in future research. First, more modality pairs will be acquired to further strengthen generalizability and performance of the proposed method. While the goal of this research project to date has been to predict a one-dimensional representation of the corresponding depth modality, the aim for the future is to generate reliable depth profiles from RGB images without paired data. Finally, the loss function will be updated by differentiable implementations of these quantitative measures in order to obtain reliable roughness information and wear values from predicted curves.

A Correlation Analysis between Reflection Images and Wear

The following section investigates correlations and patterns in the data set between the RGB images and the observed localization (TDC/BDC), highlighting the potential of transfer prediction. To this end, three well explored and commonly used approaches in unsupervised learning are used: principal component analysis [75, 76], k-means clustering [72, 77] and Gaussian mixture clustering [72]. Before the results of targeted analysis are presented, these methods are discussed briefly.

A.1 Methods

Principal component analysis (PCA) is a powerful technique that beneficially changes the basis on which the data is represented. In the first step, PCA determines the best fitting line with respect to the data, i.e. it finds a line in the high-dimensional data space that minimizes averaged squared distances to all data points. As a result, the orthogonal projections of the data onto this line explain most of the variance in the original data space and thus the line is referred to as the first principal component. In the next step, the PCA algorithm searches for another best fit line that is perpendicular to the first line. The second principal component is thus the line perpendicular to the first principal component that explains most of the remaining variance. These steps are then continued following the same procedure until the desired number of principal components has been reached. Since the first few principal components explain more of the variance in the original data than the later ones in most applications, calculation of the first few principal components is often sufficient to obtain reasonable representations in lower dimensional space.

In essence, k-means clustering is a cost minimization algorithm. The cost function is defined over the parameterized set of all possible cluster assignments of data samples x_1, \ldots, x_n . The cost function for k-means algorithm is defined as $\sum_{i=1}^k \sum_{x \in C_i} d(x, c_i)$, where C_i denotes the cluster with centroid c_i for $i = 1, \ldots, k$ and d a distance measure, e.g. the Euclidean distance. The number of considered clusters k is fixed to fit the task. Since the aim is to minimize the above cost function, the k-means clustering algorithm has been proposed. This algorithm consists of application of the following steps:

- 1. Initial centroids c_1, \ldots, c_k are randomly generated within the data domain.
- 2. Clusters C_1, \ldots, C_k are created by associating every data point with its nearest centroid with respect to a predefined distance measure, usually the Euclidean distance.
- 3. The new centroid of each cluster C_i is calculated as the mid-point of the data samples within C_i .
- 4. Steps 2 and 3 are repeated until a convergence criterion is reached.

It has been shown that k-means clustering does not increase the value of the cost function during iteration [72]. In exceptional cases, it is possible that the cost value remains constant. Nevertheless, the k-means algorithm is still widely used as it is easy to implement and performs nicely for a wide range of problems [77].

A Gaussian mixture (GM) is a function that consists of a superposition of several Gaussian kernels, each defined by a (multivariate) center and a corresponding covari-

ance matrix. In the one-dimensional case, the Gaussian kernel takes the form $g(x) = (2\pi\sigma^2)^{-1/2} \exp\left(-(x-\mu)^2/(2\sigma^2)\right)$, where μ denotes the center and σ^2 the variance. In comparison to the k-means algorithm, the GM algorithm does not conduct a strict cluster assignment. Instead, a soft assignment for each data sample is determined, i.e. a probability is calculated that the sample belongs to a certain cluster. Final labeling is achieved by assigning each data sample to the cluster with the highest probability. It has been shown that the GM algorithm converges to a local maximum of the likelihood function [78]. The GM algorithm is given by the following iterative process:

- 1. Initial Gaussian centers μ_1, \ldots, μ_k and corresponding covariance matrices are randomly generated.
- 2. For each data sample x_n , the probability $P(x_n \mid i)$ that the sample belongs to the *i*-th cluster is calculated.
- 3. The posterior probability $P(i \mid x_n)$ for the *i*-th cluster and data sample x_n is determined deploying Bayes' formula [79].
- 4. The posteriors are used to update the corresponding means and covariance matrices.
- 5. Steps 2, 3 and 4 are repeated until a convergence criterion is reached.

A.2 Unsupervised Analysis



Figure A.1: Exemplary patches taken from TDC (left) and BDC (right) images and corresponding grayscale histograms (below).



Figure A.2: Flowchart for correlation analysis.

The goal of the following liner surface analysis is to detect patterns in the data set. In this approach, input consists of patches of RGB records from the confocal microscope generated in parallel to the depth maps (Figure 3.1, right). These images have a high resolution ($\sim 3160 \times 3160$ pixels) and represent a nearly square area with a spatial range of 1.9×1.9 mm². These images are used for the data analysis and prepared in a preprocessing pipeline to save computational resources and to ensure data set enhancement. The target is to investigate if RGB images are sufficient for differentiating between worn and unworn surface measurements.

During preprocessing, the RGB images are converted to grayscale and divided into smaller subsampled patches of 150×150 pixels. The grayscaled patches and corresponding histogram representations (Figure A.1) undergo analysis. Evidently, the TDC (worn) histogram is more centered, which is due to the fact that as the amount of operating hours increases, the peaks of the honing structure steadily abrade. Therefore, peak regions of patches taken near the bottom dead center are more distinct. The PCA algorithm projects patches and histograms into a lower-dimensional space, where GM clustering and k-means clustering methods are employed to investigate how the TDC (worn) and BDC (unworn) measurements can be separated. Both clustering algorithms are not deployed on original data but on the first two principal components. The entire preprocessing pipeline and calculation steps are illustrated in Figure A.2.

A.3 Correlations



Figure A.3: Visualization of image patches and histograms projected to two-dimensional space. The axes represent the first two principal components.

Clus	•	1 T		1 B		2 T		2 B		3 T	3 B
Image patches											
1		86		334		0		293		394	50
2		314		66		400		107		6	350
Histogram representations											
1		399		8		400		96		399	59
2		1		392		0		304		1	341

Table A.1: Correlations between clustering and position of measurements (T for TDC, B for BDC) for liners 1, 2 and 3.

This subsection presents the results of the correlation analysis between RGB images and corresponding measurement positions. There are two basic measurement locations for each liner, one near the TDC and one near the BDC. While the area near the TDC should show signs of wear, the areas at the BDC position were expected to be minimally affected by wear. The data is taken from a INNIO Jenbacher GmbH & Co OG gas engine, where altogether three liners with a displacement of approximately 6 dm^3 and about 7000 operating hours each have been considered. In the following, the three liners are referred to as liner 1, liner 2 and liner 3. Image patches and histograms are projected into a two-dimensional space via PCA (Figure A.3). Qualitatively inspected, Figure A.3 suggests that histogram representations of TDC and BDC can be more efficiently separated by a well fitted parabola.

Due to the different structures of spatial image patches and histogram representations, the k-means clustering algorithm and the GM algorithm are chosen to detect clusters for image and histogram data, respectively. GM is computationally more expensive than k-means clustering and in this study both perform similarly on spatial images. Histograms provide a coarse approximation of the underlying pixel distribution of the preprocessed image data. In order to reduce the numerical complexity, the whole range of values (here rescaled to the interval [0,1]) is divided into a non-overlapping series of bins, where the height of each bin represents the amount of pixel values within the given range of the bin. The number of bins is taken as 40, since a higher amount of bins (and therefore a better approximation to pixel distribution) did not significantly change correlations between clusters and measurement positions.

The clustering results are shown in Table A.1 and indicate correlations between the amount of accumulated wear and the unsupervised detected separation into clusters of the RGB modality. A significantly better correlation is detected for liner 1. If optical patches of liner 1 are used, two clusters of data can be distinguished where 152 patches do not correlate with their measurement position. In contrast, use of the histogram representation combined with the GM algorithm yields a cluster assignment where only nine samples do not correlate with the corresponding measurement position. This successfully concludes the unsupervised analysis, suggesting that a RGB modality contains hidden structure features which are automatically detected to separate surface measurement into corresponding measured positions TDC (worn) and BDC (unworn).

A.4 Classification

The sufficient capability of the RGB modality to separate between worn and unworn surfaces in an unsupervised manner enhances the idea to develop a classifier on this issue. A supervised approach targets the functional relationship between input and corresponding output. For targeted classification, the input consists of RGB records from the confocal microscope, and corresponding output is a label denoting the measurement position. To be more exact, the classifier takes as input RGB patches $I \in [0, 255]^{M_1 \times M_2 \times 3}$, where $M_1 \times M_2$ is the number of sampling points (pixels), and estimates the corresponding measurement position, chosen from the label set {BDC, TDC}.

The afore-mentioned classifier is taken as a convolutional neural network, a basic introduction to which is given in Section 3.5. In general, there are many known architectures for classification tasks to choose from, for example VGG [70], ResNet [80] and Xception [81] architectures. These have all been evaluated on the ImageNet database [82, 83], a large-scale benchmark database for image recognition consisting of over 1.4×10^7 images. All these architectures share the property of small 3×3 convolution kernels, which results in fast computation of the hidden feature maps (cf. Figure 3.5). This study has used the VGG16 network [70] since it is the most shallow one and yields promising results on the targeted classification task.

As thoroughly discussed in Section 3.6, the choice of the loss functional is critical to the optimization of neural networks. Since this problem consists of two classes (BDC = 0, TDC = 1), the binary cross-entropy error [72] is an appropriate choice:

$$\ell_{\rm CE}(t,\hat{t}) = -t\log(\hat{t}) - (1-t)\log(1-\hat{t}), \tag{A.1}$$

where $t \in \{0,1\}$ denotes the true label (BDC or TDC), and $\hat{t} \in (0,1)$ is the probability of belonging to label 1 (TDC) estimated by the network. Since the two classes (BDC and TDC) are balanced, the actual performance of the classifier can be assessed by the accuracy metric. For ground truth labels (t_1, \ldots, t_N) and corresponding rounded network prediction probabilities $(\hat{t}_1, \ldots, \hat{t}_N)$, the accuracy is given by $\#\{\hat{t}_n = t_n \mid n = 1, \ldots, N\}/N$, where #denotes the size of the set. Training pairs (RGB patches and corresponding position labels) are divided into a training set, on which the chosen network adapts its convolution kernels, and a validation set, which is used to evaluate classification accuracy after each epoch. Here the training set is chosen to contain 70% of all data pairs, which leaves 30% for the validation set. The split is conducted randomly over all available patches. The sampling dimensions M_1, M_2 for input RGB patches equal 256.



Figure A.4: Results of the supervised classification on RGB patches and corresponding measurement positions.

The Adam optimizer [64] with an initial learning rate of 1×10^{-3} is used to optimize model weights via a stochastic gradient approach. Furthermore, to speed up minimization of the considered loss functional, model weights are not initialized randomly but taken from an equivalent VGG16 network which has already been pretrained on the ImageNet database [82, 83]. As a consequence, important features such as edges can already be distinguished at the beginning of the training process, which yields higher classification accuracy in the first epochs and subsequently faster convergence. The results after 10 epochs of training can be examined in Figure A.4. The graph on the left visualizes convergence of the loss functional in Equation (A.1) - while the graph on the right describes the corresponding classification accuracy. Obviously, the network approach achieves nearly perfect accuracy on the validation set. The smaller loss of the validation set is caused by regularization techniques used on the training set, which assists the model to prevent overfitting [72].

Funding

The authors would like to acknowledge the financial support of the "COMET - Competence Centres for Excellent Technologies" Programme of the Austrian Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK) and the Federal Ministry for Digital and Economic Affairs (BMDW) and the Provinces of Styria, Tyrol and Vienna for the COMET Centre (K1) LEC EvoLET. The COMET Programme is managed by the Austrian Research Promotion Agency (FFG) [grant number 865843].

References

[1] AVL List GmbH, AVL EPOS[™] (2020).
 URL https://www.avl.com/web/guest/-/avl-epo-1

- [2] C. Cartalemi, M. Meier, G. Sudwoj, P. Theodossopoulos, E. Tzanos, I. Karakas, A real time comprehensive analysis of the main engine and ship data for creating value to ship operators, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [3] Y. Kang, J. Kim, S. Ryu, H. Yoon, Y. Cheong, HiEMS[®], development of engine management solution based on IoT technology, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [4] T. Moeller, H. R. Olesen, S. Krogsgaard, Digitalization & IoT technologies drives development of large two-stroke marine diesel engines, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [5] C. Parikh, B. Lander, K. Lierz, C. Lau, GE transportation's remote monitoring of locomotive performance to increase availability, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [6] W. Östreicher, A. Doncic, S. Goranov, R. Hunger, A. Pennycott, C. Sidler, M. Signorile, WiCE – new engine control system for two-stroke engines, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [7] S. Venkataramanan, A. Vlietstra, R. Boom, LECM as the next-generation engine control platform for changing times in large engine market, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [8] R. Aufischer, S. Schallmeiner, A. Wimmer, M. Engelmayer, Intelligent bearings to support engine development, MTZ worldwide 80 (12) (2019) 36–41.
- [9] M. Coppo, F. Catucci, M. Ferro, M. Longhitano, Fuel injection 4.0: the intelligent injector and data analytics by OMT enable performance drift compensation and condition-based maintenance, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [10] M. Daiber, T. Spilker, Embedded turbocharger performance monitoring, in: 28th CIMAC world congress on combustion engines, 2016.
- [11] R. Teichmann, M. Abart, H. Mohr, K. Xylogiannopoulos, A. Przymusinski, R. Strasser, K. Lee, The future of condition monitoring of large engines – towards digitalization, big data tools, cloud intelligence and digital twins, in: 29th CIMAC world congress on internal combustion engines, 2019.
- [12] O. C. Basurko, Z. Uriondo, Condition-based maintenance for medium speed diesel engines used in vessels in operation, Applied thermal engineering 80 (2015) 404 – 412.
- [13] T. P. Carvalho, F. A. Soares, R. Vita, R. d. P. Francisco, J. P. Basto, S. G. Alcalá, A systematic literature review of machine learning methods applied to predictive maintenance, Computers & industrial engineering 137, 106024 (2019).
- [14] W. Zhang, D. Yang, H. Wang, Data-driven methods for predictive maintenance of industrial equipment: a survey, IEEE systems journal 13 (3) (2019) 2213–2227.
- [15] Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, A. Nandi, Applications of machine learning to machine fault diagnosis: a review and roadmap, Mechanical systems and signal processing 138, 106587 (01 2020).

- [16] W. Cao, G. Dong, Y.-B. Xie, Z. Peng, Prediction of wear trend of engines via on-line wear debris monitoring, Tribology international 120 (2018) 510–519.
- [17] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT Press, 2016.
- [18] A. Tran, J. M. Furlan, K. V. Pagalthivarthi, R. J. Visintainer, T. Wildey, Y. Wang, WearGP: a computationally efficient machine learning framework for local erosive wear predictions via nodal gaussian processes, Wear 422 (2019) 9–26.
- [19] S. Korcek, R. K. Jensen, M. D. Johnson, J. Sorab, Fuel efficient engine oils, additive interactions, boundary friction, and wear, in: Tribology series, Vol. 36, Elsevier, 1999, pp. 13–24.
- [20] R. A. Patel, K. H. Niralgikar, A review of wear in piston ring of internal combustion engine, International journal of engineering research & technology 4 (2015) 972–977.
- [21] H. Rahnejat, Tribology and dynamics of engine and powertrain: fundamentals, applications and future trends, Elsevier, 2010.
- [22] Z. Dimkovski, C. Anderberg, R. Ohlsson, B.-G. Rosén, Characterisation of worn cylinder liner surfaces by segmentation of honing and wear scratches, Wear 271 (3-4) (2011) 548– 552.
- [23] T. G. Mathia, P. Pawlus, M. Wieczorowski, Recent trends in surface metrology, Wear 271 (3-4) (2011) 494–508.
- [24] H. Donghui, W. Pingsheng, T. Weizhi, Z. Dequan, C. Guangzhou, N. Bangfa, Z. Xiuhua, L. Lin, Z. Guiying, L. Cunxiong, L. Dehong, Study on real-time wear measurement of piston-ring and cylinder-bore in an engine using thin layer activation method, Applied radiation and isotopes 66 (8) (2008) 1073–1078.
- [25] S. Rubach, T. Riemer, J. Valentin, C. Delto, Wear detection on cylinder liners with optical 3D measuring technology, MTZ worldwide 75 (3) (2014) 38–43.
- [26] P. Ilanthirayan, S. Mohanraj, M. Kalayarasan, Wear analysis of top piston ring to reduce top ring reversal bore wear., Tribology in industry 39 (4) (2017) 487–494.
- [27] S. H. R. Ali, H. H. Mohamed, M. K. Bedewy, Identifying cylinder liner wear using precise coordinate measurements, International journal of precision engineering and manufacturing 10 (5) (2009) 19–25.
- [28] R. B. Obara, Quantification of wear in cylinder bores based on wear volume profile assessment, Journal of tribology 140 (5), 051608 (2018).
- [29] R. B. Obara, A. Sinatora, Quantification of cylinder bores almost zero-wear, Wear 364 (2016) 224–232.
- [30] L. Gara, Q. Zou, B. Sangeorzan, G. Barber, H. McCormick, M. Mekari, Wear measurement of the cylinder liner of a single cylinder diesel engine using a replication method, Wear 268 (3-4) (2010) 558–564.

- [31] R. Kumar, S. Kumar, B. Prakash, A. Sethuramiah, Assessment of engine liner wear from bearing area curves, Wear 239 (2) (2000) 282–286.
- [32] S. Mezghani, I. Demirci, M. Yousfi, M. El Mansori, Running-in wear modeling of honed surface for combustion engine cylinderliners, Wear 302 (1-2) (2013) 1360–1369.
- [33] F. Cabanettes, B. G. Rosén, Topography changes observation during running-in of rolling contacts, Wear 315 (1-2) (2014) 78–86.
- [34] O. Burggraaff, N. Schmidt, J. Zamorano, K. Pauly, S. Pascual, C. Tapia, E. Spyrakos, F. Snik, Standardized spectral and radiometric calibration of consumer cameras, Optics express 27 (14) (2019) 19075–19101.
- [35] S. Kanchi, M. I. Sabela, P. S. Mdluli, Inamuddin, K. Bisetty, Smartphone based bioanalytical and diagnosis applications: a review, Biosensors and bioelectronics 102 (2018) 136–149.
- [36] C. Villani, Optimal transport: old and new, Vol. 338, Springer science & business media, 2008.
- [37] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein generative adversarial networks, in: Proceedings of the 34th international conference on machine learning - Volume 70, 2017, pp. 214–223.
- [38] A. Ramdas, N. G. Trillos, M. Cuturi, On wasserstein two-sample testing and related families of nonparametric tests, Entropy 19 (2), 47 (2017).
- [39] M. C. Salcedo, I. B. Coral, G. V. Ochoa, Characterization of surface topography with Abbott Firestone curve, Contemporary engineering sciences 11 (68) (2018) 3397–3407.
- [40] M. Bigerelle, A. Iost, A numerical method to calculate the Abbott parameters: a wear application, Tribology international 40 (9) (2007) 1319–1334.
- [41] S. H. Hill, T. C. Kantola, J. R. Brown, J. Hamelink, An experimental study of the effect of cylinder bore finish on engine oil consumption, SAE transactions 104 (1995) 1568–1576.
- [42] O. Sato, M. Takiguchi, A. Takayuki, Y. Seki, K. Fujimura, Y. Tateishi, Improvement of piston lubrication in a diesel engine by means of cylinder surface roughness, Tech. rep., SAE technical paper (2004).
- [43] J. Schmid, Optimiertes Honverfahren f
 ür Gusseisen-Laufflächen, VDI-Berichte 1906 (2005) 217.
- [44] J. Schmid, Large engine cylinder honing as a contribution to emissions reduction, MTZ industrial 3 (1) (2013) 48–53.
- [45] E. Tomanik, Liner honing quality main characteristics, Tech. rep., SAE technical paper (1992).
- [46] D. Hrdina, M. Maurizi, B. Lemm, H. Kahraman, G. S. de Faria, Variably honed cylinder liners, iron-based cast pistons and variably coated piston rings as PCU system for friction loss and TCO reduction, in: Heavy-duty-, on-und off-highway-motoren 2019, Springer, 2020, pp. 83–98.

- [47] T. Hoen, J. Schmid, W. Stumpf, Less wear and oil consumption through helical slide honing of engines by Deutz, MTZ worldwide 70 (4) (2009) 46–51.
- [48] F. Welzel, Tribologische Optimierung von Zylinderlaufflächen in Verbrennungsmotoren aus fertigungstechnischer Sicht, Ph.D. thesis, Universität Magdeburg (2014).
- [49] W. P. Dong, E. J. Davis, D. L. Butler, K. J. Stout, Topographic features of cylinder liners — an application of three-dimensional characterization techniques, Tribology international 28 (7) (1995) 453–463.
- [50] W. W. Pulkrabek, Engineering fundamentals of the internal combustion engine, 2nd ed., Journal of engineering for gas turbines and power 126 (1) (2004) 198–198.
- [51] Y. Wakuri, M. Soejima, T. Yamamoto, M. Ootsubo, Fundamental studies on scuffing between cylinder and piston ring, International journal of vehicle design 9 (2) (1988) 203–215.
- [52] Z. Ye, C. Zhang, Y. Wang, H. Cheng, S. Tung, Q. J. Wang, X. He, An experimental investigation of piston skirt scuffing: a piston scuffing apparatus, experiments, and scuffing mechanism analyses, Wear 257 (1-2) (2004) 8–31.
- [53] E. Greuter, S. Zima, Engine failure analysis: internal combustion engine failures and their causes, SAE international warrendale, 2012.
- [54] N. A. Henein, Time dependent wear and its mechanisms in engine cylinders, Tech. rep., Wayne state university Detroit - center for automotive research (1997).
- [55] Alicona infinitefocus confocal microscope, https://www.alicona.com/en/products/ infinitefocus/, accessed: 2020-12-15.
- [56] Mic-fi digitial wi-fi microscope, https://mic-fi.it/en/home/ 43-micfiuvwp-uv-light-polarised-filter-wi-fi-microscope-0806808145420. html, accessed: 2020-12-15.
- [57] Y. Hiasa, Y. Otake, M. Takao, T. Matsuoka, K. Takashima, A. Carass, J. L. Prince, N. Sugano, Y. Sato, Cross-modality image synthesis from unpaired data using cyclegan, in: International workshop on simulation and synthesis in medical imaging, Springer, 2018, pp. 31–41.
- [58] X. Han, MR-based synthetic CT generation using a deep convolutional neural network method, Medical physics 44 (4) (2017) 1408–1419.
- [59] C. Godard, O. Mac Aodha, G. J. Brostow, Unsupervised monocular depth estimation with left-right consistency, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 270–279.
- [60] D. Eigen, C. Puhrsch, R. Fergus, Depth map prediction from a single image using a multiscale deep network, in: Advances in neural information processing systems, 2014, pp. 2366–2374.

- [61] Y. Lei, J. Harms, T. Wang, Y. Liu, H.-K. Shu, A. B. Jani, W. J. Curran, H. Mao, T. Liu, X. Yang, MRI-only based synthetic CT generation using dense cycle consistent generative adversarial networks, Medical physics 46 (8) (2019) 3565–3581.
- [62] S. Ruder, An overview of gradient descent optimization algorithms (2017). arXiv:1609. 04747.
- [63] M. D. Zeiler, ADADELTA: an adaptive learning rate method (2012). arXiv:1212.5701.
- [64] D. P. Kingma, J. Ba, Adam: a method for stochastic optimization (2017). arXiv:1412. 6980.
- [65] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, X. Ding, Embracing imperfect datasets: a review of deep learning solutions for medical image segmentation, Medical image analysis101693 (2020).
- [66] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 3431–3440.
- [67] C. Angermann, M. Haltmeier, Random 2.5D U-net for fully 3D segmentation, in: Machine learning and medical engineering for cardiovascular health and intravascular imaging and computer assisted stenting, Springer, 2019, pp. 158–166.
- [68] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, M. Unser, CNN-based projected gradient descent for consistent CT image reconstruction, IEEE transactions on medical imaging 37 (6) (2018) 1440–1453.
- [69] H. Yao, F. Dai, S. Zhang, Y. Zhang, Q. Tian, C. Xu, Dr2-net: deep residual reconstruction network for image compressive sensing, Neurocomputing 359 (2019) 483–493.
- [70] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition (2015). arXiv:1409.1556.
- [71] S. Antholzer, M. Haltmeier, J. Schwab, Deep learning for photoacoustic tomography from sparse data, Inverse problems in science and engineering 27 (7) (2019) 987–1005.
- [72] S. Shalev-Shwartz, S. Ben-David, Understanding machine learning: from theory to algorithms, Cambridge university press, 2014.
- [73] J. Adler, A. Ringh, O. Öktem, J. Karlsson, Learning to solve inverse problems using wasserstein loss (2017). arXiv:1710.10898.
- [74] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, Journal of machine learning research 15 (56) (2014) 1929–1958.
- [75] S. C. Ng, Principal component analysis to reduce dimension on digital image, Procedia computer science 111 (2017) 113–119.
- [76] S. Wold, K. Esbensen, P. Geladi, Principal component analysis, Chemometrics and intelligent laboratory systems 2 (1-3) (1987) 37–52.

- [77] M. Ahmed, R. Seraj, S. M. S. Islam, The k-means algorithm: a comprehensive survey and performance evaluation, Electronics 9 (8), 1295 (2020).
- [78] L. Xu, M. I. Jordan, On convergence properties of the EM algorithm for gaussian mixtures, Neural computation 8 (1) (1996) 129–151.
- [79] K. P. Murphy, Machine learning : a probabilistic perspective, MIT Press, 2013.
- [80] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [81] F. Chollet, Xception: deep learning with depthwise separable convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 1251– 1258.
- [82] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, Li Fei-Fei, ImageNet: a large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, 2009, pp. 248–255.
- [83] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., ImageNet large scale visual recognition challenge, International journal of computer vision 115 (3) (2015) 211–252.