Computers in Biology and Medicine 151 (2022) 106224

Contents lists available at ScienceDirect

Computers in Biology and Medicine

journal homepage: www.elsevier.com/locate/compbiomed

Classification and regression of stenosis using an in-vitro pulse wave data set: Dependence on heart rate, waveform and location

Alexander Mair [a,1], Michelle Wisotzki [a,1], Stefan Bernhard [a,b,*]

[a] Technische Hochschule Mittelhessen, Department Life Science Engineering, Wiesenstrasse 14, 35390 Gießen, Germany
[b] Freie Universität Berlin, Institute of Mathematics, Berlin, Germany

ARTICLE INFO

Keywords:
Stenosis
Machine learning
Inverse problem
In-vitro simulator
Classification
Regression

ABSTRACT

Background: Data-based approaches promise to use the information in cardiovascular signals to diagnose cardiovascular diseases. Considerable effort has been undertaken in the field of pulse-wave analysis to harness this information. However, the inverse problem, inferring arterial properties from waveform measurements, is not well understood today. Consequently, uncertainties within the estimation hinder the diagnostic application of such methods.

Method: This work contributes a publicly available data set measured at an in-vitro cardiovascular simulator, focusing on a set of input conditions (heart rate, waveform) and stenosis locations. Furthermore, a first attempt is undertaken to perform classification and regression on this data set using standard machine learning methods on features extracted from four peripheral pressure signals.

Results: The locations of six different stenoses could be distinguished at high accuracy of 93%, where transfer function-based features outperformed features based solely on signal shape in almost all cases. Furthermore, regression on the stenosis position could be performed with a root mean square error of 2.4 cm along a 20 cm section of the arterial system using a shallow neural network. However, the performance difference between shape and transfer function features was not clear for this task.

Conclusion: The data set contains 800 measurements and allows investigating the influence of different heart boundary conditions, such as heart rate and waveform shape, on classification and regression tasks. Extracting features that minimise this influence is a promising way of improving the performance of these tasks.

1. Introduction

Recently developed data-based approaches allow to tackle hard problems in many areas of biomedical technology, including analysis of images, biomarkers and biomedical signal processing [1]. Concerning biosignal processing, photoplethysmography (PPG) is a data source that is not fully utilised at the moment. However, recently a lot of studies have been undertaken, trying to diagnose diseases from these signals [2]. The form of a PPG-signal is influenced by structural and elastic properties of the human arterial system. Therefore this signal source might be a promising candidate for non-invasive determination of cardiovascular risk factors and diseases in an early stage.

A variety of diseases cause pathological conditions in the arterial system. Abnormal vascular geometry can manifest in stenosis (reduced diameter) or aneurysms (increased diameter). In addition to the arterial geometry, the elasticity of the vessel wall plays an important role in aneurysms in the presence of hypertension. In clinical practice, geometrical variations can be assessed with common imaging techniques (e.g. computed tomography, magnetic resonance imaging or ultrasound) routinely used to diagnose aortic aneurysms [3]. In contrast, stenosis or, more specifically, the common peripheral arterial diseases (PAD), can also be diagnosed by the ankle-brachial-index (ABI) [4]. However, each of these methods has some shortcomings. For example, there is a large barrier to using imaging techniques, as they are often expensive and unavailable at the family practitioner level. Stenosis screening using ABI returns a single value that does not immediately pinpoint the exact location of the stenosis. Furthermore, ABI cannot be applied in non-peripheral locations that do not allow for a cuff-based blood pressure measurement, such as carotid stenosis. Thus, imaging techniques must be used [5]. A non-invasive and easily applicable method for the early screening of such diseases is desirable.

Because the structural vessel parameters of diameter, length and elasticity all influence the shape of the pulse wave, it has been a long-standing idea to infer these parameters from pulse wave analysis. Initially, the investigation examined whether properties of the

* Corresponding author at: Technische Hochschule Mittelhessen, Department Life Science Engineering, Wiesenstrasse 14, 35390 Gießen, Germany.
E-mail address: stefan.bernhard@lse.thm.de (S. Bernhard).

[1] Alexander Mair and Michelle Wisotzki are joint first authors and contributed equally to this article.

https://doi.org/10.1016/j.compbiomed.2022.106224
Received 27 April 2022; Received in revised form 18 September 2022; Accepted 15 October 2022
Available online 21 October 2022

arterial tree can be inferred from waveform measurements focused on the input impedance (relationship between pressure and flow at the aortic root). This problem was first stated in [6] as a haemodynamic inverse problem, and it was shown that the solutions are not unique. Multiple measurements at different locations could, however, provide the possibility of solving a constraint version of this problem by either classifying the presence of certain geometrical variations or even estimating related parameters like the diameter or location. To date, it is not clear how much information from different measurement locations is required to accomplish these tasks with a certain precision. Recently, the interest in this approach arose again because data-driven approaches allow tackling the problem using machine learning techniques. The required data basis is usually either generated by computer simulations (in-silico), physical simulators (in-vitro) or clinical measurements (in-vivo). While the first two approaches enable direct access to pressure and flow waveforms, in-vivo studies try to acquire similar information (mostly) non-invasively by the measurement of a photoplethysmogramm (PPG).

Much of the work in previous studies on aneurysms was based on simulation data. While most approaches apply machine learning in the time domain [7–9], there is also recent work where the features are based on the Fourier series representation [10], resulting in a reduced parameter space when a limited number of modes is used. However, the underlying model assumptions often reduce the complexity of the generated data, including the number and statistical distribution of parameters that are varied in a virtual patient cohort. The authors in [11] studied a simple model of an arterial bifurcation varying many geometrical parameters, resulting in a reduction of accuracy compared to other studies mentioned in this paragraph. Interestingly, while there is not much literature using in-vitro or in-vivo data for classification, parameter estimation is done on a wealth of data sources, sometimes including all three major data sources listed above (in-silico, in-vitro, in-vivo). Parameter estimation studies [12–15] have been done for a variety of parameters like peripheral resistances and arterial wall properties [16], while [14] established estimates of geometrical properties in the cerebrovascular system. However, estimates for the diameters or locations of stenosis in the major arteries are not yet performed. Detailed investigations on the physical effect of aneurysms on the waveform were performed on in-vitro data [17]. Finally, there are some in-vivo studies on stenosis [18–20] and aneurysm detection [21]. In [18] for example, multiple PPG signals were analysed to compute features depending on the phase shift; [21] used a transfer function approach as an attempt to exclude dominant boundary conditions.

Within this work, we try to close some gaps in the classification and regression of stenosis using in-vitro data. There are no publicly available large in-vitro data sets to validate algorithms for classification and regression of structural variations within the arterial system. Most publicly available pulse wave databases are generated by computer simulations [8,22], which offer flexibility in generating a large amount of data. The development of new algorithms gives hope for applications on real-world data, since real-world data is accompanied by specific challenges that must be tackled and are difficult to infer from simulation-only data. Measurement uncertainties and simplifications of the model could hinder the direct transfer of algorithms from in-silico to in-vivo. On the one hand, this study provides a data set to improve and test algorithms for stenosis detection and localisation on physical in-vitro data as an intermediate step. On the other hand, specific investigations on the performance and accuracy of classification and regression for different stenosis locations under different boundary conditions, such as heart waveforms, heart rates and degrees of stenosis, are presented.

2. Materials and methods

A well-formed data set should contain variations similar to expected inter-patient variations. However, the distributions of these parameters in the general population are not known. Algorithms that work on a data set with large variations should also work if the real variations turn out to be smaller. Thus if two classes are compared, the variations in these classes should be considered; otherwise, a classification task is trivial. In this work, different stenosis locations are referred to as classes, and one possibility of generating variations within each class is to change the input boundary condition. This setting also allows for investigating the influence of different heart conditions on the pressure waves throughout the vascular system for both healthy and diseased conditions. Severe stenoses cause a significant translesional pressure drop and are trivial to detect by comparison of absolute pressure values (as addressed by the ABI). Consequently, we mainly focus on the detection limit of mild stenosis (which would be helpful for early diagnosis), but we also compare the accuracy to that of severe stenosis detection.

The data are acquired and analysed in MATLAB (The MathWorks Inc., Natick, MA, USA) version 2020b. All source code to generate the figures starting from the data set is available in a GitLab repository [23].

2.1. Data set generation

The data set was acquired using the in-vitro simulator MACSim (Major Arterial Cardiovascular Simulator), including the major 33 arteries. A detailed description is given in [24]. The general conditions (vascular geometry, fluid properties, sensor setup, etc.) outlined there are also applicable to this work. The total arterial compliance and the total peripheral resistance was $C = (2.42 \pm 0.09) \cdot 10^{-9}$ m^3/Pa and $R_p = (1.94 \pm 0.02) \cdot 10^8$ Pa/(m^3/s), respectively. The density of the fluid (approx. 40/60 weight % glycerin–water mixture) at room temperature ($\theta = 22.4$ °C) was $\rho = (1.094 \pm 0.002) \cdot 10^3$ kg/m^3, and the kinematic viscosity was $\eta = (3.88 \pm 0.09) \cdot 10^{-3}$ Pa s.

The data set was acquired by performing a parameter sweep over the following parameters:

- 16 different stenosis ‘states’
- 5 different heart rates
- 10 different waveforms

In, total the data set contains 800 ($= 16 \cdot 5 \cdot 10$) independent measurements. While initially acquired to an SQL database, the data were saved to mat-files publicly available at [25]. In the following, the design of the data set is described in more detail:

Variation of stenosis degree and location. The stenosis locations were used to split the data set into two parts, one for classification and one for regression. Five different locations (see Fig. 1 I-V and Table 1) with different degrees of stenosis and a measurement series without stenosis were used to generate the classification data. The regression locations were generated by separating a 20-cm long arterial segment into ten equidistant positions, each 2 cm apart. The regression positions are located on the same arterial segment with similar diameters; thus, the same stenosis degree is maintained. All stenoses were implemented using a 3D-printed clamp squeezing the respective arteries. A closer description is given in [24]. The degrees of stenosis are described in Table 1.

Variation of the heart rate. The heart rate was varied in the range between $HR = 50..70$ bpm using an increment of $\Delta HR = 5$ bpm. For each increment, the maximum ventricle pressure was increased to mimic a similar ejection volume of the ventricular assist device (VAD) used to mimic the heart. The pressure amplitudes, p_A, were varied from 220 mmHg at 50 bpm to 250 mmHg at 70 bpm in 7.5-mmHg steps. Note that this is not the maximum pressure because there is an offset of $p_O = -100$ mmHg.

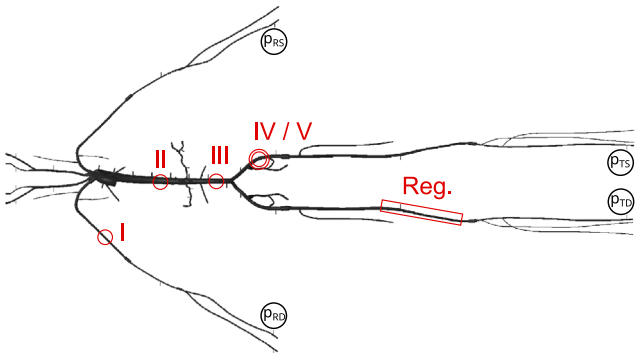

Fig. 1. Overview of the arterial network in the simulator, including the stenosis locations investigated in this work. Locations I-V refer to the measurement scenarios described in Table 1. In the following, node numbers Ref. [24, Figure 14]). Location I is located approx. 32 cm downstream to the aortic arch (between Node 9 and 10). Location II is approx. 7 cm upstream from the renal arteries (Node 28), and Location III is approx. 5 cm downstream from the renal arteries (between Node 34 and Node 35). Locations IV/V are approx. 6 cm downstream from the aortic bifurcation (between Nodes 55 and 59). Finally, the Reg. location is approx. 34 cm to 52 cm downstream from the aortic bifurcation (between Nodes 44 and 45). p_{RS}, p_{RD}, p_{TS} and p_{TD} denote peripheral pressure sensors corresponding to the anatomical locations (radials sinistra, radialis dextra, tibialis sinistra and tibialis dextra). Note: the full data set contains additional signals from 17 pressure sensors and one flow sensor.

Table 1
Description of the locations shown in Fig. 1.

No.	Anatomical location	δ	A_2/A_1
I	Brachialis dextra	12.5%	23.4%
II	Aorta abdominalis suprarenal	42.2%	66.7%
III	Aorta abdominalis infrarenal	29.0%	49.6%
IV	Iliaca communis sinistra I	25.0%	44.1%
V	Iliaca communis sinistra II	3.1%	6.2%
Reg.	(Popliteal dextra)	11.7%	21.9%

It is assumed that the application of the clamp creates a cross-sectional rectangular shape with two half circles attached. δ is the fraction of the distance of the clamp jaws (minus two times the wall thickness) to the initial inner diameter of the artery segment with a circular cross-section. The fraction of cross-sectional inner area A_2/A_1 for this geometry is given by $A_2/A_1 = 2\delta - \delta^2$. A more detailed derivation of this relationship is given in [24].

Variation of heart waveform. A trapezoidal function was chosen as the heart waveform due to the significant change in the waveform shape by adapting the curves' ascending/descending slopes (see Fig. 2). For technical reasons, the waveform was first generated on a normalised timescale $\tilde{t} = t/T$, where T is the time period of the heart rate, and then rescaled to the real period. This gives the benefit that the parameters that determine the shape can be chosen largely independent of the heart rate. The trapezoidal function at normalised time $\tilde{t}$ is given by:

$$p_{in}(\tilde{t}) = \begin{cases} p_O & 0 \le \tilde{t} < \tilde{t}_{a,1} \\ p_O + \frac{\tilde{t}-\tilde{t}_{a,1}}{\tilde{t}_{a,2}-\tilde{t}_{a,1}} p_A & \tilde{t}_{a,1} \le \tilde{t} < \tilde{t}_{a,2} \\ p_O + p_A & \tilde{t}_{a,2} \le \tilde{t} < \tilde{t}_{d,1} \\ p_O + p_A - \frac{\tilde{t}-\tilde{t}_{d,1}}{\tilde{t}_{d,2}-\tilde{t}_{d,1}} p_A & \tilde{t}_{d,1} \le \tilde{t} < \tilde{t}_{d,2} \\ p_O & \tilde{t}_{d,2} \le \tilde{t} < 1 \end{cases} \quad (1)$$

A linear ascent is created between $\tilde{t}_{a,1}$ and $\tilde{t}_{a,2}$ followed by a plateau and a descent between $\tilde{t}_{d,1}$ and $\tilde{t}_{d,2}$. The parameters for p_A and p_O are already described above. $\tilde{t}_{a,1}$ is always chosen to be 0.1, while $\tilde{t}_{a,2} \in [0.15, 0.55]$, $\tilde{t}_{d,1} \in [0.45, 0.85]$ and $\tilde{t}_{d,2} \in [0.5, 0.9]$ all start at the beginning of their intervals and are shifted by $0.4/9 \approx 0.044$ s for each consecutive waveform. Finally, the curve is smoothed with MATLAB's *smoothdata* function with a window of 0.1 before temporal rescaling. The trapezoidal function and the resulting aortic pressure curves in the ascending aorta are shown in Fig. 2. Approaching waveform 10, the peaks become flatter as the pressure valve can completely follow the curve. While this shape is not 'natural', the alteration of the shape can still serve as a basis to investigate the influence of a variation in input conditions. The measurement process was partially automated by scripting the parameter variation of the heart boundary condition and placing the stenosis clamp in the correct position.

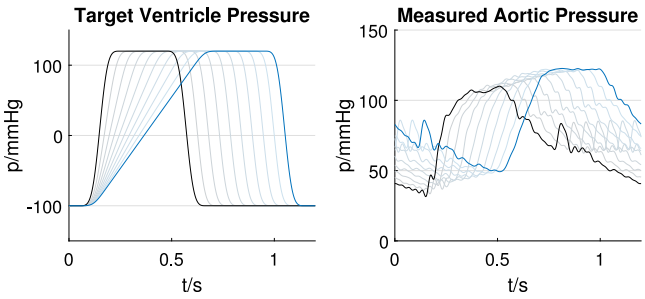

Fig. 2. Exemplary pressure curves for a pulse of 50 bpm (≈ 0.83 Hz). Waveforms from left to right (black to blue) for each plot are waveforms 1–10. The measured pressure from the aorta is filtered with a sharp 30 Hz lowpass filter and averaged over 40 periods starting from the fifth period.

2.2. Preprocessing, feature extraction and data augmentation

To stay as close as possible to clinically measurable quantities, only the pressure curves on four peripheral locations shown in Fig. 1 are used. While it is clear that the pressure curves do not resemble the PPG at these locations, they are the easiest signal to acquire on this hardware simulator. The data are processed in the following steps:

For each of the four pressure curves shown in Fig. 1, ($p_{TS}(t)$, $p_{TD}(t)$, $p_{RS}(t)$, $p_{RD}(t)$), which we will denote as $p_X(t)$:

1. Cut the data $p_X(t)$ from 5 s to 57.5 s to obtain $p_{X,c}(t)$
2. Split the remaining part $p_{X,c}(t)$ in five 10.5-s intervals $p_{X,i}(t)$ with $i = 1..5$
3. Fit a Fourier series to each $p_{X,i}(t)$ with $n = 1..5$ modes to obtain the coefficients $a_{X,i,n}, b_{X,i,n}$ (we discard $a_{X,i,0}$)
4. Create complex coefficients $c_{X,i,n} = a_{X,i,n} - ib_{X,i,n}$ and normalise the coefficients by the amplitude of the first harmonic $\hat{c}_{X,i,n} = c_{X,i,n}/\|c_{X,i,1}\|$

Then either:

5a. Shift the coefficients of each separate signal based on the phase of the first harmonic of that signal. Combine the real and imaginary parts of all signals in a feature array.
5b. Combine the coefficients in pairwise transfer functions and append the real and imaginary parts of these transfer functions to a feature array.

This is similar to the procedure used in [21] on clinical PPG signals but differs on some key points. Each step is described in more detail below.

Step 1. The data are cut before 5 s because the system needs some time to equilibrate, whereas the cut after 57.5 s is made because the signal length differed in the original recording (between 60 s and 70 s).

Step 2. Splitting the intervals serves two purposes. First, it is a form of data augmentation that can be beneficial for the consecutive training of machine learning algorithms. Second, one can, in principle, discard non-valid intervals by removing outliers. For example, this can happen for in-vivo data if there are extrasystoles. The period of 10.5 s is chosen because it does not have an exact match to any frequency of the parameter sweep; exact matches can lead to problems when extracting an estimation of the fundamental frequency in the next step.

Step 3. Features are extracted by fitting a Fourier series of degree $N = 5$ to the measurement data. Thus the data are approximated by:

$$p_{X,i}(t) \approx a_{X,i,0} + \sum_{n=1}^{N} (a_{X,i,n} \cos(2\pi f_0 nt) + b_{X,i,n} \sin(2\pi f_0 nt)) \tag{2}$$

where f_0 is the fundamental frequency of the signal. This procedure contains multiple substeps. First, an approximation of the fundamental frequency is extracted by finding the first real maximum of the circular autocorrelation. Then the normal equation is used to obtain a first estimation of the Fourier coefficients given the estimated frequency. As a last step, MATLAB's *fit* function is used to perform a nonlinear fit (including f_0). In contrast to [21], we do not perform any lowpass filtering or baseline correction on the data. Lowpass filtering does not change the results of the Fourier series fit significantly, and baseline filtering does more harm than good in this case because the data should not exhibit significant baseline drift.

Step 4. Casting the features to the complex representation makes it easier to perform the normalisation and consecutive steps. The information about the absolute amplitude is deliberately taken out of the signal because it is not accessible by non-invasive methods like the PPG.

Step 5a. Because all coefficients have a different phase based on how the signal is cut, we must shift them to make them comparable. A global time shift Δt is converted to a phase shift for different harmonics by $\Delta\phi_n = 2\pi f_0 n \Delta t$. The operation performed in the complex space by

$$\hat{c}'_{X,i,n} = \hat{c}_{X,i,n} / \hat{c}^n_{X,i,1} \cdot \|\hat{c}_{X,i,1}\|^n \tag{3}$$

automatically takes this relationship into account, creating a time shift such that the phase ϕ_1 of the first harmonic $\hat{c}'_{X,i,n}$ is always zero. The last multiplication, preserving the amplitude of the complex coefficient, could be omitted because $\|\hat{c}_{i,1}\|$ is already normalised to 1 in Step 4. After shifting the signals for all four locations in this manner, we append their complex and imaginary values to one line, including also the estimated frequency f_0, as the frequency could bear important information because the response of the system strongly depends on it.

$$[\hat{c}'_{TS,i}, \hat{c}'_{TS,i}, \hat{c}'_{RS,i}, \hat{c}'_{RD,i}, f_0] \tag{4}$$

The third index n being omitted in this case means that each entry represents all modes. As the first mode $n = 1$ does not contain any information, these entries are removed. These adjustments yield a total of 33 real numbers used as features for this approach.

Step 5b. Instead of looking at the signal shape directly, we compute features $h_{X_1X_2,i,n}$, which are labelled as a transfer function in this work, by dividing the complex coefficients for each harmonic (same frequency) for two locations labelled as (X_1, X_2). The relative information between the two signals is captured in this way. Strictly, this is a 'transfer function' between two outputs and not an input/output, for which this terminology is used more widely.

$$h_{X_1X_2,i,n} = \hat{c}_{X_1,i,n} / \hat{c}_{X_2,i,n} \tag{5}$$

This is done four times for the pairs (X_1, X_2) being (RS,TS), (RD,TD), (RS,RD) and (TS,TD). The real and imaginary parts of all modes n of the transfer functions are then appended to a single line and used as features, including the estimated frequency.

$$[h_{RSTS,i}, h_{RDTD}, h_{RSRD,i}, h_{TSTD,i}, f_0] \tag{6}$$

As above, omitting n means that real and imaginary parts of all modes are stacked in the array side by side. Although the first mode is also normalised, it still contains phase information in this case; thus, these coefficients are not removed. Overall, 41 real numbers are used as features here.

The comparison between 5a and 5b allows us to investigate the different importance of the signal shape and the relative changes to the signal because the transfer function features contain information about a time delay and amplitude ratio for each mode between the different measurement positions. In principle, it would be possible to include all information in a feature set, for example, by shifting the signals of different locations by the phase of the signal from one location. This would most likely need a more complex classifier/regression algorithm to make use of the available information. In this article, we want to focus on comparing the signal shape and transfer function features.

2.3. Classification

The subset of data used for classification comprises the locations I-V and a measurement series without stenosis, thus 300 (= 6 · 5 · 10) measurements. We use the augmented data for all our classification problems, effectively having 1500 data points. Cross-validation is not random but chosen in a way that allows us to judge the interpolation and extrapolation abilities with regard to heart rate and heart waveform. For one, we look at a cross-validation of different heart rates where the data for one heart rate are always in the validation set, and the data for the remaining four heart rates are in the training set. Second, we make a split across the waveforms, keeping two consecutive waveforms in the validation set (e.g. waveforms 1 and 2, or waveforms 3 and 4) and the rest in the training set. Thus 80% of the data is used for training and 20% for validation. The way the cross-validation splits are created (non-randomly) already guarantees that data points from the same measurement are either fully in the training set or fully in the validation set. Because the focus of this investigation is specifically on the performance of these splits and not on the overall performance, no separate test data set is used.

Little is known about the structure of the feature space for this problem. Therefore, we compare the most commonly used classifier that we hope can reveal something about this structure. Many other powerful machine learning methods are available, including deep learning; however, only a few easy-to-understand methods were chosen for this first investigation.

K-nearest neighbour (KNN). The KNN classifier works by computing the distance feature space and choosing the k-nearest data points [26, Chapter 13.3]. The label is then decided by a majority decision. We use the Euclidean distance with equal weights (no dimension is prioritised), no weighing by distance and K = 10 neighbours. This is a standard choice in MATLAB using *fitcknn* and fits well with the dimensions of our data. Based on the results of these classifiers, it is possible to tell if the extracted features of the same stenosis location are close to each other in the feature space.

Support vector machine (SVM). A linear support vector machine is used as a second classifier. SVMs are binary classifiers [26, Chapter 12]; thus, a special method is needed to distinguish the six classes in our case. MATLAB's function *fitcecoc* implements by default a one vs one approach training 15 classifiers for the six classes in our case. Each classifier is trained to distinguish two classes from each other. Afterwards, the results of all classifications are evaluated in a unified loss function [27] to determine the predicted class. The results of this classifier can give an indication of whether the data are linearly separable by a hyperplane.

Neural network. Feed-Forward Neural Networks (see e.g. [26, Chapter 11]) are known for their properties as universal approximators and thus can in theory represent any function describing even a complex feature space. However, this is only true for either an arbitrary width of a network with one hidden layer [28], or an arbitrary depth for a network with hidden layers of width greater than the input dimension plus one [29]. In this work, only a shallow network, with one hidden layer and 20 nodes, is used to see how the performance compares to the

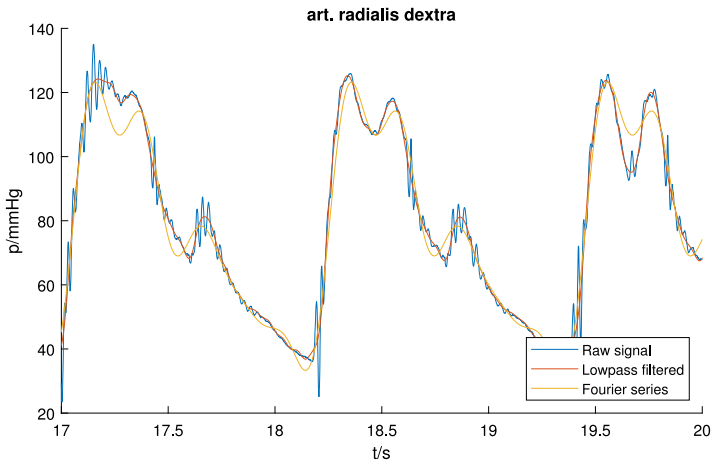

Fig. 3. Example of a raw signal, a filtered version (15 Hz sharp lowpass, see the code for further details [23]) and the Fourier series fit with $N = 5$. The fit process also averages the results as can be seen in the first and third periods; the signal can sometimes either dip below or above the average curve form. Filtering is only used for this figure but nowhere else in the results. High-frequency oscillations are most likely caused by mechanical vibrations in the experimental setup.

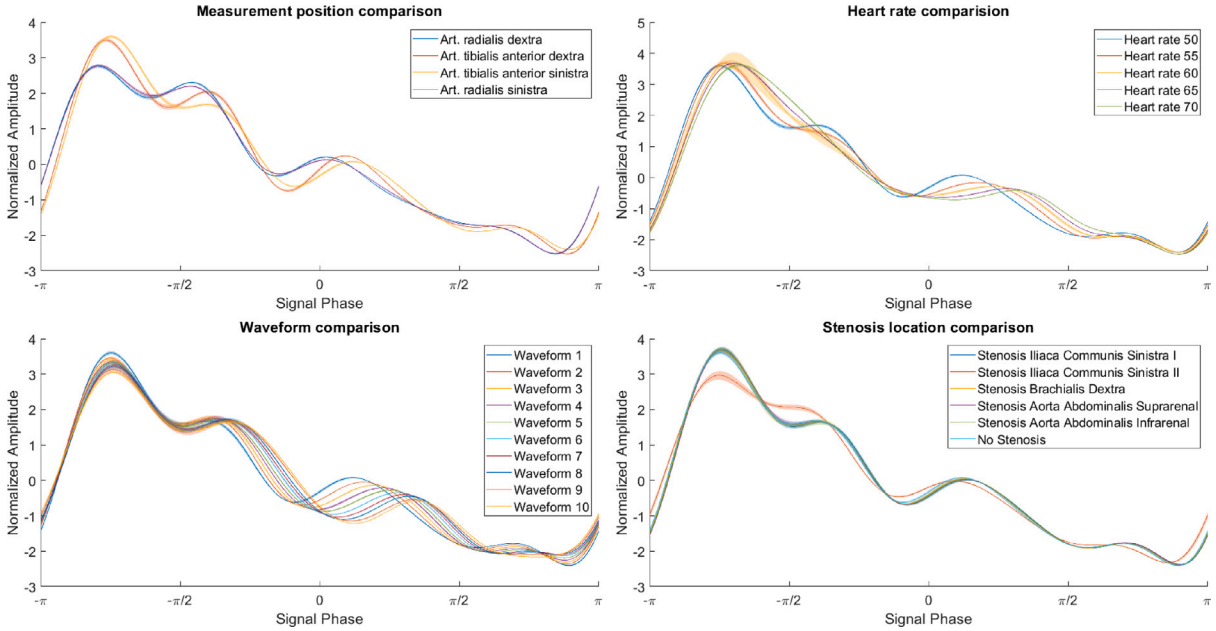

Fig. 4. Comparison of the signal shape obtained by transforming the pre-extraction Fourier coefficients $\hat{c}'_{X,i}$ back to the time domain. For comparison, this transformation always happens with the same number of samples, disregarding the different frequencies. Thus, the signal phase is used as an x-axis instead of a physical time axis. The base point for this comparison is the measurement with no stenosis, at heart rate 50, waveform 1 measured at the arteria tibialis anterior sinistra. All four panels then compare the signal shape along the dimension mentioned in the title. Because the coefficients are extracted over five intervals in the data augmentations steps, five independent time series (i = 1..5) can be extracted. Plots show the mean curves +/− the pointwise standard deviations (shaded). Note that the relative phase between signals cannot be seen in this figure, as we removed it deliberately during the feature extraction.

other methods. We hope this yields a more general approach than the KNN or SVM, even given the limitations that there are no theorems to date about trainability, and we can, of course, not choose an arbitrary size or number of layers. In Section 2.2, we have outlined the different number of features (41/33) for each feature extraction approach. Thus the transfer function features with layers of size 41/20/6 use more weights than the shape features with layers of size 33/20/6. MATLAB's function *patternnet* is used for all calculations.

2.4. Regression

Regression is performed for the stenosis locations labelled as Reg. in Fig. 1. The position was chosen after the only flow sensor in the system (which is not depicted in the figure but included in the public data set). First, the stenosis was placed directly after the flow sensor (see the red square in Fig. 1; a more detailed location can be seen in [24, (Node 45)]. Afterwards, the sensor was moved downstream in 2-cm steps. The distance d was, however, measured from the change of tube diameter to the right of the rectangle in Fig. 1; thus, positions 1–10 correspond to distances 22,20, . . . ,2 in our predictions. A total of 500 measurements are in this part of the data set. With data augmentation, this gives a size of 2500.

Similar to Section 2.3, the cross-validation is chosen to investigate heart rate and waveform dependency, using the identical logic.

Linear regression. Linear regression aims to find a simple relationship between the features x_i and the target y with a proportionality coefficient β_i and an offset β_0 [26, Chapter 3.2]. In our case, x_is are the real and imaginary parts of the Fourier coefficients and the fundamental frequency described in Section 2.2, and y the above-described distance

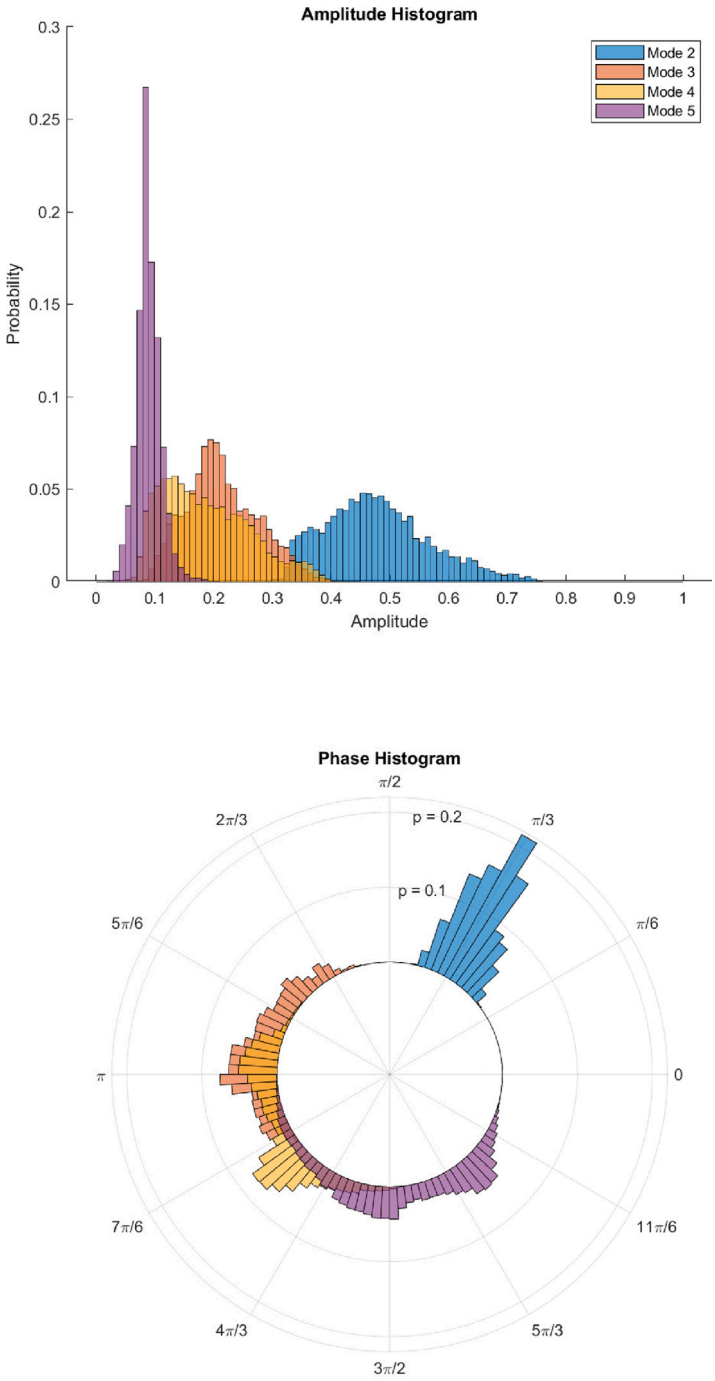

Fig. 5. Histograms of the features describing the signal shape (see Section 2.2 Step 5a). Amplitudes are the complex amplitudes of the respective Fourier coefficients ($\sqrt{Re(c)^2 + Im(c)^2}$), and the phase is the respective complex phase ($\arctan 2(Re(c), Im(c))$). The legend of the amplitude histograms also applies to the phase histogram. Both histograms are normalised with respect to the probability. The binwidth of the amplitude histogram is 0.1, and the binwidth of the phase histogram is $2\pi/100$.

d.

$$\hat{y} = \beta_0 + \sum_{i=1}^{N_f} \beta_i x_i \tag{7}$$

For a given set of N_f data points, this problem is solved with least square optimisation yielding the coefficients β_0 and β_i, which gives a total of 34 coefficients for the shape features and 42 coefficients for the transfer function features. MATLAB's *fitlm* is used with default parameters to fit the model.

Linear regression + interaction. While linear regression provides an easy model, it fails to account for possible relationships between the variables. As a further approximation, one can thus consider combinations of the feature by multiplying them [26, Chapter 3.2]. The complete regression is described by:

$$\hat{y} = \beta_0 + \sum_{i=1}^{N_f} \beta_i x_i + \sum_{i=1}^{N_f} \sum_{j=i+1}^{N_f} \beta_{i,j} x_i x_j \tag{8}$$

Matlab's *fitlm* is used with the *interactions* option. The number of coefficients in this case grows quadratically with the number of data points N_f to $1 + N_f + N_f(N_f - 1)/2$.

Neural network. As mentioned in the classification section, neural networks are universal approximators and can thus, in principle, represent any real valued function when the input and target space are bounded.

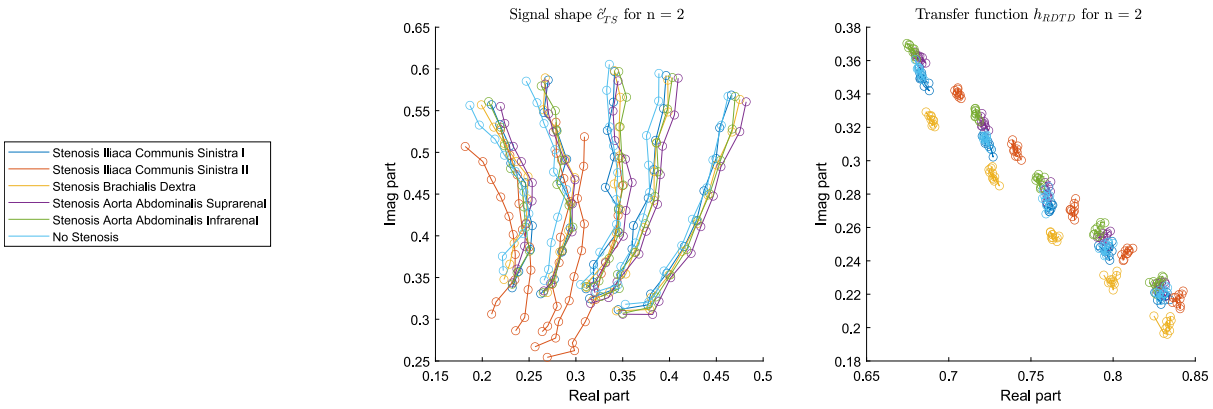

Fig. 6. Scatter plot of features in selected dimensions. Shape features are shown in the middle plot compared to transfer function features in the right plot. To remove noise, an average is shown in this plot. Therefore $\hat{c}'_{TS,i,2}$ and $h_{RDTD,i,2}$ are averaged over the features extracted from five consecutive intervals of the same signal (i = 1..5). Circles show these averaged points; lines connect features with the same waveform. One can see that transfer function features are more concentrated in a single spot. The points in the right plot are also connected by lines that are not visible because they collapse in this representation. Note that these features come from different sides of the body.

These properties suggest using neural networks to see whether it is, in principle, possible to fit the data for the regression problem. In this work, MATLAB's *fitnet* function is used with a hidden layer of size 20.

3. Results

3.1. Data set generation

The recorded signals have a signal-to-noise-ratio, computed with MATLAB's *snr* function, of (17.22 ± 1.62) dB (we always use the mean $\pm$ standard deviation in the following). Fitting a Fourier series to the signal as depicted in Fig. 3 works sufficiently well. In all cases, the correct frequency is reached in the non-linear fitting process with an average deviation of $(0.05 \pm 0.04)\%$. To show the difference between the filtered signal and the Fourier series approximation, a sharp lowpass filter was designed with MATLAB's *designfilt* function at 15 Hz and applied with MATLAB's *filtfilt* function to avoid phase shifting (details can be found in the code [23]). Filtering the signal with a lowpass filter before the fitting process only leads to negligible changes in the resulting Fourier coefficients, so this process was omitted for the final results. Hence lowpass filtering is only used once, for illustration in Fig. 3, but not in the rest of the paper. One can see that there is higher frequency noise (15–30 Hz), possibly due to vibrations in the experimental setup. These oscillations also lead to a relatively high pointwise fitting error of on average (3.95 ± 0.47) mmHg. However, much of this noise can be removed or is irrelevant due to the implicit averaging process when fitting the Fourier series.

The final effect of this averaging process can be seen in Fig. 4, where the features are transformed back to the time domain. Note that in this process, information about the absolute phase, amplitude and frequency is removed from the signals. Differences in the shape of the waveform are continuous if the underlying parameter is continuous. For example, in the waveform comparison at $pi/8$, waveform 1 has the highest value, then waveform 2, and so forth until waveform 10 has the lowest value. Waveforms created with an initial rise are spread more widely at this location and have a broader initial peak (a double peak with a dicrotic notch). The location at the arteria tibialis sinistra was chosen to show the significant difference in the signal shape for iliaca communis sinistra II, the most severe stenosis we applied in our measurements. The difference between the other stenosis locations can barely be seen in the time domain, especially when the standard deviations are included. The signals will be shown to be more separable in the feature space in Section 3.2. Comparing the values across the heart rate dimension shows a large standard deviation for a heart rate of 60. This could be due to oscillations in the experimental setup that align with this frequency.

To illustrate the distribution of the extracted Fourier coefficients, histograms of amplitude and phase are depicted in Fig. 5 for the shape features. The distributions include the Fourier coefficients of the signal shape in the arms and the legs. This creates a bimodal distribution that is only visible for mode 4 and mode 5 in the phase histogram. Higher modes seem to be more spread out, consistent with, for example, linear windkessel theories. In these theories, the elasticity of the arteries is modelled as a capacitance, and the influence of such a capacitance scales with the frequency.

3.2. Classification

Visually, it seems that the distribution of different classes in the feature space cannot be described by simple rules. An exemplary scatter plot is shown in Fig. 6. As mentioned, these plots do not originate from the same side of the body. The transfer function-based features h_{RSTS}, related to the location of the iliaca communis inistra II, are so far away from the other features that the difference between the other classes could not be observed well in this representation. Using the transfer function for feature extraction collapses most of the variation associated with different waveforms in this dimension, and thus, we would expect a better classification result. Note that this is not necessarily the case for higher dimensions, as shown later for regression in Fig. 10. The shape features form a grid in this 2D projection. At first glance, it would be very hard to separate the elements that make up this grid. However, because many classifiers are able to consider much higher dimensional relationships, we cannot judge which features are better suited for classification from these plots. Heart rate and waveform seem to create changes largely orthogonal in this dimension.

As an example, Fig. 7 shows the confusion matrix of only one classification method for the waveform split. One can see that severe stenosis (brachialis dextra and iliaca communis sinistra II) can be separated very well from the other states. Troubles arise only for weak stenosis (compare with Table 1). Unsurprisingly, the aorta abdominalis suprarenal and infrarenal states are often mistaken for each other. However, they are also sometimes classified as no stenosis or mistaken for the weak stenosis on iliaca communis sinistra I.

More detailed accuracies of the waveform split for different features and methods are shown in Fig. 8. Accuracy is computed by the number of correctly classified instances divided by the total number of instances in the test set. The baseline accuracy for this six-class problem is 100/6, or roughly 17%. One can see a clear dip on the borders that suggests that extrapolation is harder for this case than interpolation. Transfer function features perform better than the extracted shape features for all classifiers. The linear SVM classifier reaches the highest accuracy together with the shallow neural network, both with an average accuracy

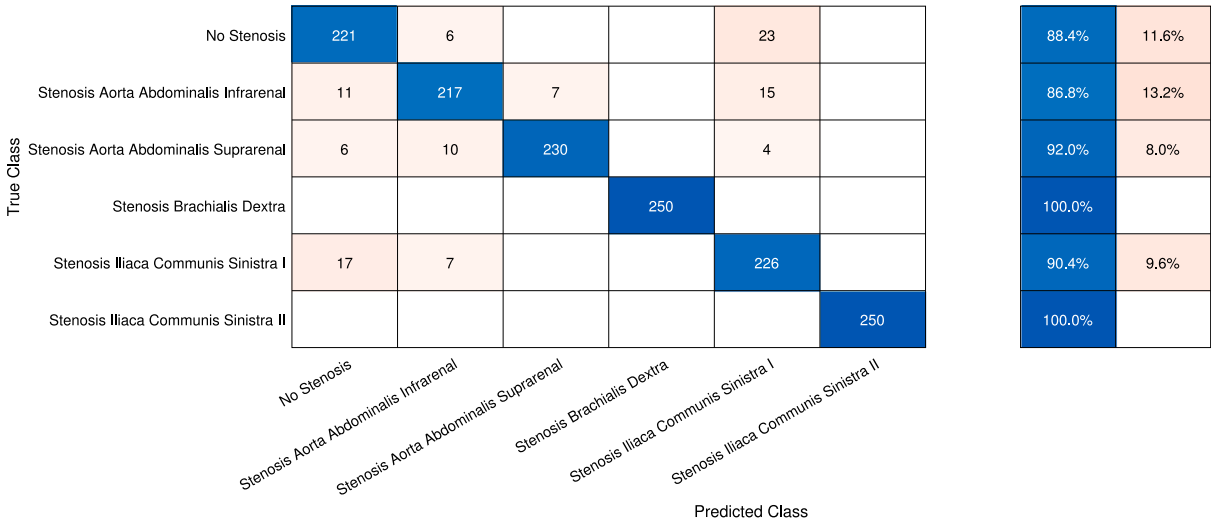

Fig. 7. Confusion matrix using the waveforms as a cross-validation split and a linear SVM classifier on the transfer function features.

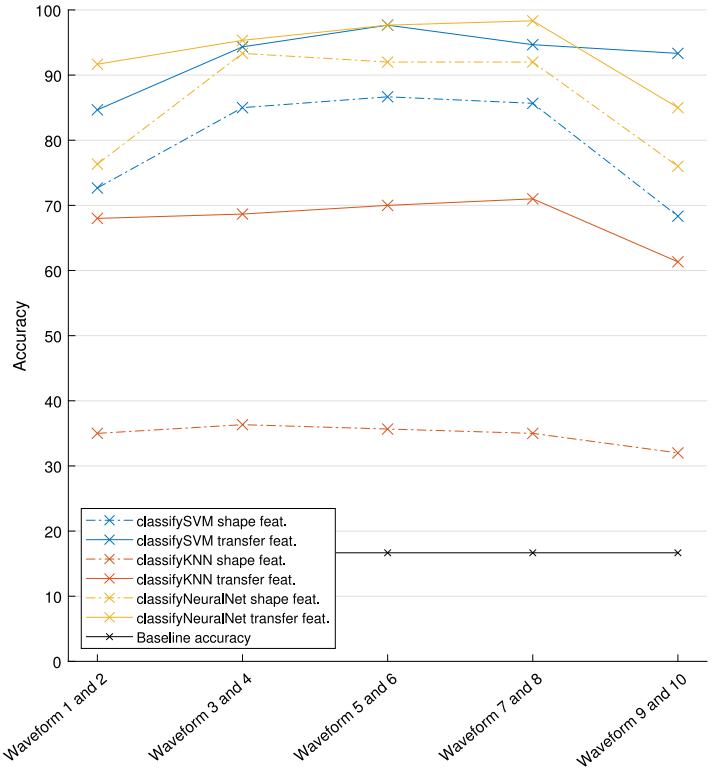

Fig. 8. Accuracy across different folds of the cross-validations. The waveforms listed in the label are used as the validation set, while all other waveforms are used for training.

of around 93%. There is a significant difference between the transfer and the waveform features of around 10% for the linear SVM. This is somewhat intuitive given the separation shown in Fig. 6. The neural network, in contrast, does not make such a substantial distinction between the shape and transfer function features, indicating that the information is present in both sets but is not as easily separable as the SVM on shape features. The KNN performs well on the transfer function features but not so well on the shape features. This is also expected from Fig. 6; for the waveform split, many neighbours are always present in the blobs for each heart rate, but the distances seem less important for the shape features.

Finally, the heart rate split shown in Fig. 9 shows an even stronger dip on the borders than the waveform split. This makes sense for the transfer function features because the distances in the feature space are much larger between different waveforms than between different heart rates. This also explains why the KNN performs significantly worse on that split even for the transfer function features. There is an additional dip in the middle of the curve at 60 bpm. As indicated in Section 3.1 and Fig. 4, there seem to be more disturbances for this heart rate, possibly due to resonance frequencies causing stronger vibrations in the experimental setup. This also affects the accuracy of the classification in this case.

3.3. Regression

Scatterplots in Fig. 10 show a more or less continuous trajectory in the feature space, especially for the transfer function features. However,

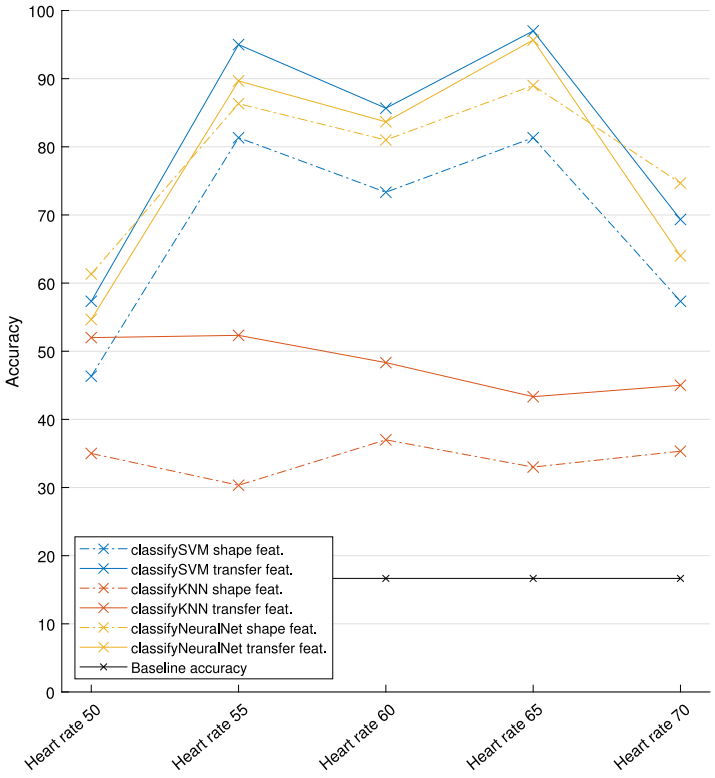

Fig. 9. Comparison of accuracy across different folds of the cross-validation based on heart rate. As in Fig. 8, the labels on the x-axis describe the subset used for validation, while all other data are used for training.

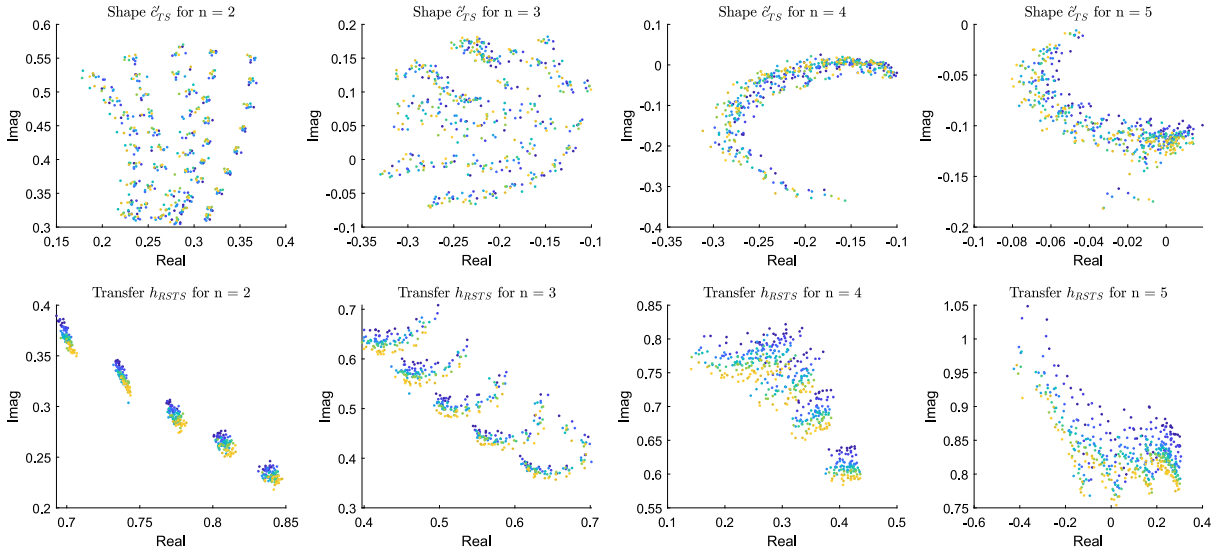

Fig. 10. Scatter plot of the regression data set. The first row shows the shape-based features, while the second row shows the transfer function-based features. The columns represent coefficients that belong to a different Fourier mode ($n = 2..5$ multiples of the fundamental frequency). A blue-to-yellow colourmap (parula) is used to distinguish between different locations, where blue is the most proximal regression point (Position 1), and yellow is the most distal (Position 10). The plots in the first column have a similar shape to the plots in Fig. 6. In the leftmost figure of the bottom row, the five visible clusters correspond to different heart rates.

while the influence of the waveform can be taken out for the second mode ($n = 2$), it is less evident if this is still the case for higher modes. The spread could be due to higher measurement deviations for these modes, or the transfer function approach might not remove the waveform effect as reliably. Note that the data of the first mode of the transfer function still contains information but is not shown in this figure because it is spread on a circle with a radius of 1 and cannot be distinguished well.

The best regression prediction result was reached with linear interaction regression on shape features at an average root mean square error (RMSE) of 2.15 cm. The RMSE is computed by $\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}$, where N is the number of data points in the test

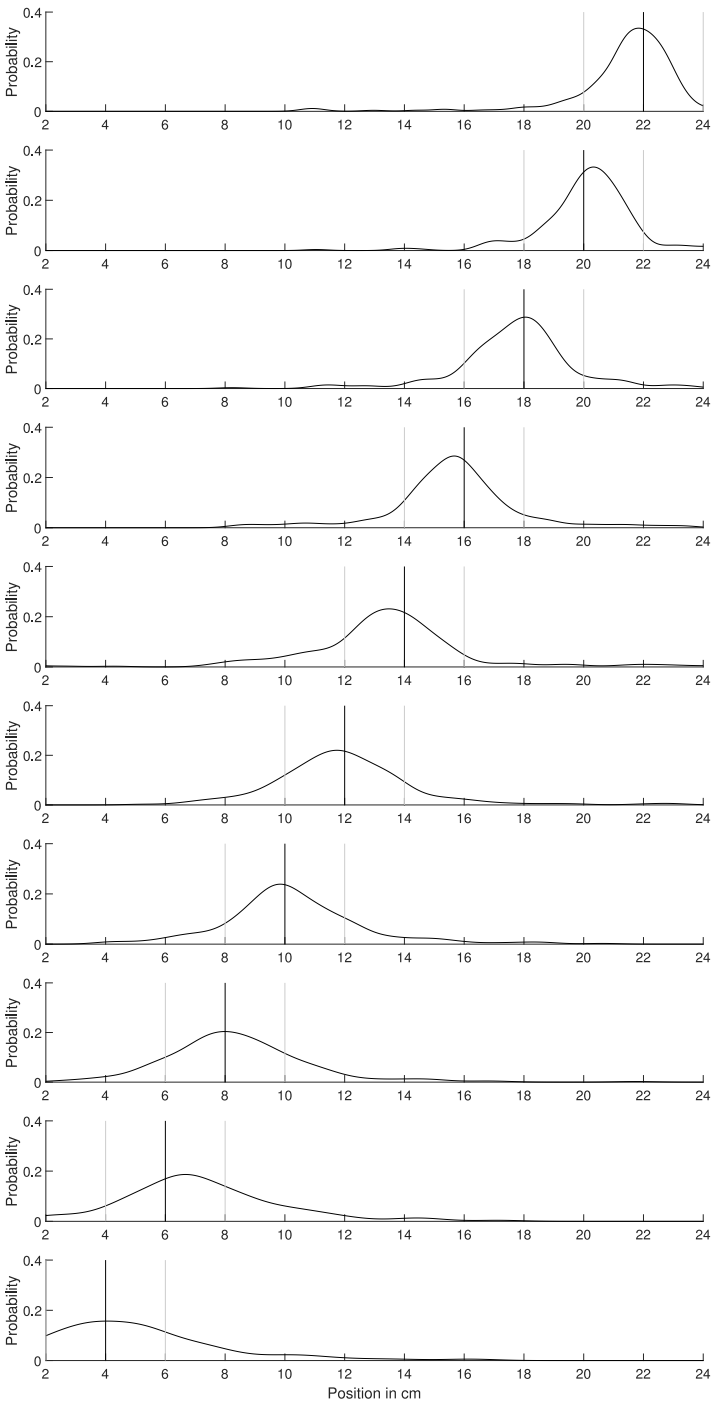

Fig. 11. Distribution of predictions for each regression location. Curves are generated by fitting a probability density function (PDF) (kernel distribution) to the data. Each distribution contains 250 (five heart rates, 10 waveforms, five splits) data points. This figure is based on the waveform cross-validation and uses a shallow neural network for regression on the transfer function features.

set, y_i the real position of the stenosis, and $\hat{y}_i$ the regression result. The results of a shallow neural network on transfer function features are illustrated in Fig. 11 and reach an average RMSE of 2.38 cm. Distributions are unimodal but have rather long tails. Because the distributions are composed of all parts of the cross-validations discussed below, the tails are likely caused by validation sets for which extrapolation would be necessary, which generally have poorer performance.

A comparison across the waveform cross-validation is shown in Fig. 12. The performance for the classification tasks does not differ greatly for neural net regression. Linear regression is remarkably stable everywhere, even for the validation sets that require extrapolation. There is no significant difference between the transfer and shape features in this case. Finally, the linear interaction method only outperforms linear regression for the validation sets that require interpolation (three points in the middle of Fig. 12) but is unstable on the edges.

Cross-validation with different heart rates gives much worse results than waveform cross-validation. As shown in Fig. 13, there are even more significant jumps at the outer folds of the cross-validation (note that the y-axis scale is different in this figure). Surprisingly in this case, the shape features with linear regression give the best performance with an average RMSE of 3.65 cm. Meanwhile, in the middle of the

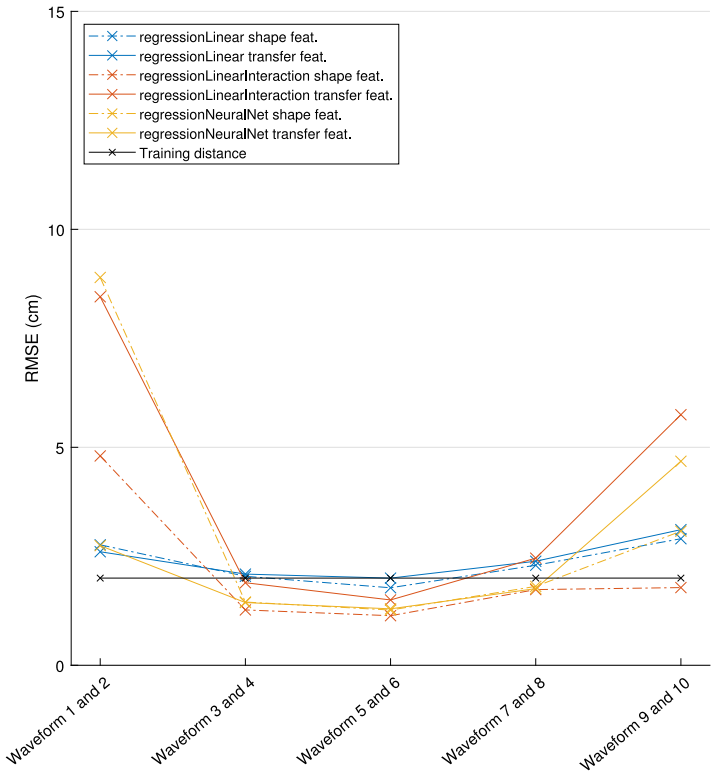

Fig. 12. Comparison of performance of transfer and shape features for three different classifiers (linear regression, linear interaction and SVM). Labels of the x-axis show the parts of the data set used for validation; the rest is used for training. For this figure, the data are split by the heart waveform.

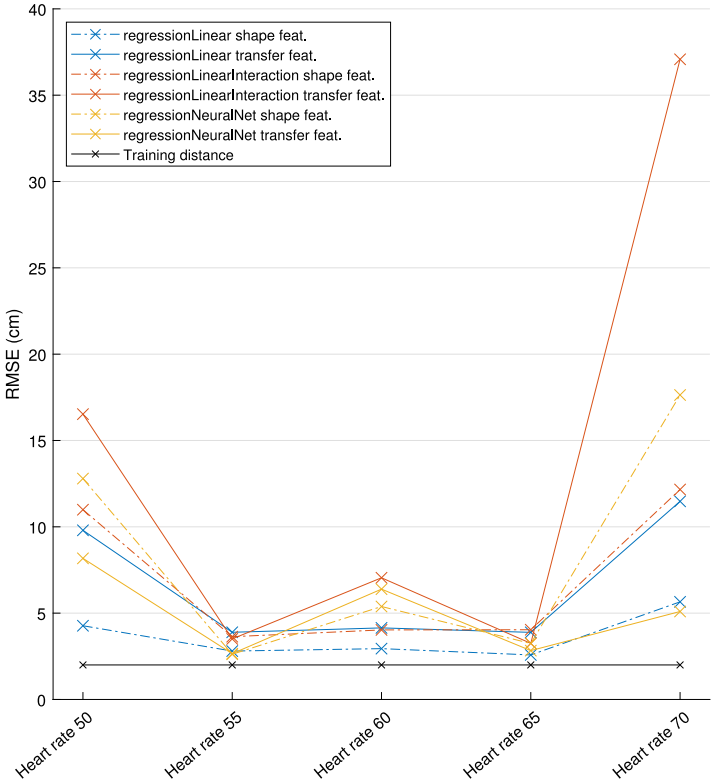

Fig. 13. Results of heart rate cross-validation per fold. Again, the performance of linear regression, linear interaction and SVM is compared for the shape and transfer feature sets. In this case, the split is done for different heart rates.

split, the transfer function features are only worse by ≈ 1 cm, and the edges are much more stable for the shape features. The neural network's performance for transfer function features is comparable to linear regression but not acceptable for the shape features. This effect could be due to the training not converging properly. Linear interaction performs poorly on this cross-validation. Values are comparable to linear regression in the interpolation part, but the jumps at the border are extraordinarily high.

4. Discussion

The created data set shows a reasonable distribution of features and all properties required to investigate the influence of the heart boundary condition for the classification and regression tasks. Features could be extracted reproducibly and with adequate precision, as shown in Fig. 4. It was also shown that the influence of our artificially constructed boundary conditions and heart rate on the waveform shape exceeds the influence of mild changes in arterial geometry. Apart from the most severe stenosis (iliaca communis sinistra II), the waveforms cannot be distinguished visually in the time domain by their shape. However, a closer look at the Fourier coefficients suggests the possibility of distinguishing the stenosis locations. This is further supported by the best classification result achieving 93% accuracy on the six-label classification problem. The substantial separation of features for different heart rates indicates a clear need to include the heart rate as a feature for both classification and regression problems. However, the maximum pressure was raised in conjunction with the heart rate; thus, it is not clear if the significant changes can be attributed to the frequency change only.

Transfer function features can also reduce the influence of the boundary conditions, more precisely the heart waveform, in an experimental setup. The scatter plots in Figs. 6 and 10 show that one dimension can be collapsed, even though this seems to not work as well for the higher Fourier modes. This effect pays off for the classification problem where transfer function features perform in general better than shape features. However, we could not find a clear effect in the case of the regression problem, where, for different methods, either shape or regression features performed better. We have also tried classification and regression on features where relative and shape information was contained using only one signal as a reference to shift the shape features. Because the performance of this feature was not significantly better or worse than the other two methods, the results were not included in this work for clarity of presentation.

Classical machine learning methods like SVM or linear regression work well on both problems. However, the consistently competitive performance of shallow neural networks for both feature types gives an indication that methods that can map more complex boundaries in the features space (e.g. deep learning-based methods) could be preferable in the long run. This is especially the case because, in our data set, there is no significant variation of geometrical parameters except the included stenosis. Including such variation would most likely require a better understanding of the feature space or methods powerful enough to map these complexities. Pre-training on large amounts of in-silico data and later transfer to an in-vitro data set like the one presented here could be a promising path in the future. Further, methods that accurately quantify uncertainties already in the measurement could be investigated on such experimental data sets.

5. Conclusion

The created data set with 800 measurements allows us to investigate the influence of heart rate and waveform on the classification and regression of stenosis. With the Fourier series coefficients as features, the different input conditions create a wider spread in the feature space than mild stenosis. Both classification and regression could be performed with good results but have room for improvement. Transfer functions seem to have an advantage over shape-based features in classification tasks by collapsing one dimension of the input variation but do not perform better for the investigated regression task. The importance of selecting appropriate features to minimise the effect of patient variations in heart conditions is highlighted.

Declaration of competing interest

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

Acknowledgement

The authors would like to thank Urs Hackstein for the fruitful discussion.

References

[1] R. Zemouri, N. Zerhouni, D. Racoceanu, Deep Learning in the Biomedical Applications: Recent and Future Status, Appl. Sci. 9 (8) (2019) 1526, http://dx.doi.org/10.3390/app9081526, URL https://www.mdpi.com/2076-3417/9/8/1526.

[2] H.W. Loh, S. Xu, O. Faust, C.P. Ooi, P.D. Barua, S. Chakraborty, R.-S. Tan, F. Molinari, U.R. Acharya, Application of photoplethysmography signals for healthcare systems: An in-depth review, Comput. Methods Programs Biomed. 216 (2022) 106677, http://dx.doi.org/10.1016/j.cmpb.2022.106677, URL https://www.sciencedirect.com/science/article/pii/S0169260722000621.

[3] S. Aggarwal, A. Qamar, V. Sharma, A. Sharma, Abdominal aortic aneurysm: A comprehensive review, Exp. Clin. Cardiol. 16 (1) (2011) 11–15, URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3076160/.

[4] K. Ouriel, Peripheral arterial disease, Lancet 358 (9289) (2001) 1257–1264, http://dx.doi.org/10.1016/S0140-6736(01)06351-6, URL https://www.sciencedirect.com/science/article/pii/S0140673601063516.

[5] R. Gokaldas, M. Singh, S. Lal, R.J. Benenstein, R. Sahni, Carotid Stenosis: From Diagnosis to Management, Where Do We Stand? Curr. Atheroscler. Rep. 17 (2) (2015) 1, http://dx.doi.org/10.1007/s11883-014-0480-7.

[6] C.M. Quick, W.L. Young, A. Noordergraaf, Infinite number of solutions to the hemodynamic inverse problem, Am. J. Physiol.-Heart Circ. Physiol. 280 (4) (2001) H1472–H1479, http://dx.doi.org/10.1152/ajpheart.2001.280.4.H1472, URL https://journals.physiology.org/doi/full/10.1152/ajpheart.2001.280.4.H1472.

[7] U. Hackstein, S. Krickl, S. Bernhard, Estimation of ARMA-model parameters to describe pathological conditions in cardiovascular system models, Inform. Med. Unlocked 18 (2020) 100310, http://dx.doi.org/10.1016/j.imu.2020.100310, URL https://www.sciencedirect.com/science/article/pii/S2352914819303752.

[8] T. Wang, W. Jin, F. Liang, J. Alastruey, Machine Learning-Based Pulse Wave Analysis for Early Detection of Abdominal Aortic Aneurysms Using In Silico Pulse Waves, Symmetry 13 (5) (2021) 804, http://dx.doi.org/10.3390/sym13050804, URL https://www.mdpi.com/2073-8994/13/5/804.

[9] N.K. Chakshu, I. Sazonov, P. Nithiarasu, Towards enabling a cardiovascular digital twin for human systemic circulation using inverse analysis, Biomech. Model. Mechanobiol. 20 (2) (2021) 449–465, http://dx.doi.org/10.1007/s10237-020-01393-6.

[10] Z. Li, W. He, Stenosis diagnosis based on peripheral arterial and artificial neural network, Netw. Model. Anal. Health Inform. Bioinform. 10 (1) (2021) 13, http://dx.doi.org/10.1007/s13721-021-00290-x.

[11] G. Jones, J. Parr, P. Nithiarasu, S. Pant, A proof of concept study for machine learning application to stenosis detection, Med. Biol. Eng. Comput. 59 (10) (2021) 2085–2114.

[12] D. Lombardi, Inverse problems in 1D hemodynamics on systemic networks: A sequential approach, Int. J. Numer. Methods Biomed. Eng. 30 (2) (2014) 160–179.

[13] S. Pant, C. Corsini, C. Baker, T.-Y. Hsia, G. Pennati, I.E. Vignon-Clementel, Inverse problems in reduced order models of cardiovascular haemodynamics: aspects of data assimilation and heart rate variability, J. R. Soc. Interface 14 (126) (2017) 20160513.

[14] R. Lal, F. Nicoud, E.L. Bars, J. Deverdun, F. Molino, V. Costalat, B. Mohammadi, Non Invasive Blood Flow Features Estimation in Cerebral Arteries from Uncertain Medical Data, Ann. Biomed. Eng. 45 (11) (2017) 2574–2591, http://dx.doi.org/10.1007/s10439-017-1904-7.

[15] L.O. Müller, A. Caiazzo, P.J. Blanco, Reduced-order unscented kalman filter with observations in the frequency domain: application to computational hemodynamics, IEEE Trans. Biomed. Eng. 66 (5) (2018) 1269–1276.

[16] A. Caiazzo, F. Caforio, G. Montecinos, L.O. Muller, P.J. Blanco, E.F. Toro, Assessment of reduced-order unscented Kalman filter for parameter identification in 1-dimensional blood flow models using experimental data, Int. J. Numer. Methods Biomed. Eng. 33 (8) (2017) e2843, http://dx.doi.org/10.1002/cnm.2843, URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.2843.

[17] A. Swillens, L. Lanoye, J. De Backer, N. Stergiopulos, P.R. Verdonck, F. Vermassen, P. Segers, Effect of an Abdominal Aortic Aneurysm on Wave Reflection in the Aorta, IEEE Trans. Biomed. Eng. 55 (5) (2008) 1602–1611, http://dx.doi.org/10.1109/TBME.2007.913994.

[18] Y.-C. Du, A. Stephanus, A Novel Classification Technique of Arteriovenous Fistula Stenosis Evaluation Using Bilateral PPG Analysis, Micromachines 7 (9) (2016) 147, http://dx.doi.org/10.3390/mi7090147, URL https://www.mdpi.com/2072-666X/7/9/147.

[19] N. Mueller, J. Streis, S. Müller, H. Pavenstaedt, T. Felderhoff, S. Reuter, V. Busch, Pulse wave analysis and pulse wave velocity for fistula assessment, Kidney Blood Press. Res. 45 (4) (2020) 576–588.

[20] J. Allen, H. Liu, S. Iqbal, D. Zheng, G. Stansby, Deep learning-based photoplethysmography classification for peripheral arterial disease detection: a proof-of-concept study, Physiol. Meas. 42 (5) (2021) 054002, http://dx.doi.org/10.1088/1361-6579/abf9f3.

[21] U. Hackstein, T. Krüger, A. Mair, C. Degünther, S. Krickl, C. Schlensak, S. Bernhard, Early diagnosis of aortic aneurysms based on the classification of transfer function parameters estimated from two photoplethysmographic signals, Inform. Med. Unlocked 25 (2021) 100652, http://dx.doi.org/10.1016/j.imu.2021.100652, URL https://www.sciencedirect.com/science/article/pii/S2352914821001404.

[22] G. Jones, J. Parr, P. Nithiarasu, S. Pant, A physiologically realistic virtual patient database for the study of arterial haemodynamics, Int. J. Numer. Methods Biomed. Eng. (2021) http://dx.doi.org/10.1002/cnm.3497.

[23] A. Mair, M. Wisotzki, S. Bernhard, Publication GitLab. URL https://gitlab.com/agbernhard.lse.thm/publication_macsim_machinelearning.

[24] M. Wisotzki, A. Mair, P. Schlett, B. Lindner, M. Oberhardt, S. Bernhard, In-vitro Major Arterial Cardiovascular Simulator to generate Benchmark Data Sets for in-silico Model Validation, 2022, arXiv:2204.10005 [physics]. arXiv:2204.10005. URL http://arxiv.org/abs/2204.10005.

[25] S. Bernhard, M. Wisotzki, A. Mair, In-vitro dataset for classification and regression of stenosis: dependence on heart rate, waveform and location, 2022, http://dx.doi.org/10.5281/zenodo.6421497.

[26] T. Hastie, R. Tibshirani, J.H. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Vol. 2, Springer, 2009.

[27] E.L. Allwein, R.E. Schapire, Y. Singer, Reducing multiclass to binary: A unifying approach for margin classifiers, J. Mach. Learn. Res. 1 (Dec) (2000) 113–141.

[28] G. Cybenko, Approximation by superpositions of a sigmoidal function, Math. Control Signals Systems 2 (4) (1989) 303–314, http://dx.doi.org/10.1007/BF02551274.

[29] Z. Lu, H. Pu, F. Wang, Z. Hu, L. Wang, The Expressive Power of Neural Networks: A View from the Width, in: Advances in Neural Information Processing Systems, Vol. 30, Curran Associates, Inc., 2017, URL https://proceedings.neurips.cc/paper/2017/hash/32cbf687880eb1674a07bf717761dd3a-Abstract.html.