Identifying Pathophysiological Intracranial Pressure Waveforms via Fully Convolutional Classification

[1]O. Vrabie, [2]R. Faltermeier, [2]N.O. Schmidt, [2]A. Brawanski and [1] E. W. Lang
[1]CIML Group, Biophysics, University of Regensburg,
93040 Regensburg, Germany
[2]Department of Neurosurgery , University Hospital Regensburg,
93042 Regensburg, Germany
Email: *oleg.vrabie@biologie.uni-regensburg.de*

November 5, 2020

Abstract

It is generally assumed that the analysis of intracranial pressure (ICP) waveforms could be used for the detection of multiple cerebral pathophysiologies. In this study we will demonstrate that ICP waveform analysis utilizing a convolutional neural network (CNN) can be used to distinguish between an intact cerebral compliance and a diminished one. A main obstacle for the analysis of ICP waveforms is given by the large variation of their generating signal, the arterial blood pressure (ABP). Using extended principal component analysis (PCA) we will demonstrate that it is possible to distinguish between ICP waveforms generated by pathophysiological ABP wave forms, for example in case of a heart failure, without loosing the information about the state of the cerebral compliance.

1 Introduction

During treatment of patients suffering from Traumatic Brain Injury (TBI) at an intensive care unit multiple vitals like ICP, ABP and heart rate are recorded continuously and with high resolution. TBI is a category of brain damages that result from external forces applied to the skull. It is classified into primary, i.e. immediate to the initial insult, and secondary, which is an indirect consequence of the initial trauma [24]. Primary TBI is irreversible, whereas secondary TBI can happen from hours to days after the impact and presents the highest risk to the victim. Secondary TBI can be prevented by detecting patterns in continuously recorded vital signs, which indicate an upcoming brain swelling. Early enough detection of these patterns can help in preventing serious brain damages and/or saving patients' lives.

One way to detect these patterns is to study the ICP across different modalities, from exploiting the low frequency variations to inspection of individual pulsatile waves. Low-frequency ICP was investigated in patients with intracranial hypertension [2], TBI [8] and other injuries [25]. On the opposite end of the spectrum, high-frequency pulsatile ICP waves proved to contain morphological patterns correlated with hydrocephalus [5], TBI [26] and other pathologies.

Extracting and examining pulsatile ICP waves rapidly can become an overly demanding task for a human. In this regard considerable efforts have been spent to automatize both extraction and categorization tasks. One prominent example is the morphological clustering and analysis of intracranial pressure (MOCAIP) algorithm [12]. In their study the authors split the analysis problem into three tasks: *delineation and recognition of each individual pulse*; *detection of all peaks in*

each non-artifactual ICP pulse and *designation of the three characteristic peaks assigned to the detected ICP peak candidates* [12]. The MOCAIP algorithm contains five components: The first three focus on solving the first task encompassing the beat-by-beat pulse detector, the pulse clustering and the non-artifactual pulse recognition components. Furthermore, the other two tasks are completed by the peak detection and optimal designation components, ultimately delivering a sequence of locations of three characteristic peaks on the pulse wave. In a follow-up study Hu et al. [11] improved the analysis by defining 24 MOCAIP metrics to describe each ICP pulse, which subsequently were fed to a classifier as features. However, the 24 manually picked metrics contain redundant information and quickly become very time-consuming to be analysed.

At the other end of the spectrum the efforts are focused on leaving the entire feature extraction procedure to a non-biased "agent". An example of this paradigm was carried out in a more recent study where authors developed and clinically evaluated an artificial neural network (ANN) based analysis method for classification of cerebrospinal fluid pulse pressure waveforms (CSFPPW) [22]. They defined four distinct morphological classes of the CSF pulses ranging from normal to pathological. Based on only the three characteristic peaks on the pulse wave, labelled 486 CSFPPWs from 160 patients with different neurosurgical conditions, which formed the training set of the ANN. Prior to feeding the pulses to the ANN, they normalized them and mapped each pulse onto an embedding vector of radial basis function coefficients. After training the network, 60 new CSFPPWs were used to evaluate it's accuracy by comparing the performance of the model with a human expert. They reached 88.3% classification concordance, an unquestionably good result, considering that the model was exposed to less

than 500 examples.

These two studies as well as many others that have not been mentioned rely significantly on human experts. This happens either at the feature definition and selection phases, or during the training examples picking and labelling. Even the choice of the analysis tool, which has a strong impact onto the results, is done without discussing the deeper implications that may lead to biased results, errors or idiosyncratic pictures of the problems to be solved. Safely excluding the human expert from the analysis can be a near-to impossible task but would help bypassing the above mentioned issues. It could be done in phases, as was demonstrated in the automatic feature selection for the MOCAIP algorithm [11], letting the model find patterns on its own using, for example, an ANN-based classifier like in Nucci et al. [22], or combining them towards a more generalizable approach.

In this study we marched towards the same goal - analysis and classification of pulsatile ICP. At the same time we tried to overcome the subtle issues associated with human-in-the-loop paradigm by making use as much as possible of human-detached methods for data extraction, analysis and classification. More specifically, the study pursued two major goals:

- First, we made use of an analysis tool for data extraction called Selected Correlation Analysis (SCA) [9, 10, 23]. It detects brain swellings directly from synchronized raw ABP and ICP recordings. It does that by calculating a so-called **selected correlation (sc)**-index which varies from "selected correlation positive" (scp), which implies a brain swelling with failing autoregulation system, to "no selected correlation" (nsc), which signals no swelling with intact autoregulation. The sc-index was exploited as an ob-

jective label of the extracted pulses. A detailed presentation of the tool set will be revealed below in Section 2.2. Furthermore, instead of hand-picking representative pulses per category we searched for morphological categories in a very low-dimensionality space provided by the patient-wise principal component projection of the pulses. Finally, having defined a dataset, we were able to train and test successfully a Fully Convolutional Neural (FCN) classifier.

- The second focus of this study was directed towards the influence of another vital sign, which gets little to no attention when discussing the shape of ICP pulses, namely the ABP. Because ICP and ABP data is collected simultaneously, and because their pulses occur jointly, one being determined by the other, we were able to extract each ICP-ABP pulse pair. Having chosen three distinct ABP pulse shapes, i. e.

  - one corresponding to a heart dysfunction,
  - one corresponding to reflections created by stiff blood vessels, and finally
  - one that resembles the most common pulse configuration

  ICP pulses could be separated by means of both the sc-index and the type of the input ABP waveform. The FCN model was trained on the six-class dataset reaching a very high accuracy on a hold-out test set. In addition to high classification accuracy, the use of an FCN allowed us to construct class activation maps, which helped visualize exclusively regions in the ICP pulses that determined the pulse classification.

2 Materials and Methods

2.1 Participants and data acquisition

The data was collected from a cohort consisting of seven patients suffering from traumatic brain injury (TBI). The ICP time series were recorded continuously in the ICU using an intraparenchymal piezoelectric sensor, which was inserted into the brain tissue at a depth of around $2\ cm$. The ABP time series were simultaneously recorded at the radial artery. Both quantities were measured continuously with a sampling rate of $1\ kHz$ [10, 16, 28].

The study was conducted in accordance with the ethical guidelines of the Institutional Review Board of the University of Regensburg. Informed consent was obtained from the patient's relatives and all study results were stored and analysed in an anonymized fashion.

2.2 Selected Correlation Analysis

Selected Correlation Analysis (SCA) is a set of mathematical tools that reliably detects negative and positive correlations between arterial blood pressure (ABP) and intracranial pressure (ICP) data, recorded during multimodal monitoring, in a time-resolved way [10]. By splitting both time series into isochronous, fixed length segments, the multi-taper coherence spectrum and the multi-taper power spectra of the segments are calculated. From this spectral information an index called selected correlation (sc) is deduced, which reflects the strength of correlation of the data segments with respect to a specific frequency window. Additionally, the mean Hilbert phase difference of the segments is calculated to reflect the

phasing of the data segments. Using a specific statistical test against false positives the data segments are subsequently labeled as selected correlation positive (scp), selected correlation negative (scn) or no selected correlation (nsc). Scp implies that the data a positively correlated, indicating a diminished compliance with defective autoregulation, whereas scn implies a negative correlation indicating a diminished cerebral compliance with an intact cerebral autoregulation. Nsc means that no significant correlation is detectable indicating no desturbance of cerebral compliance or autoregulation.

As it was shown by Faltermeier et al. [10], SCA can be used to predict the outcome of the analysed patients suffering from TBI. The authors showed that the value of scp was highly correlated ($p < 0.0001$) with the Glasgow Outcome Scale (GOS) [14], an scale ranging from $1 \rightarrow 5$, which measures the outcome of patients with brain injuries. Apart from that they revealed that scp is an independent factor for patients' outcomes ($p < 0.001$), when accounting for age, diagnosis, scp and Glasgow Coma Scale (GCS) in a multiple regression analysis.

In order to asses the relevance of SCA, a further study was realised comparing SCA results with a well-established marker of impaired brain autoregulation called pressure reactivity index (PRx) [23]. SCA indexes were used as labels for the segments they analysed and in result they found that scp-marked portions presented significantly higher mean PRx values than nsc portions. They proved that SCA-computed indexes can be used as labels for the analysed segments, thus accounting for the status of the brain swelling and autoregulation mechanisms as follows:

- scp: brain swelling with failing autoregulation (high mean-PRx),

- scn: brain swelling with intact autoregulation,
- nsc: no brain swelling with intact autoregulation (low mean-PRx).

In the current study, the SCA was used to analyse and categorize raw data, resulting in segments of ICP and ABP recordings described with different values of the three classification numbers. Only scp- and nsc-marked segments were selected from the recordings in order to ease interpretability. In the following section, the utilization of a pre-processing routine is described.

2.3 Pre-processing and cutting of individual pulses

The ICP time series are composed of three main components: slow, respiratory and pulse waves that overlap in the time domain, but can be separated in the frequency domain [2]. In this study, we focused only on the pulse waves, which have a frequency equal to that of the heart rate. Based on the assumption that one could extract physiological patterns that can lead to a swelling of the brain by only analysing the shapes of individual ICP pulses, we proceeded further to the problem extraction of individual pulses from raw recordings.

Utilizing a Matlab package developed by Achhammer [1], the representative pulses were extracted by averaging over a $10\ s$-long window of SCA-catalogued ICP segments. In the first step, the routine is cutting each pulse from minimum to minimum (Fig. 1 panel B). This is followed by the alignment of raw pulses by shifting the curves in order to minimize the Euclidean distance between them (Fig. 1 panels C-D). The mean of the aligned pulses is computed (Fig. 1E), and the resulting prototype pulse is normalized (Fig. 1F). An additional step was included in order to ease the further analysis. It preformed the resizing of the pulses to

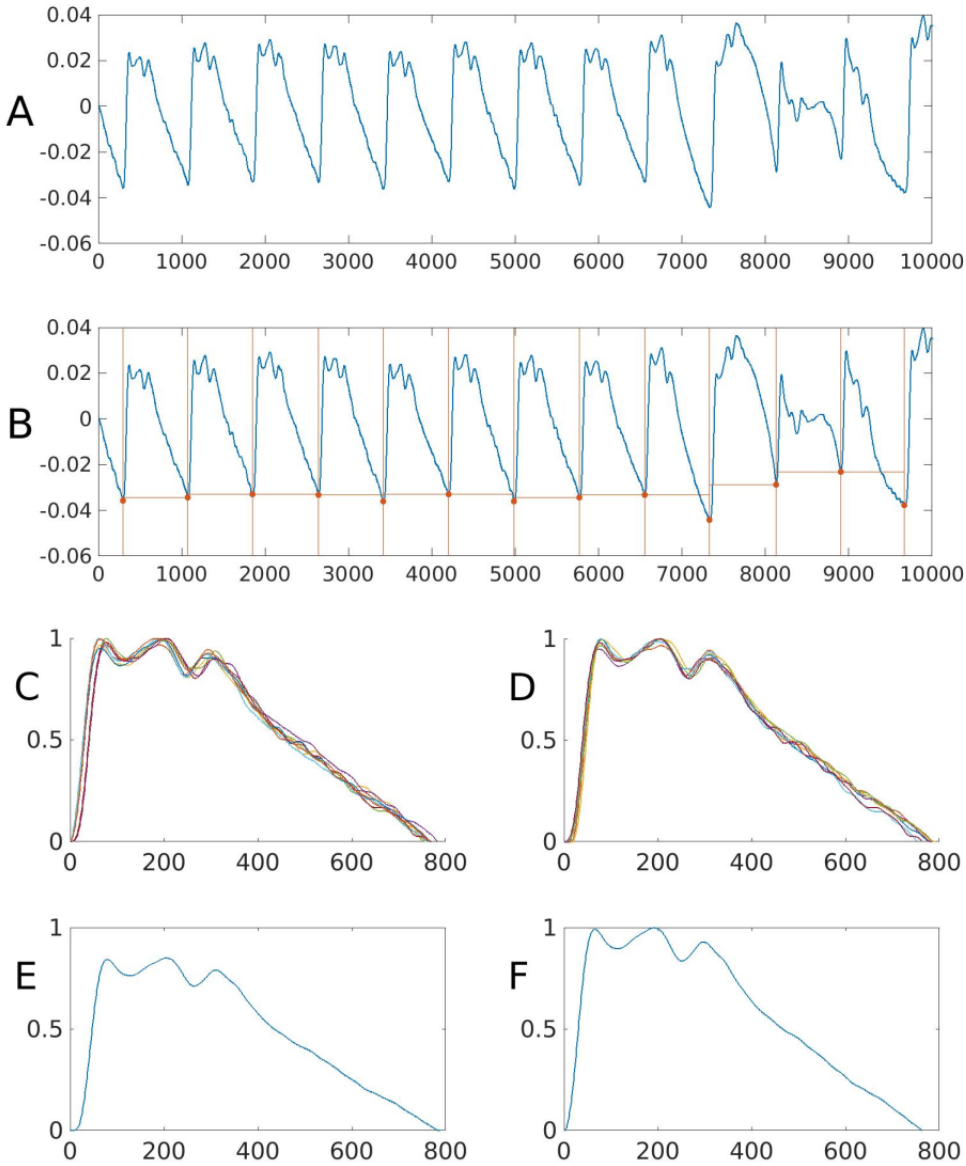

Figure 1: The process of creating a representative ICP pulse wave from a $10s$-long window . The raw sequence (**A**) is cut into pieces, from minimum to minimum (**B**), and then stacked together (**C**). The curves are shifted along the x-axis s.t. the Euclidean distance between pulses gets minimal (**D**). The resulted stack of waves is averaged (**E**) and normalized (**F**). On the x-axis the time is captured in milliseconds. On the y-axis (**A-B**) the ICP amplitude is captured in mmHg, whereas in the following panels is the normalized amplitude. Figure obtained from Achhammer [1].

an average length of 780 sample points. As a result, we ended up with a database (Table 1) consisting of scp and nsc ICP pulses from seven human subjects' recordings.

Table 1: Number of pulses extracted from raw ABP and ICP time-series for each human subject.

	Number of pulses	
Patient ID	scp	nsc
1	9408	2702
2	6053	2258
3	7255	3418
4	3970	1941
5	4181	4941
6	6171	2575
7	6835	3986

2.4 Extracting ICP modes using PCA

In order to visualize the structure of the our data, the principal component analysis (PCA) was used as a means of dimensionality reduction. Given a dataset, this technique transforms to a new space of uncorrelated variables, or principal components (PCs), that contain the most of the variation present in the original variables [15]. The dataset to be fed into a PCA encompassed the extracted ICP pulses from our cohort.

The projection in Fig. 2 of the entire ICP dataset onto the first three PCs revealed that there is no major structural difference that could be responsible for the variability in pulse shape of nsc and scp data. Therefore we proceeded to a patient-wise extraction of modes by means of clustering. Two clustering techniques were used to reveal the hidden structure of the dataset: Gaussian Mixture

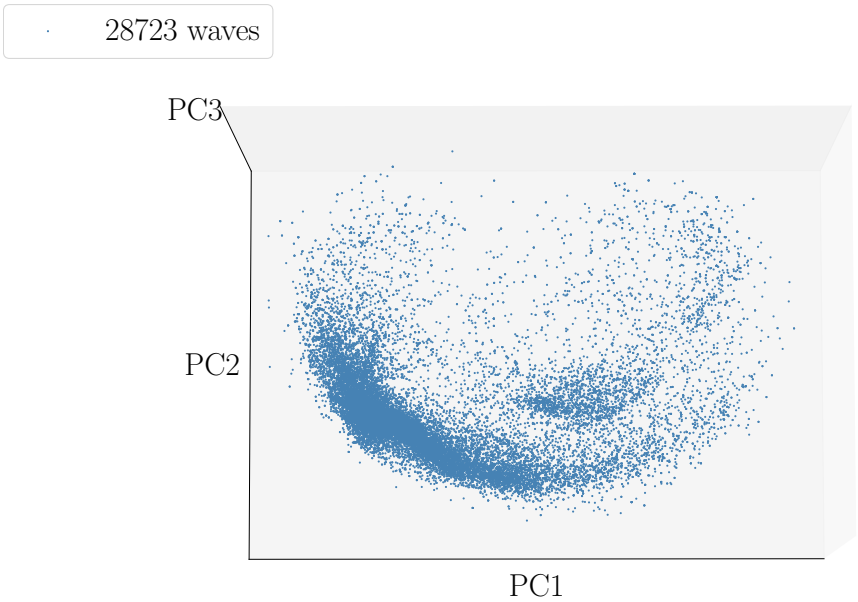

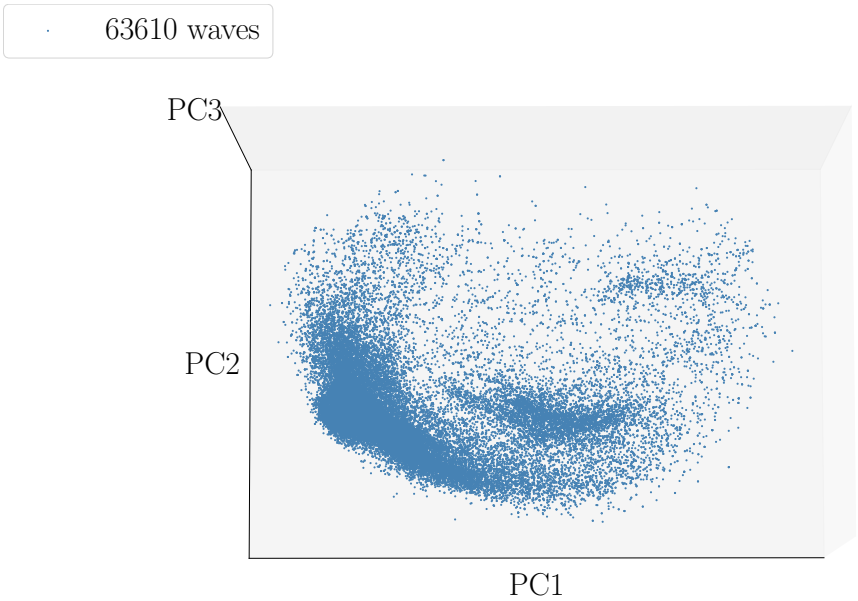

Figure 2: Projections of concatenated nsc (a) and scp (b) ICP pulses from all the seven patients onto first three PCs.

Models (GMMs) and HDBSCAN, both chosen for their flexibility and highly customizable parameters.

2.5 Clustering ICP modes

GMMs represent a probabilistic method that attempts to find a mixture of multi-dimensional Gaussian probability distributions that best fit a dataset. The parameter search for these models is usually based on a maximum likelihood estimation (MLE), using the Expectation-Maximization (EM) algorithm [3]. The advantage of using GMMs for clustering comes from the ability of the user to define the number of components into which to decompose the dataset. HDBSCAN, on the other hand, is an advanced density-based clustering algorithm based on hierarchical density estimates [4, 20]. The main advantage of using HDBSCAN is that it allows clusters with varying densities of points.

The results of the patient-wise clustering of pulsatile ICP are presented in grid plots in Fig. S2 for both nsc and scp. Surprisingly, both selected correlation segments amounted to almost the same number of clusters which they could be separated into, i.e. 46 for nsc and 45 for scp. Apart from regular ICP modes there are also a few very noisy modes in the data. Some of them correspond to movements of the patient lying in bed, others catch the plugging and unplugging of the sensors. Finally, there are also some modes that correspond to improperly cut pulses during data preprocessing.

Another conclusion that can be drawn from the shape of the extracted prototype pulses is that there is a large biological variability of pulse shapes across the patient cohort. This is most probably due to physiological differences between

patients, like age, mean ICP and ABP, but also the severity of the head trauma they suffered from. Nonetheless, our assumption was that these ICP waveforms resemble brain states that are prototypical to a majority if not all patients suffering from a TBI. Therefore, linking the pulse shapes with SCA labels deemed natural, thus we proceeded further to classify these data.

2.6 ICP pulse classification

Following the paradigm used by Nucci et al. [22], we pursued the classification task as follows: With the help of an expert, we selected a list of modes and separated them into four classes to distinguish between the following pulse modes: **nsc**, **scp**, **intermediate** and **noisy**. However, there will be no comparison between the performances of the human expert and an algorithm. Instead, the the SCA labels are considered the reference for the classification. This way we were able to keep the human expert out of the analysis loop, providing an objective, largely automatic and less biased classification procedure.

As choice for the classification algorithm we selected a Fully Convolutional Neural Network (FCN) for its performance and ease of training. A short description of the model together with the details of the architecture are listed in the next subsection.

2.6.1 Fully Convolutional Network for time-series classification

Time series classification falls into the group of problems that innately demand a special treatment which is based on sequence modelling. In the field of Deep Learning, sequence modelling problems are frequently attributed to the realm of

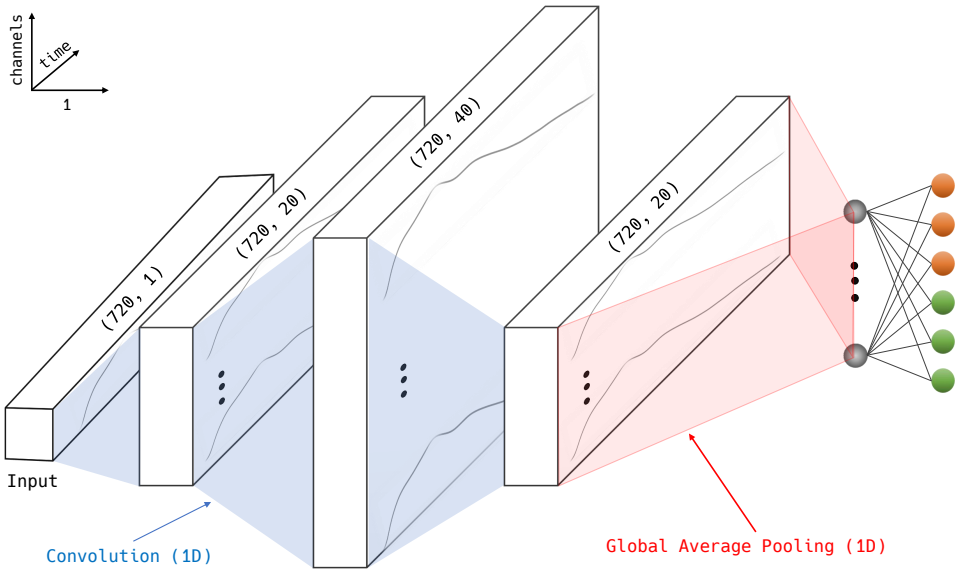

Figure 3: Schematic of the FCN architecture. The 1D input is subjected to three successive convolutions. The weights of the last convolutional layer undergo GAP. This is followed by the last, fully-connected `softmax` layer with 6 nodes, each representing a class (orange for scp; green for nsc).

Recurrent Neural Network (RNN) architectures. However the majority of RNN architectures is still very hard to train, requiring a large amount of compute power and/or compute time.

With that in mind, we transformed a sequence modelling problem into a computer vision problem by treating the time dimension as an additional image dimension similar to, e.g. a colour channel. Computer vision comes with a long and well established list of feed-forward architectures that are well documented, straightforward to implement and easy to train. The go-to architecture in that respect is the Convolutional Neural Network (CNN) [17].

For our purposes, we used a variation of CNNs called Fully Convolutional Neural Networks (FCNs) that have been recently popularized in [19] for semantic segmentation. It recently achieved state-of-the-art performance from scratch

in time series classification [27]. Apart from proving easy to train from scratch, FCNs offer additional significant advantages like, for example, the invariance of the number of parameters in convolutional layers. The difference between a typical CNN and an FCN rests on two main points: (1) there are no local pooling or downsampling layers in between convolutions and (2) the number of parameters is decreased drastically due to the replacement of the final fully connected (FC) layer with a global average pooling (GAP) operation [18].

GAP layers were first proposed by Lin et al. [18] as an alternative to traditional FC layers. Instead of creating a further downsampled feature map, it takes the feature maps of size $(height \times width \times n_channels)$ of the last convolutional layer and computes a single (average) number per channel resulting in $(1 \times 1 \times n)$ channels. This procedure enforces a correspondences between feature maps and categories. Additionally there is no parameter to optimize in the GAP layer, further limiting the risk of overfitting.

Besides the mentioned advantages of using GAPs, these also allow for building class activation maps (CAMs), which highlight regions in the input contributing to the classification, therefore aiding interpretability of the model [13, 27, 29]. In principle, they work as follows: "specific visual patterns in the receptive field of a convolutional layer determine the activation of specific units". These activations are further upsampled to the size of the input image building the input's respective CAM, a procedure which helps identifying most relevant regions per category. See the complete definition of CAMs in 5.2.

The architecture of our model is depicted in 3. It was implemented in Keras [6] and was inspired by the FCN model designed in the review study of [13]. The one-dimensional CAMs that have been used in this study were introduced by

Wang et al. [27]. This allowed us to visualize which regions in the ICP pulses contributed the most to label assignment by our trained FCN in Figs. S7 and S8.

The model was trained on a 42711 sample points long set of waves from the four chosen classes with proportions: $5018 : 21946 : 4486 : 11261$ representing nsc, scp, intermediate and noise, respectively. After a small amount of training epochs, the model converged to a reasonable validation accuracy. The final test was performed on a hold-out test set, consisting of 10678 waves with class proportions: $1274 : 5483 : 1193 : 2728$. This model reached $93.1\ \%$ accuracy on this set.

2.7 Extracting input ABP modes

The fact that the *input* to each ICP pulse is the ABP pulse traveling to the head motivated us to search for identifiable patterns in the ICP pulse shape that are determined by the shape of the input ABP pulse. To accomplish that, we performed the same pre-processing on the ABP pulses as was described in previous sections for the ICP.

After pulse-cutting we ended up with one ABP pulse for each ICP pulse studied earlier. The collective PCA-projection (e.g. Fig. 4) of ABP pulses from all patients revealed a more complex internal structure in comparison with that of the ICP. The high number of compact structures depicted in the projection subspace pointed towards the high number of pulse morphologies that reside in the ABP time series. This led us to go through ABP pulses in a patient-wise manner using the same clustering techniques.

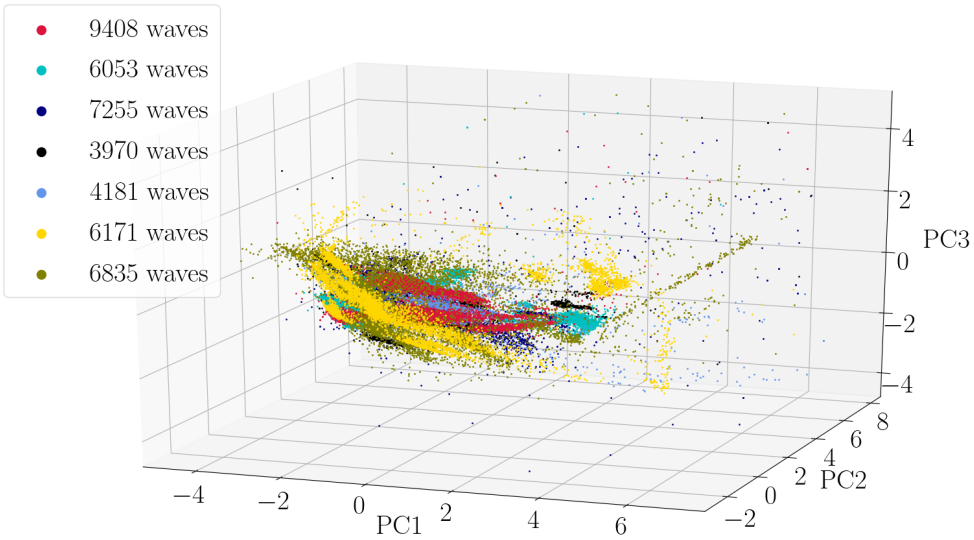

Figure 4: PCA projection of the ABP pulses from **scp** segments. $PC1, PC2$ and $PC3$ represent the first three principal components onto which the ABP pulses were mapped. Different colors correspond to different patients.

2.7.1 Clustering ABP modes

The result of the patient- and selected correlation-wise clustering of pulsatile ABP is presented in grid plots for both nsc and scp in Figs. S3 and S4. At first glance, each patient has at least six different morphologies for both selected correlation types. There are a few noise components, e.g. nsc components $16, 20, 30$ and 44 (Fig. S3); scp components $11, 37, 44, 49$ and 54 (Fig. S4), that probably were caused by external mechanical perturbations. There is moderate variation between pulses coming from the same patient, but there is a strong variation in pulse shape between individuals. Another aspect to be mentioned is the fact that there is a larger number of scp clusters in the ABP data, namely 75 in comparison with 61 nsc clusters, resulting in an unbalaced dataset.

2.7.2 Pathological ABP waveforms

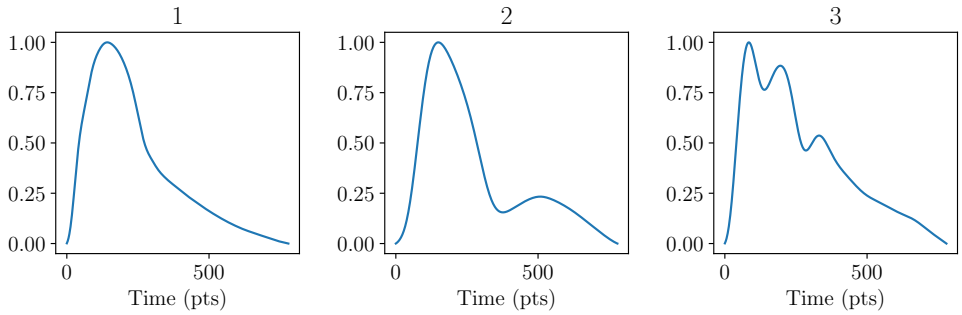

Figure 5: Chosen ABP WFs from our recordings. First representative WF is taken from patient 7 nsc and corresponds to a normal (most common in our cohort) ABP wave; second wave is from patient 10 scp and corresponds to a heart failure; the third comes from patient 37 scp and corresponds to multiple reflections, i.e. vasoconstriction or stiff blood vessels.

Assuming that different ABP pulse morphologies determine different output ICP morphologies, we searched for sources of structural changes in the ABP pulse shape. Some specific studies revealed noticeable changes in the ABP pulse shape that happen during pathological states. One such characteristic type of ABP waveforms that we could identify in our cohort (e.g. see patient 37, component nr. 60) in Fig. S3) has been previously reported and was assigned to *multiple wave reflections* [21]. It was considered to result from abnormal blood vessel structure or vasoconstriction. An exemplar of these waveforms was chosen as reference in Fig. 5 (panel 3). Another example of characteristic waveform changes was recorded in patients with left ventricular systolic dysfunction, a condition that leads to *heart failure* [7]. Subjects suffering from this disease tend to have a notch in the middle of the ABP pulse (see panel 2 in Fig. 5). By visually inspecting the prototypical waveforms from Figs. S3 and S4, we collected, as a function of sc-index, all ABP clusters with prototypes similar to selected waveforms given in

Table 2: The chosen ABP components from figures S4 and S3 divided into three physiological categories: regular, heart failure and multiple reflections, and selected correlation types: positive (scp) and negative (nsc). Notation: index of the ABP component from each selected correlation type.

sc-index	regular	heart failure	reflections
scp	1, 9, 17, 45, 50	7, 15	65, 8
nsc	5, 6, 36, 46	17	31, 59

Fig. 5. The selected ABP components are listed in Table 2.

In order to visualize the shape of ICP pulses corresponding to these chosen ABP components, we aggregated ABP pulses from chosen clusters by pathology category and sc-index (Fig. 6a), and extracted the corresponding to them ICP pulses (Fig. 6b). Furthermore, to analyse the variety of ICP pulses that constructed the means in Fig. 6b, we turned back to PCA. The first three principal components were extracted and used to project the pulses for a further clustering. The task was realised by the HDBSCAN algorithm with the results for scp in Fig. 7 and for nsc in Fig. 8. By calculating the ratios of ABP waveforms (*wf1, wf2 and wf3*) that "influenced" each cluster of ICP pulses, we could identify which resulted component is built-up mainly by a single ABP WF type.

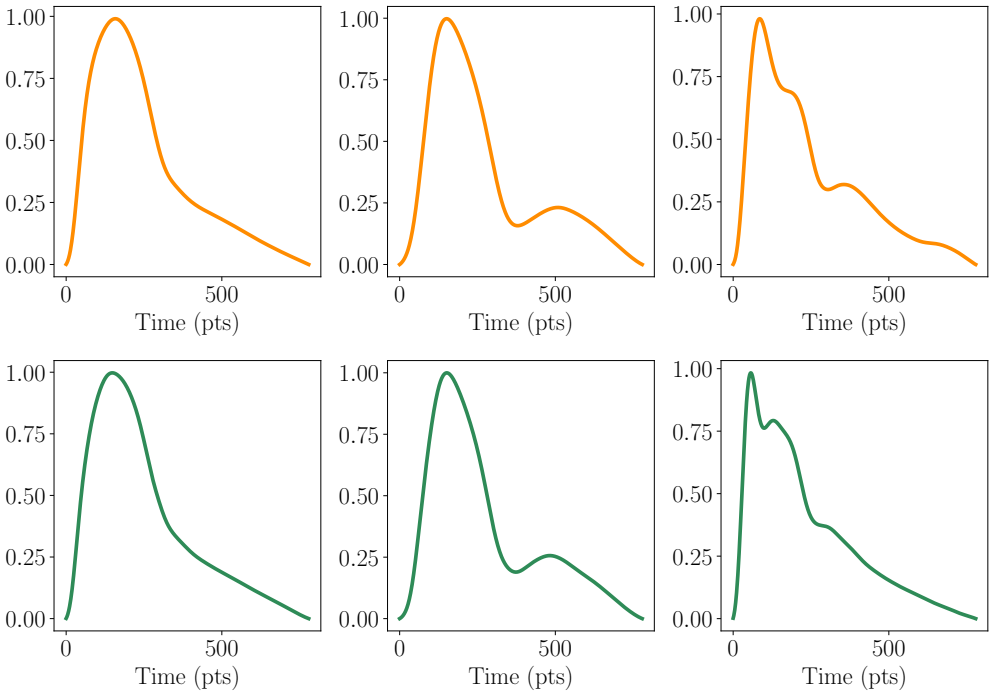
1.00
0.75
0.50
0.25
0.00
0
500
Time (pts)

(a)

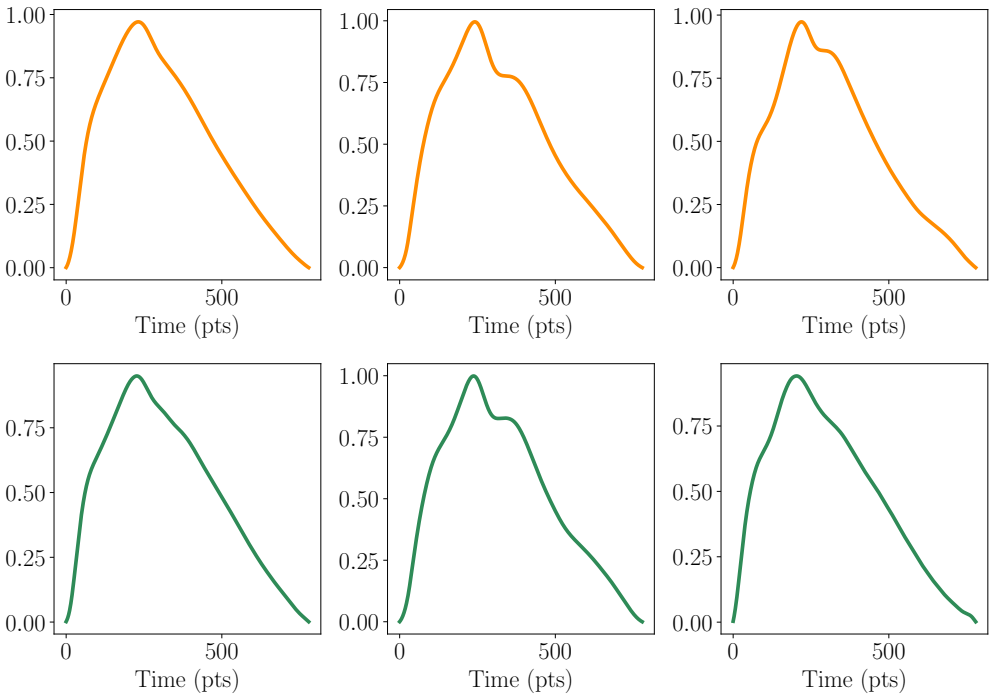
1.00
0.75
0.50
0.25
0.00
0
500
Time (pts)

(b)

Figure 6 *(previous page)*: (**a**) The means of concatenated chosen ABP components by WF and selected correlation. First column corresponds to common ABP WF; second column – to heart failure; third column – to multiple reflections. First row (orange) corresponds to scp and the second (green) to nsc. (**b**) The means of concatenated ICP waves corresponding to chosen ABP WFs by selected correlation. Each column corresponds to each chosen WF, and each row corresponds to selected correlation type: first (orange) : scp, second (green) : nsc.

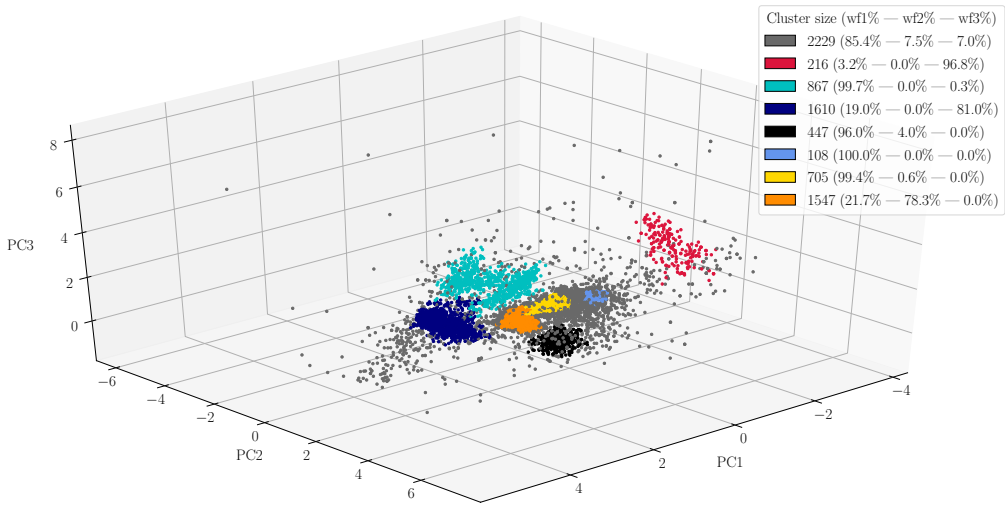

Figure 7: Projection of **scp** ICP waves corresponding to chosen WFs onto first three PCs. Clusters identified by HDBSCAN are marked by different colours. The gray points correspond to noise points considered by the algorithm. In the legend are captured the ratios of each WF that contributed to the found cluster.

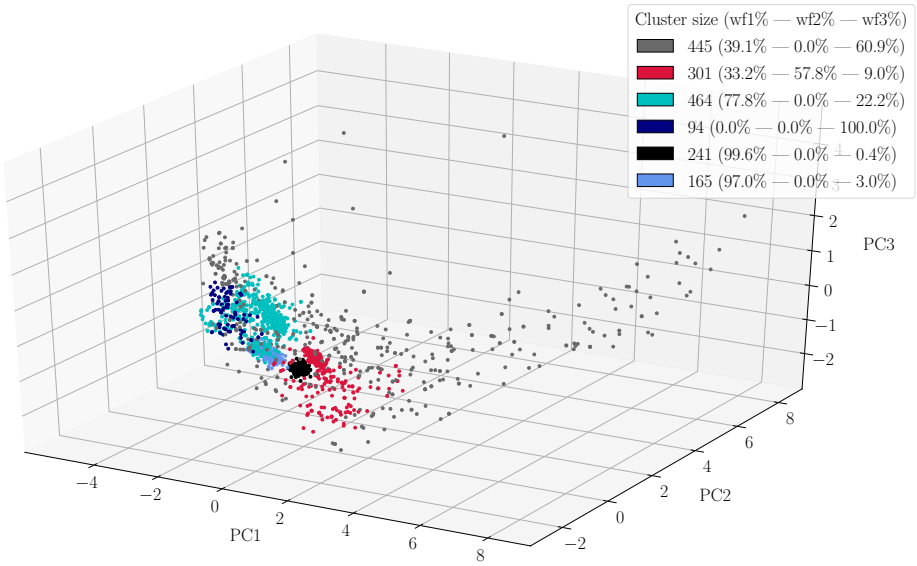

Figure 8: Projection of **nsc** ICP waves corresponding to chosen WFs onto first three PCs. Clusters identified by HDBSCAN are marked by different colours. The gray points correspond to noise points considered by the algorithm. In the legend are captured the ratios of each WF that contributed to the found cluster.

3 Results and Discussion

3.1 Inter-patient ICP pulse classification

The premise of this part of the study was that inspecting data from seven different patients will reveal injured brain states that could potentially capture the transitions between "ignorable" and "dangerous" conditions. Projecting ICP pulses from additional patients unveiled, as anticipated, an increased variety of waves residing in agglomerated clusters. This led us to define an inter-patient dataset that was used to train a classifier, which reached a level of accuracy exceeding 90 % on a test set.

3.2 Selected correlation or input ABP

Given the ratios as shown in Figs. 7 and 8, the next question is: *are we able to distinguish between ICP waves created by the same type of input ABP waveform, but belonging to a different selected correlation type?*. In other words, is the difference between selected correlations dominating the influence of the input ABP waveform? The idea was to project together identified ICP clusters one-by-one, from both selected correlations, and see (1) if these clusters are at all separable and, if not completely, (2) what is the rate of pulses from each selected correlation that build up those additional clusters.

An example of such a projection is depicted in Fig. S5, where the 867 **scp** ICP pulse cluster and the 241 **nsc** ICP pulse cluster are both predominantly ($99.7\ \%$ and $99.6\ \%$) created by ABP waveform 1. A GMM with 5 components was able to very clearly separate the clusters. The prototype components (cluster means) are plotted in Fig. S6.

Further plotting the projections of the cluster components onto the first two principal components (PCs) revealed, without exception, that all components were separable in the PC subspace using only GMM clustering. The complete results of this pair-wise analysis are presented in Table 3.

3.3 Classification

The data arranged in Table 3 was further exploited as a dataset for training and testing the FCN classifier described in Section 2.6.1. The six types of ICP pulses (Fig. 6b) that are assigned to regular waveform, heart failure and reflective input ABP waveforms, and which were separated by their selected correlation type

Table 3: Additional clustering results of pair-wise (scp-nsc) selected ICP clusters corresponding to chosen ABP waveforms.

ABP WF	scp cluster	nsc cluster	Additional clustering			
			Algorithm	Size	scp (%)	nsc (%)
regular	867	241	GMM(5)	241	0	100
				178	100	0
				121	100	0
				242	100	0
				326	100	0
		165	GMM(3)	165	0	100
				416	100	0
				451	100	0
	447	241	GMM(2)	241	0	100
				447	100	0
		165	GMM(2)	165	0	100
				447	100	0
heart failure	1547	301	GMM(4)	169	1.2	98.8
				124	2.4	97.6
				335	99.4	0.6
				1220	99.1	0.9
reflections	216	94	GMM(2)	216	100	0
				94	0	100
	1610		GMM(4)	94	0	100
				775	100	0
				534	100	0
				301	100	0

come naturally as class labels for the training data (Table 4). Even though the dataset is highly imbalanced, after fitting the FCN on the training set ($80\ \%$ of the dataset) it reached an average accuracy of $96\ \%$ on the hold-out test set.

With the purpose of understanding what determined the model to classify each type of pulses, a part of the test set was subjected to the CAMs technique (see Figs. S7 and S8), as described in Section 2.6.1. One common characteristic among the majority of these CAMs is the weight with which the initial slope of the pulse

Table 4: Training dataset consisting of six classes, divided into **scp** and **nsc**, additionally split into regular, heart failure and multiple reflections input ABP waveforms.

Selected correlation type	input ABP-WF	Nr. of ICP pulses
	regular	2181
scp	heart failure	1555
	reflections	1826
	regular	406
nsc	heart failure	293
	reflections	95

contributes to its classification. This phenomenon is especially striking in the **scp-common** class (see first panel in Fig. S7), where the majority of the contribution is focused right at the beginning of the pulse sample, since one would expect that the label will be assigned based on the ripples recorded in $150 - 500$ segment.

One definite conclusion that can be drawn from these results it that the ABP waveforms, viewed as an input to the ICP, definitely introduce detectible changes in the shape of ICP waves. However, these alterations are not as substantial as to smear out the differences between nsc and scp. Moreover, the handful of techniques applied, namely PCA for dimensionality reduction and GMMs with HDBSCAN for clustering, proved to be very accurate and practical in uncovering patterns and correlations in the ICP and ABP pulsatile data never observed before.

4 Conclusions

Starting from a simple time series preprocessing algorithm developed by Achhammer [1], we established a procedure for analysis and discrimination of ICP waves by means of PCA decomposition, concomitant dimensionality reduction and clustering. This approach was applied to seven TBI patients' ICP data, uncovering

various pulse modes previously not mentioned in the literature. These modes can now be further analysed by medical experts in order to explain the states and phenomena they represent.

We hypothesized that as the TBI is exclusively affecting the brain, the ICP pulses describing brain swelling must be independent of the shape of ABP pulses. With that in mind, we proved that ABP waveforms indeed introduce significant changes to the shape of ICP waves, yet these alterations are not as severe as to mask the indications for a brain swelling. These findings provide a proof of concept that could serve as a starting point to further develop this type of analysis and apply to more individuals and more pathologies. Undoubtedly this will result in an increasingly higher accuracy of a potential online-diagnosis system built on the methodology developed in this study.

Even though we succeeded in detecting and identifying several hitherto unknown types of ICP and ABP pulses, we still lack decent knowledge about their interpretation and their significance in a patient-oriented treatment. Our entire methodology can be drastically improved in further studies, in order to arrive at an end-to-end mobile application that can be included in a monitor or alert system at the ICU. The hope is that medical institutions will collaborate more intensely and, more importantly, share multivariate data that will help training models to identify and discriminate between different states with a higher accuracy and a higher confidence.

References

[1] Thomas Achhammer. “An Empirical Mode Decomposition to investigate Intracranial Pressure Profiles to distinguish brain disease states”. In: (2017).

[2] M. Balestreri, Marek Czosnyka, L. A. Steiner, E. Schmidt, P. Smielewski, B. Matta, J. D. Pickard, C. S. Robertson, L. T. Dunn, and I. R. Chambers. “Intracranial hypertension: What additional information can be derived from ICP waveform after head injury?” In: *Acta Neurochirurgica* (2004).

[3] Christopher M. Bishop. *Pattern Recoginiton and Machine Learning*. Springer-Verlag New York, 2006.

[4] Ricardo J.G.B. Campello, Davoud Moulavi, and Joerg Sander. “Density-based clustering based on hierarchical density estimates”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 7819 LNAI.PART 2 (2013), pp. 160–172.

[5] E. R. Cardoso, J. O. Rowan, and S. Galbraith. “Analysis of the cerebrospinal fluid pulse wave in intracranial pressure”. In: *Journal of Neurosurgery* 59.5 (1983), pp. 817–821.

[6] François Chollet. *Keras*. 2015.

[7] Scott J. Denardo, Ramavathi Nandyala, Gregory L. Freeman, Gary L. Pierce, and Wilmer W. Nichols. “Pulse wave analysis of the aortic pressure waveform in severe left ventricular systolic dysfunction”. In: *Circulation: Heart Failure* 3.1 (2010), pp. 149–156.

[8] Celeste Dias, Isabel Maia, António Cerejo, Georgios Varsos, Peter Smielewski, José-Artur Paiva, and Marek Czosnyka. “Pressures, Flow, and Brain Oxygenation During Plateau Waves of Intracranial Pressure”. In: *Neurocritical Care* 21.1 (Aug. 2014), pp. 124–132.

[9] Rupert Faltermeier, Martin A Proescholdt, Sylvia Bele, and Alexander Brawanski. “Parameter Optimization for Selected Correlation Analysis of Intracranial Pathophysiology”. In: *Computational and Mathematical Methods in Medicine* 2015 (2015).

[10] Rupert Faltermeier, Martin A Proescholdt, Sylvia Bele, and Alexander Brawanski. “Windowed multitaper correlation analysis of multimodal brain monitoring parameters”. In: *Computational and Mathematical Methods in Medicine* 2015 (2015).

[11] Xiao Hu, Thomas Glenn, Fabien Scalzo, Marvin Bergsneider, Chris Sarkiss, Neil Martin, and Paul Vespa. “Intracranial pressure pulse morphological features improved detection of decreased cerebral blood flow”. In: *Physiological Measurement* (2010).

[12] Xiao Hu, Peng Xu, Fabien Scalzo, Paul Vespa, and Marvin Bergsneider. “Morphological clustering and analysis of continuous intracranial pressure”. In: *IEEE Transactions on Biomedical Engineering* (2009).

[13] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre Alain Muller. “Deep learning for time series classification: a review”. In: *Data Mining and Knowledge Discovery* 33.4 (2019), pp. 917–963. arXiv: 1809.04356.

[14] Bryan Jennett and Michael Bond. “Assessment of outcome after severe brain damage. A Practical Scale”. In: *The Lancet* (1975).

[15] I T Jolliffe. “Principal Component Analysis, Second Edition”. In: *Encyclopedia of Statistics in Behavioral Science* (2002).

[16] Andreas Jung. “Statistical analysis of biomedical data”. In: (2003).

[17] Yann Lecun. *Generalization and network design strategies*. Ed. by R Pfeifer, Z Schreter, F Fogelman, and L Steels. Zurich, Switzerland, 1989.

[18] Min Lin, Qiang Chen, and Shuicheng Yan. “Network In Network”. In: (Dec. 2013). arXiv: 1312.4400.

[19] Jonathan Long, Evan Shelhamer, and Trevor Darrell. “Fully Convolutional Networks for Semantic Segmentation”. In: (Nov. 2014). arXiv: 1411.4038.

[20] Leland McInnes and John Healy. “Accelerated Hierarchical Density Based Clustering”. In: *IEEE International Conference on Data Mining Workshops, ICDMW* 2017-Novem (2017), pp. 33–42. arXiv: 1705.07321.

[21] Mahesh Nirmalan and Paul M. Dark. “Broader applications of arterial pressure wave form analysis”. In: *Continuing Education in Anaesthesia, Critical Care and Pain* 14.6 (2014), pp. 285–290.

[22] Carlotta Ginevra Nucci, Pasquale De Bonis, Annunziato Mangiola, Pietro Santini, Marco Sciandrone, Arnaldo Risi, and Carmelo Anile. “Intracranial pressure wave morphological classification: automated analysis and clinical validation”. In: *Acta Neurochirurgica* 158.3 (2016), pp. 581–588.

[23] Martin A. Proescholdt, Rupert Faltermeier, Sylvia Bele, and Alexander Brawanski. “Detection of impaired cerebral autoregulation using selected correlation analysis: A validation study”. In: *Computational and Mathematical Methods in Medicine* 2017.February (2017).

[24] Peter L. Reilly. “Brain injury: The pathophysiology of the first hours. ’Talk and Die revisited’”. In: *Journal of Clinical Neuroscience* 8.5 (2001), pp. 398–403.

[25] L. A. Steiner and P. J.D. D Andrews. “Monitoring the injured brain: ICP and CBF”. In: *British Journal of Anaesthesia* 97.1 (2006), pp. 26–38.

[26] Mark E. Wagshul, Per K. Eide, and Joseph R. Madsen. “The pulsating brain: A review of experimental and clinical studies of intracranial pulsatility”. In: *Fluids and Barriers of the CNS* 8.1 (2011), pp. 1–23.

[27] Zhiguang Wang, Weizhong Yan, and Tim Oates. “Time Series Classification from Scratch with Deep Neural Networks: A Strong Baseline”. In: *CoRR* abs/1611.0 (2016). arXiv: `1611.06455`.

[28] Angela Zeiler. “Weighted Sliding Empirical Mode Decomposition and its Application to Neuromonitoring Data”. In: (2012).

[29] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. *Learning Deep Features for Discriminative Localization*. 2015. arXiv: `1512.04150 [cs.CV]`.