IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received August 17, 2021, accepted September 9, 2021, date of publication September 16, 2021,
date of current version September 28, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3113055

Deep-Learning Approach for Tissue Classification Using Acoustic Waves During Ablation With an Er:YAG Laser

CARLO SEPPI[1,2], ANTAL HUCK[1,2], HERVÉ NGUENDON KENHAGHO[1,3], EVA SCHNIDER[1,2], GEORG RAUTER[1,4], (Member, IEEE), AZHAR ZAM[1,3], (Member, IEEE), AND PHILIPPE C. CATTIN[1,2], (Member, IEEE),

[1]Department of Biomedical Engineering, University of Basel, 4123 Allschwil, Switzerland
[2]Center for medical Image Analysis & Navigation (CIAN), 4123 Allschwil, Switzerland
[3]Biomedical Laser and Optics Group (BLOG), 4123 Allschwil, Switzerland
[4]Bio-Inspired Robots for Medicine-Laboratory (BIROMED-Lab), 4123 Allschwil, Switzerland

Corresponding author: Carlo Seppi (carlo.seppi@unibas.ch)

This work was supported by Werner Siemens Foundation through the Minimally Invasive Robot-Assisted Computer-Guided LaserosteotomE (MIRACLE) Project.

ABSTRACT Today's mechanical tools for bone cutting (osteotomy) lead to mechanical trauma that prolong the healing process. Medical device manufacturers continuously strive to improve their tools to minimize such trauma. One example of such a new tool and procedure is minimally invasive surgery with laser as the cutting element. This setup allows for tissue ablation using laser light instead of mechanical tools, which reduces the post-surgery healing time. During surgery, a reliable feedback system is crucial to avoid collateral damage to the surrounding tissues. Therefore, we propose a tissue classification method that analyzes the acoustic waves produced during laser ablation and show its applicability in an ex-vivo experiment. The ablation process with a microsecond pulsed Erbium-doped Yttrium Aluminium Garnet (Er:YAG) laser produces acoustic waves that we captured with an air-coupled transducer. Consequently, we used these captured waves to classify five porcine tissue types: hard bone, soft bone, muscle, fat, and skin tissue. For automated tissue classification of the measured acoustic waves, we propose three Neural Network (NN) approaches: A Fully-connected Neural Network (FcNN), a one-dimensional Convolutional Neural Network (CNN), and a Recurrent Neural Network (RNN). The time- and the frequency-dependent parts of the measured waves' pressure variation were used as separate inputs to train and validate the designed NNs. In a final step, we used Grad-CAM to find the frequencies' activation map and conclude that the low frequencies are the most important ones for this classification task. In our experiments, we achieved an accuracy of 100 % for the five tissue types for all the proposed NNs. We tested the different classifiers for their robustness and concluded that using frequency-dependent data together with a FcNN is the most robust approach.

INDEX TERMS Acoustic feedback, laser ablation, tissue classification, neural network.

I. INTRODUCTION

Minimally invasive procedures demonstrate a significant step towards accelerated recovery after surgery [1], [2]: replacing the mechanical tools from open osteotomies with laser-based ablation [3] shows a further reduction in recovery time [4], [5]. Mechanical tools – which are still the standard in conventional osteotomy – induce thermal and mechanical trauma due to mechanical friction. Replacing mechanical tools with lasers can reduce this trauma [3], [6], [7].

While exposing the tissue to a microsecond pulsed Er:YAG laser, the water in the tissue is heated until it vaporizes. This process takes place within microseconds and builds up pressure that is released in a series of micro-explosions. The explosions ablate a small portion of the tissue [8] and thus produce an acoustic wave [9]. A transducer can then measure the resulting pressure variation. Carbonization causes thermal damage and reduces the cutting efficiency. Therefore,

The associate editor coordinating the review of this manuscript and approving it for publication was Salvatore Surdo.

the ablated tissue needs to be constantly re-hydrated and cooled down [3], [6], [7], [10]–[12].

The goal of the recently founded project MIRACLE[1] is to improve laser osteotomy, by integrating the advantages of robot-assisted laser surgery into an endoscope [13]–[15]. This way, the surgeon may perform laser-based osteotomy by inserting an endoscope into the body through a small incision or a natural orifice. Information on the endoscope's surroundings, e.g., the type of the ablated tissue, can help the surgeon to avoid cutting the wrong tissue. Multiple approaches have been considered to discern porcine tissue types as feedback for tissue ablation with laser, e.g., using Optical Spectroscopy [16]–[18] or Optical Coherence Tomography (OCT) [19], [20]. The authors of [21] proposed that optoacoustic imaging can be used for differentiating different types of hard dental tissue. The authors of [22], [23] have investigated optical and acoustic signals during Er:YAG laser osteotomy. They have proposed a heuristic that decides when the laser needs to be switched off to prevent damaging nerves. In contrast, our goal is to only use the acoustic signal for tissue classification to prevent the laser from continuing the cut when detecting tissue that should not be damaged. Similar approaches have been proposed in [11], [24], [25]. A different approach [26], [27] used acoustic waves in a 2D simulation to infer the acoustic density within a region of interest. This information subsequently can be used to classify the underlying tissue.

For this research, we used supervised deep learning to train NNs that can infer the tissue from the acoustic waves. In a simplified ex-vivo experiment, we prepared specimen from porcine tissues such that each tissue type can be ablated without the interference of others.

Neural Networks [28] found their way into numerous related applications, e.g., medical image classification problems [29]–[31] or speech and signal processing [32]–[35]. Similarly, we used NNs to classify different porcine-tissue types. In our case, we used the pressure variation emitted during the ablation to feed our classifying NNs. We either used the time-dependent pressure variation or its Fourier-transform, but we used the same network architecture in both cases. We compared an FcNN of three fully-connected (FC) layers, a one-dimensional CNN with one convolutional layer followed by three FC layers, and a bidirectional RNN followed by three FC layers. To further analyze the frequency domain, we applied Grad-CAM to find activation maps. Grad-CAM [36] can be used to compute the activation maps that highlight the essential part of the data for a specific classification task. Therefore, we applied Grad-CAM to our proposed CNN with frequency-dependent data and found the frequency domain's corresponding activation map.

The proposed new approach showed superior results when compared to the original method [25] on the same data.

[1]MIRACLE (Minimally Invasive Robot-Assisted Computer-guided LaserosteotomE), *01.08.2020*, https://dbe.unibas.ch/en/research/flagship-project-miracle

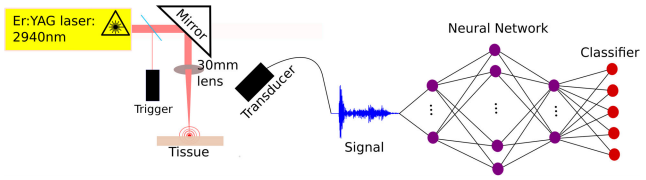

FIGURE 1. Setup contained an Er:YAG laser with a wavelength of 2940 nm, where the pulses had a repetition rate of 2 Hz and the energy per pulse was 940 mJ. The tissue was placed at the laser's focal point, at a distance of 30 mm from the lens. The acoustic wave created during the ablation was measured by a transducer at a distance of 5 cm, with an angle of 45°. We used the measured signal as the input of our NNs to classify different types of tissues.

Furthermore, we performed a robustness analysis of the different NNs, compared the performance on time- and frequency-dependent data, and finally discussed our results.

II. MATERIAL

In this section, we describe the setup and the data acquisition of our experiments performed in [25], to which we also refer for more details.

A. SETUP

Figure 1 visualizes the setup used in this research. We used an Er:YAG laser (Syneron Candela, litetouch LI-FG0001A) with a wavelength of 2940 nm that produces 400 μs pulses with an energy of 940 mJ. A CaF_2 mirror was placed at a small angle in front of the laser's head, such that it splits the laser light into two parts: 96% transmitted and 4% reflected. The reflected light is captured by a fast PbSe photodiode (PbSe Fixed Gain Detector, PDA20H, 1500 − 4800 nm), used as a triggering signal. This triggering signal activates the measurement of the acoustic signal received by the transducer. The custom-made air-coupled piezoelectric transducer[2] with a diameter of 15 mm, a frequency range of 0.1 MHz − 0.8 MHz, and the resonance frequency at 0.4 MHz captured the acoustic waves produced during ablation and records it in a 0.82 ms time window. The experiment was performed in wet conditions, using a distilled water spray with a flow rate of 0.1 ml s^{-1}, which reduces carbonization during ablation. This transducer was placed at a distance of 5 cm at an angle of 45° to the specimen. The transducer converted the measured pressure variation into a digital signal with a sampling rate of 10 MHz. The measured data was then used as input to our proposed tissue classifier.

B. DATA

The data was obtained from ablating fresh specimens, namely, hard bone (compact bone fragment), soft bone (spongious bone), muscle, fat, and skin tissue, with a size of $10 \times 50 \times 5\,mm^3$. All the tissues were rinsed in distilled water before the experiment was performed. The tissues were dissected from five porcine proximal and distal femurs, which were bought on different days. The laser then ablated the

[2]Provided by Tomas E. Gomez Alvarez-Arenas in the ITEFI-Instituto de Tecnologias Fisicas y de la Information, CSIC, Madrid, Spain.

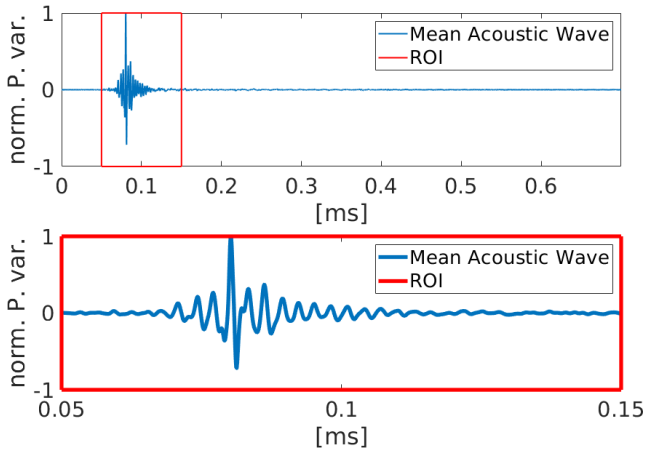

FIGURE 2. **Normalized mean value of all the acoustic wave f. The** 0.1 ms **region of interest (ROI) is marked red.**

tissue samples with the following setup: Of each tissue type, 10 specimens were probed, where each specimen was ablated 10 times, producing a vertical hole with respect to the bone surface. The individual holes were spaced at least 4 mm apart from each other. Each of those laser ablations consisted of 180 laser pulses with a repetition rate of 2 Hz. Consequently, 1800 measurements of ablation-induced pressure variations were made for each specimen, giving a total of 18 000 measurements per tissue type. The examination of five different tissue types results in an overall of 90 000 measurements. The data acquisition window for each acoustic wave was 0.82 ms. Taking the normalized mean value of all acoustic waves, we conclude that the primary information of the wave was in a 0.1 ms time window, as is visualized in Figure 2. In the proposed approach, we used that time window of size 0.1 ms; and with a frame rate of 10 MHz each acoustic wave was therefore represented as a 1000×1 -dimensional array.

III. METHODS

Our goal is to interpret the acoustic waves produced during ablation to classify the ablated tissue, with an end-to-end neural network approach. We use three network designs in both the time and the frequency domain, namely an FcNN, a CNN, and an RNN. The FcNN (top of Figure 3) consists of two hidden FC layers with 1000 neurons each. Each hidden layer is followed by a *ReLU* activation function. The neurons on the final FC layer correspond to the number of different tissues for the classification task, i.e. five. The layers of the CNN (middle part of Figure 3), consist of one convolutional layer and three FC layers. The convolutional layer has 6 output channels, followed by a *ReLU* activation function and a subsequent *Maxpool* layer with a kernel size of 2. The following FC layer maps the output to one channel of size 1000. Again, *ReLU* serves as an activation function. The RNN (bottom of Figure 3), consists of a bidirectional RNN layer, where the number of features in the hidden state corresponds to the input size with the *Tanh* activation function. Again, it is followed by three FC layers, where

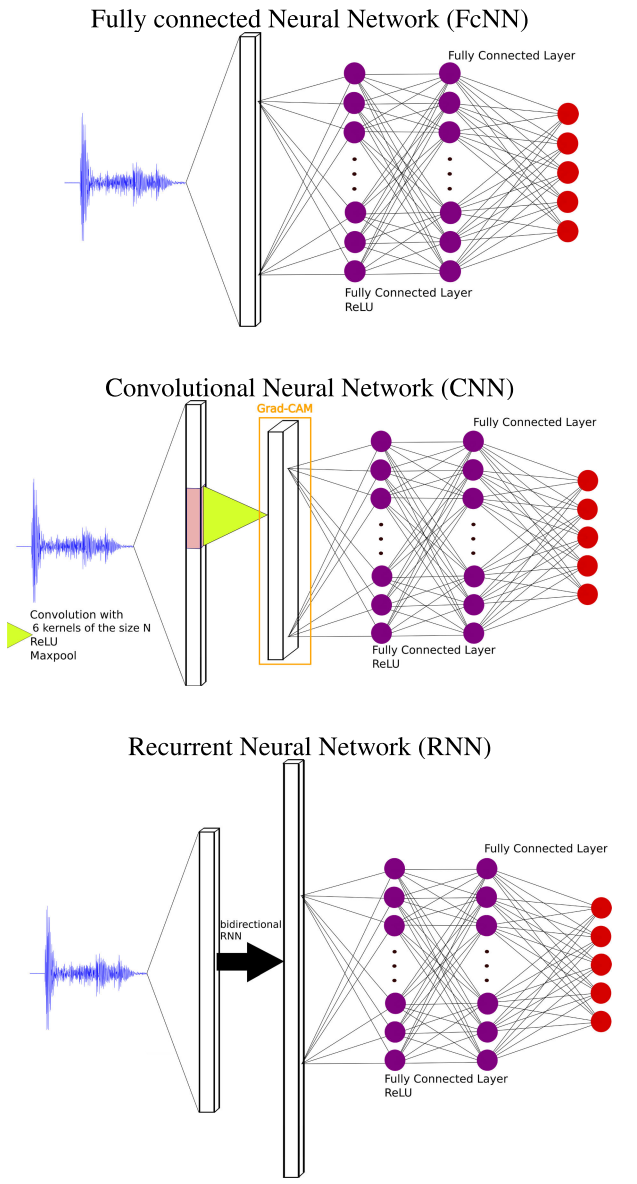

FIGURE 3. ***Top*: Visualization of the FcNN. We have two hidden 1000-way FC layers, each followed by a *ReLU* activation function. The final FC layer has 5 neurons, which coincides with the 5 types of tissue in the classification problem. *Center*: Visualization of the CNN. The input is the measured pressure variation in the time domain or the frequency domain. First, a convolutional layer is applied, followed by a *ReLU* activation function and a subsequent *Maxpool* layer. The CNN's kernel size with the time-dependent data is $N_t = 200$, and the kernel size of the frequency-dependent data is set to $N_f = 2$. The convolutions are followed by two 1000-way FC layers with *ReLU* activation functions. Since we classify 5 types of tissue, the final FC layer has an output dimension of 5. By applying Grad-CAM to the CNN's convolutional layer with the frequency-dependent data, we obtain the activation map in the frequency domain, which provides an evaluation of the influence of each frequency for classification. *Bottom*: Visualization of the RNN. The first hidden layer is a bidirectional RNN layer, where the number of features in the hidden state corresponds to the input size with the activation function *Tanh*. It is followed by two hidden 1000-way FC layers with *ReLU* activation function. Again, the final FC layer has the output dimension of 5.**

the two hidden FC layers have 1000 neurons each and are followed by a *ReLU* activation function. The final layer has 5 neurons, corresponding to the number of different tissues in the classification task. We choose *Adam* [37] as an optimizer

for the training phase with a learning rate of 10^{-3} and a batch size of 16 in combination with the *Cross-Entropy* loss.

We divided the data set into three disjoint subsets: 20 % (2 specimens of each tissue type) were used for training, 20 % (2 specimens of each tissue type) for validation, and the remaining 60 % (6 specimens of each tissue type) as test data. After each training epoch, i.e., an iteration of all training data used to train the network, the algorithm's performance was estimated on the validation data. To detect overfitting and ensure our network's robustness, we exclusively used previously unseen data for testing, i.e., measurements from a given specimen were only used in one subset. To evaluate the robustness and variability of our approach, we conducted five-fold cross-validation. To this end, we split the data into five disjoint subsets, e.g. $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$, where each subset contains 2 specimens of each tissue type. The first network used subset $\mathcal{A}$ for training, $\mathcal{B}$ for validation, and $\mathcal{C} \cup \mathcal{D} \cup \mathcal{E}$ for testing. Note that with this construction, the training, validation, and test data sets were all disjoint. The second network used subset $\mathcal{B}$ for training, $\mathcal{C}$ for validation, and $\mathcal{D} \cup \mathcal{E} \cup \mathcal{A}$ for testing. This continues in a rotating fashion until the fifth network used $\mathcal{E}$ for training, $\mathcal{A}$ for validation, and $\mathcal{B} \cup \mathcal{C} \cup \mathcal{D}$ for testing.

We note that we used only 20% of the data for training, i.e. two specimens of each tissue type. The maximum accuracy of the network is achieved with little data in all folds of the cross-validation experiment. Using more training data does not improve the performance with respect to the accuracy.

A. TIME-DEPENDENT DATA

In the left column of Figure 4, we visualize the pressure variation of exemplary measurements. As demonstrated in this figure, the absolute values between the different tissues may vary drastically. Therefore, we apply a Hamming window [38] and normalize the resulting data (dividing it by the maximum of the absolute value). Hence, the normalized pressure variation (see Figure 4, middle column) varies between -1 and 1. The resulting size of the array of the preprocessed measurements remains at 1000×1. For the CNN with time-dependent data, we choose a convolutional layer (bottom part of Figure 3) with a kernel size of $N_t = 200$ and a padding size of 0. This reduces the trainable parameters at the transition from the convolutional layer to the FC layer from $6 \cdot 500 \cdot 1000$ to $6 \cdot 400 \cdot 1000$. The number of parameters of the NN are 2 007 005 for the FcNN, 3 408 211 for the CNN, and 7 011 005 for the RNN.

B. FREQUENCY-DEPENDENT DATA

To transform the time-dependent data into the frequency domain, we perform the following steps: First, we apply a Hamming window to reduce the leakage in the Fast Fourier Transformation (FFT) [38]. Then, we normalize the resulting data so that the magnitudes vary between -1 and 1 (middle column of Figure 4) and apply the FFT to the normalized data. The experiments conducted in [25] were measured with an air-coupled piezoelectric transducer, which was limited to the frequency between 0.1 MHz and 0.8 MHz. Therefore, we only use the spectrum's magnitudes of this given range as the input of the NNs (right column of Figure 4). In a final step, we scale the resulting pressure variation of the frequency domain with a factor $\alpha = 6$, such that the maximum magnitudes of the input data were close to 1. The advantage of using a constant α, instead of normalizing the pressure variation, is to enable the comparison of the activation maps of the different frequency subsets (top left of Figure 6). The frequency-dependent data has a much smaller input array of the size of 70×1 in comparison to the size of the time-dependent data of 1000×1. Therefore, we choose a smaller kernel size of $N_f = 2$ for the convolutional layer of the CNN. The number of parameters of the NN is 1 077 005 for the FcNN, 1 211 023 for the CNN, and 1 166 885 for the RNN.

C. GRAD-CAM

Selvaraju *et al.* [36] have introduced a gradient-based localization, called Grad-CAM. They proposed that the convolutional layers' gradients highlight the parts needed for classification; we refer to it as the *activation map*. Since Grad-CAM requires a convolutional layer, we can only apply it to the proposed CNN but not to the FcNN and the RNN. Applying Grad-CAM to the frequency-dependent data enables us to highlight the essential frequencies of our classifier. In particular, the higher the activation of a frequency, the more important this frequency is for the network's classification task. We apply Grad-CAM to the convolutional layer of the trained CNN with the frequency-dependent data, as depicted at the bottom of Figure 3.

To compute the activation map, Grad-CAM needs to compute the gradient of the convolution layer. To this end, we compute the CNN with a given measurement. Then we apply back-propagation to find the gradients. Since our convolution layer has 6 output channels, our gradient has 6 channels as well. First, we apply the *ReLU* activation function to highlight the gradient's positive impact, and then took the gradients' sum. Finally, we compute the mean value over all training data and all cross-validations to find the most important frequencies for all the tissues. Because of the dimension reduction after the first layer, we interpolate the resulting vector to the same length as the initial frequency. Ultimately, we normalize the vector by its maximum value.

D. ROBUSTNESS

We analyzed the robustness of the NNs by augmenting the data at test time and measuring the resulting accuracy.

First, we assume that the angle between the microphone and the tissue is fixed. Therefore, varying distances lead to a shift in the time frame. We augment such shifts by transposing the measured acoustic wave window between -10 to 10 frames for the time-dependent data and -150 to 150 for the frequency-dependent data. Assuming that the speed of sound is 343 m s^{-1} and the measured frame rate was 10 MHz, shifts of 10 and 150 frames correspond to transitions between the microphone and the specimen of approximately 0.03 cm

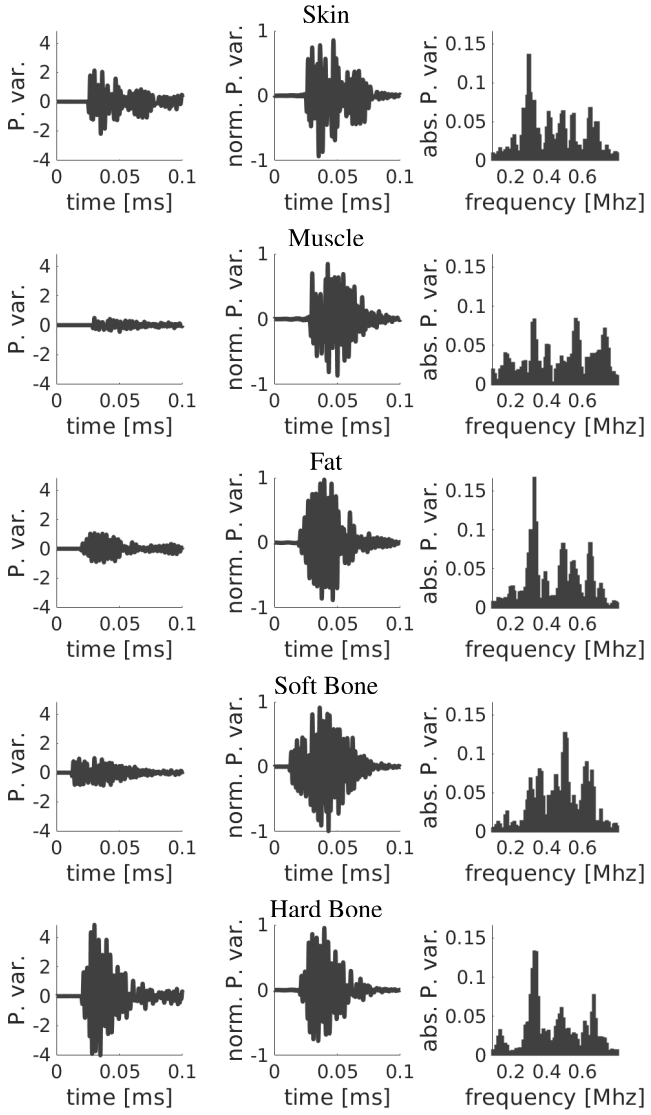

FIGURE 4. Five exemplary measurements and their preprocessed versions as inputs for the NNs. ***Left column:*** Measured pressure variation (P. var.), with a time window of $0.1\,\mathrm{ms}$. ***Middle column:*** First Hamming window and then normalization are applied to the inputs of the NNs with time-dependent data. ***Bottom column:*** Absolute value of the FFT of the Hamming-normalized pressure variation (input for the NNs with frequency-dependent data).

and 0.5 cm, respectively. We note that a much smaller time window for the time-dependent NNs is required as they prove to be less robust than the frequency-dependent NNs – see Section IV.

Second, changing the distance of the microphone to the ablation point, in theory, will change the magnitude of the measured pressure variation. Since within our method we normalize with respect to the absolute maximum value of the measured pressure variation, a linear scaling will have no effect. Therefore, we apply the nonlinear scaling:

$$p^* = p \cdot (1 + \beta \cdot \exp(-|p|)), \tag{1}$$

where p is the measured pressure variation, β a value between -1 and 1, and p^* the augmented pressure variation.

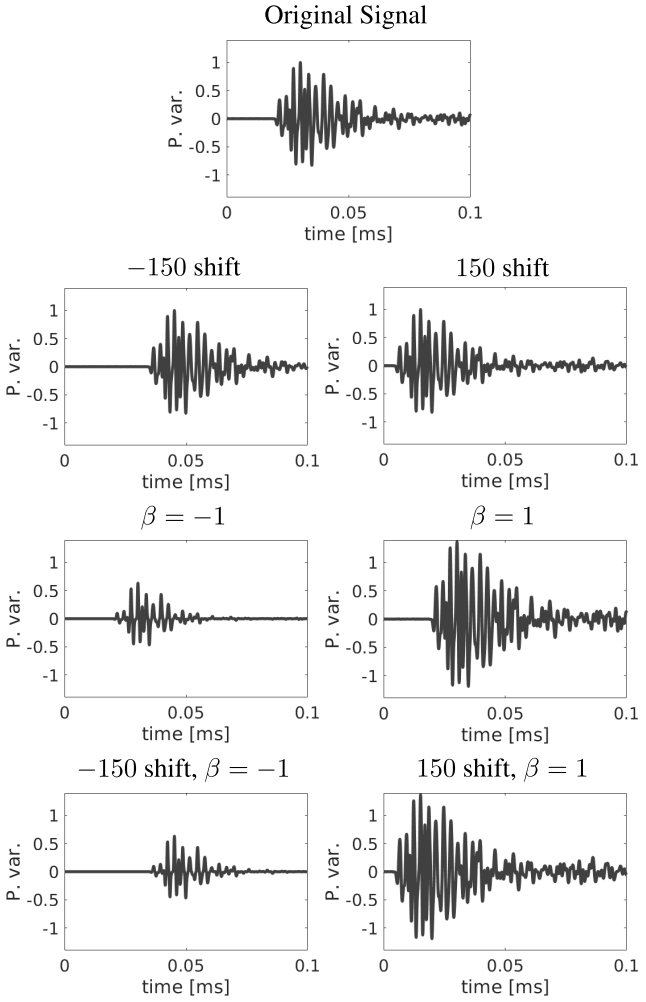

FIGURE 5. Augmented shift of the distance between the microphones and the tissue of approximately $\pm 0.5\,\mathrm{cm}$ on an exemplary acoustic signal. ***Top to bottom***: example of a shift of ±150 frames and nonlinear pressure variance amplification using the equation 1 with $\beta = \pm 1$. ***Bottom:*** combination of the frameshift and the amplification.

An exemplary visualization of the time shift, the nonlinear magnitude variation, and a possible shift of $\pm 0.5\,\mathrm{cm}$ are visualized in Figure 5.

We compare the robustness score of the different networks by evaluating the network on a subset of the test data and compare the networks' mean accuracy resulting from the various time shifts and β.

IV. RESULTS

We implemented the networks in *Python (3.6.9)* using *PyTorch (1.5.1)* [39] and trained them until they reached an accuracy of 100 % on the validation data, meaning, all validation samples were assigned to the correct tissue class. In addition, we applied five-fold cross-validation, where we permuted the training, validation, and test data as described in the previous section. In our cross-validation experiments, the NNs with the time-dependent data had to be trained between 14 and 42 epochs, and the NNs with the frequency-dependent data between 22 and 98 epochs. In a

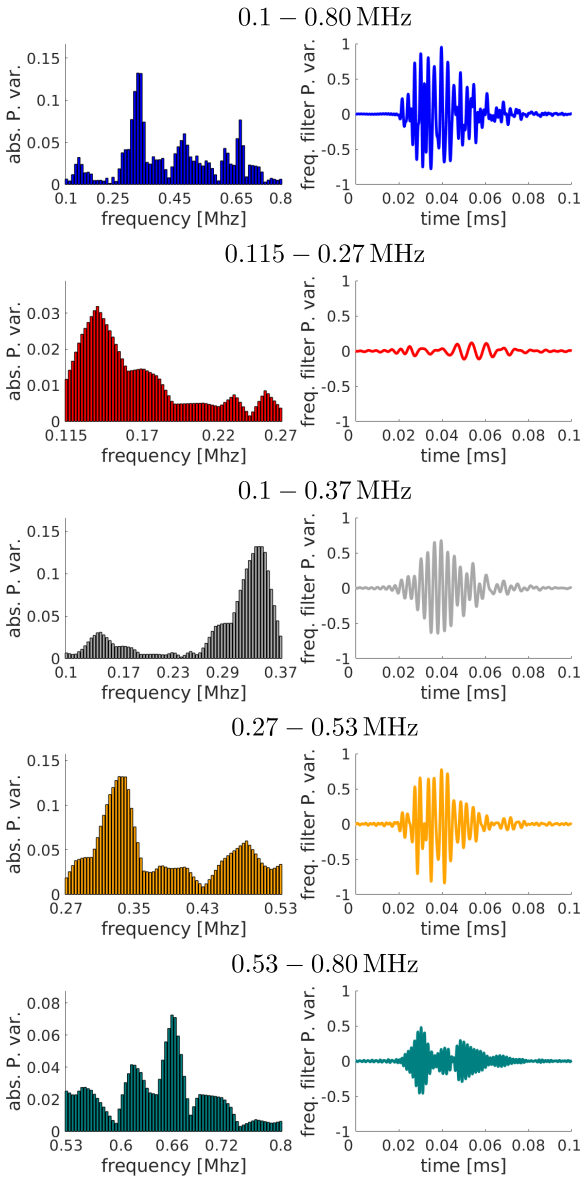

FIGURE 6. ***Top to bottom*: An exemplary ablation signal of a hard bone sample with six subsets of frequency ranges. *Left*: Exemplary data used as input for training the CNN with frequency-dependent data. *Right*: Frequency filtered (inverse FFT) pressure variation (P. var.) used as input of the CNN with time-dependent data.**

final step, we used the previously unseen 54 000 measurements (60 %, 6 specimens for each tissue type) to test our network. This allowed assessment of the robustness and the generalization capabilities of the network on unseen data. For both the time- and the frequency-dependent data, we achieved a classification accuracy of 100 % on the test data.

We compare our results to those from [25], which are based on the same data and which proposes three methods: a Quadratic-SVM, Gaussian-SVM, and a shallow artificial neural network (ANN) with one single hidden layer consisting of 10 neurons. For all those methods, first, a principal component analysis (PCA) was applied to the Fourier spectrum and subsequently, the scores of the first three principal components were used for further processing. The ANN

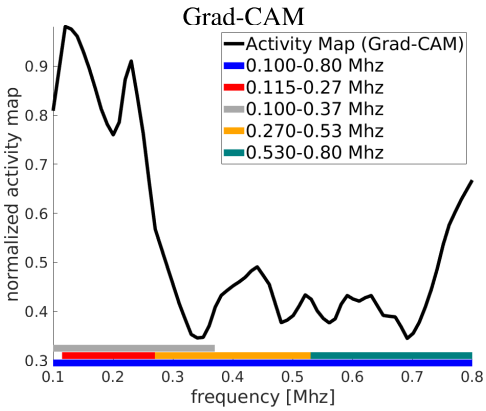

FIGURE 7. **The activation map, acquired with Grad-CAM, is represented by the mean values over all cross-validation networks. It suggests that the low frequencies were the most important ones for our classification task.**

approach performed best, and therefore we limit our analysis to the comparison between the ANN and our networks.

Similar to [25], we divided the frequency range into five different subranges: 0.1 MHz − 0.8 MHz, 0.115 MHz − 0.27 MHz, 0.1 MHz − 0.37 MHz, 0.27 MHz − 0.53 MHz, and 0.53 MHz − 0.8 MHz. Reducing the frequency range, also reduces the length of the input array. Therefore, we interpolated the input arrays such that they matched the input sizes of the NNs with the frequency-dependent data (70×1), compare the left column of Figure 6. For the time-dependent data, we applied the inverse FFT to the normalized pressure variation as a frequency filter, solely using the given subranges, as is visualized on the right side of Figure 6.

Also, when training only subranges of frequencies, all the resulting NNs still achieved an accuracy of 100 % on the validation and test data, as it is presented in Table 1. In Table 2, we compare our results to ANN, where the best score was achieved with the frequency subrange of 0.115 MHz − 0.27 MHz with 90.88 %. The bandwidth of 0.1 MHz−0.8 MHz and the subrange of 0.1 MHz−0.37 MHz was not tested by the authors of [25]. We note that they used a time window of 0.82 ms, while we solely used a time window of 0.1 ms. This reduces the dimension of our input for the NNs.

In a final step, we compared the execution time of our NNs to the shallow ANN presented in [25]. The minimal execution time of the ANN approach they have reported to be 11.2 ms. Executed on a system with an *Intel(R) Xeon(R) CPU E5-2680 v4* @ 2.40 GHz and 94 GB 2400 MHz DDR4 memory, our approach had an execution time between 0.8 ms − 1.2 ms, which was approximately ten times faster than the execution time of the shallow ANN (see Table 3).

A. GRAD-CAM

We applied Grad-CAM to the CNN, which was trained on the frequency domain. This revealed the activation map, visualized in Figure 7. The activation map shows that the lower frequencies are more important to the classification process

TABLE 1. Comparison of the five different frequency ranges and the number of epochs needed to reach an accuracy of 100 % on the validation data. In all cases, the networks reach perfect accuracy on the test data. The minimum and the maximum number of epochs needed are shown. The top table represents the results of the FcNN, the center represents the results of the CNN, and the bottom table represents the results of the RNN.

Fully connected Neural Network (FcNN)

Freq. [MHz]	Time-dependent data # of epochs	accuracy	Freq.-dependent data # of epochs	accuracy
0.1 − 0.80	14 − 25	100 %	28 − 29	100 %
0.115 − 0.27	16 − 25	100 %	68 − 73	100 %
0.10 − 0.37	18 − 24	100 %	52 − 54	100 %
0.27 − 0.53	15 − 19	100 %	48 − 50	100 %
0.53 − 0.80	14 − 16	100 %	56 − 59	100 %

Convolutional Neural Network (CNN)

Freq. [MHz]	Time-dependent data # of epochs	accuracy	Freq.-dependent data # of epochs	accuracy
0.1 − 0.80	24 − 29	100 %	35 − 37	100 %
0.115 − 0.27	23 − 37	100 %	94 − 98	100 %
0.10 − 0.37	26 − 35	100 %	62 − 64	100 %
0.27 − 0.53	35 − 46	100 %	42 − 46	100 %
0.53 − 0.80	28 − 40	100 %	70 − 79	100 %

Recurrent Neural Network (RNN)

Freq. [MHz]	Time-dependent data # of epochs	accuracy	Freq.-dependent data # of epochs	accuracy
0.1 − 0.80	27 − 39	100 %	22 − 27	100 %
0.115 − 0.27	24 − 35	100 %	36 − 47	100 %
0.10 − 0.37	24 − 36	100 %	29 − 35	100 %
0.27 − 0.53	26 − 40	100 %	30 − 31	100 %
0.53 − 0.80	25 − 42	100 %	37 − 40	100 %

TABLE 2. Comparison our results from our NNs to the shallow ANN of [25] in different subranges.

[MHz]	FcNN / CNN / RNN Freq.	Time	ANN Freq.
0.1-0.8	100 %	100 %	-
0.115-0.27	100 %	100 %	90.88 %
0.10-0.37	100 %	100 %	-
0.25-0.53	100 %	100 %	85.38 %
0.53-0.80	100 %	100 %	78.85 %

TABLE 3. Comparison of the different execution times of our NNs and the ANN of [25]. The execution times of the FcNNs and CNNs (time- and freuency-dependent input) were calculated as the mean execution times of the forward passes of all the NNs associated to the different subranges. For the ANN, solely the frequency range of 0.115 – 0.27 MHz was measured.

	FcNN		CNN		RNN		ANN
	Freq.	Time	Freq.	Time	Freq.	Time	Freq
[ms]	0.8	0.8	1.0	1.1	1.0	1.2	11.2

than the high frequencies. This supports the results found in [25], where they achieved the highest accuracy within a low-frequency subrange of 0.115 − 0.27 MHz, using a shallow neural network classifier.

Another indication that the activation map highlights the important frequencies is presented in the first epochs in Figure 8: during the first few epochs, the accuracy of the low-frequency subrange of 0.115 − 0.27 MHz exceeded those of the subranges 0.27 − 0.53 MHz and 0.53 − 0.8 MHz. But, the networks need a longer training time for final convergence. A reason for this can be the poor choice of the learning rate.

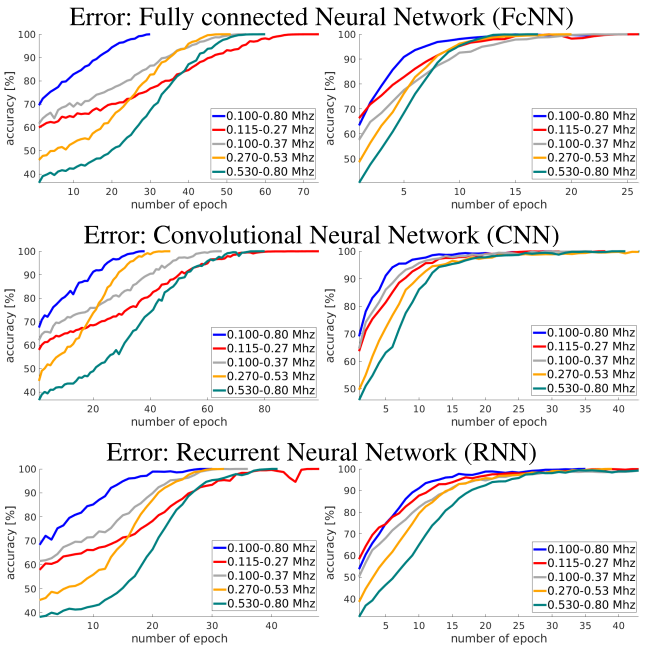

FIGURE 8. Accuracy of the network applied to the validation data of the five different frequency subranges, represented by the mean values over all cross-validation networks. *Top*: FcNN, *center*: CNN, *bottom*: RNN. *Left*: Error of the NNs with frequency-dependent data. *Right*: Error of the NNs with time-dependent data.

B. ROBUSTNESS

We tested the robustness of the frequency and the time-dependent NNs on 5 % of the test data (2700 measurements). To this end, we performed a grid search over β, as described in Equation 1, and the time frameshift. In Figure 9, we visualize for all networks and frequency ranges the accuracy for each grid point.

We see that the time-dependent networks were much more sensitive with respect to the time frameshift. Therefore, we chose, for the frequency-dependent networks, a time shift between ± 150 and for the time-dependent networks a time shift between ± 10. In addition, the frequency-dependent networks are slightly more sensitive to the choice of β, in comparison to the time-dependent networks. For all the networks, we chose β in the range of ± 1. Figure 9 clearly visualizes that the frequency-dependent networks are more robust than the time-dependent networks.

To compare the different frequency ranges, we computed the robustness accuracy – the mean value over all combinations of the time frameshift and β – as is summarized in Table 4. We note that the FcNN and the RNN show similar results and better robustness on the frequency-dependent data. The CNN performs better in the robustness test for the time-dependent data. The highest robustness was achieved when using the whole frequency range (0.1 − 0.8 MHz) of the transducer with the frequency-dependent networks. In terms of network architectures, the highest robustness was achieved by the FcNN with a robustness score of 92.5 %. The RNN network has a similar robustness score of 91.4 %, followed by the CNN network with a robustness score of 85.2 %. For the time-dependent network, all three networks have the highest

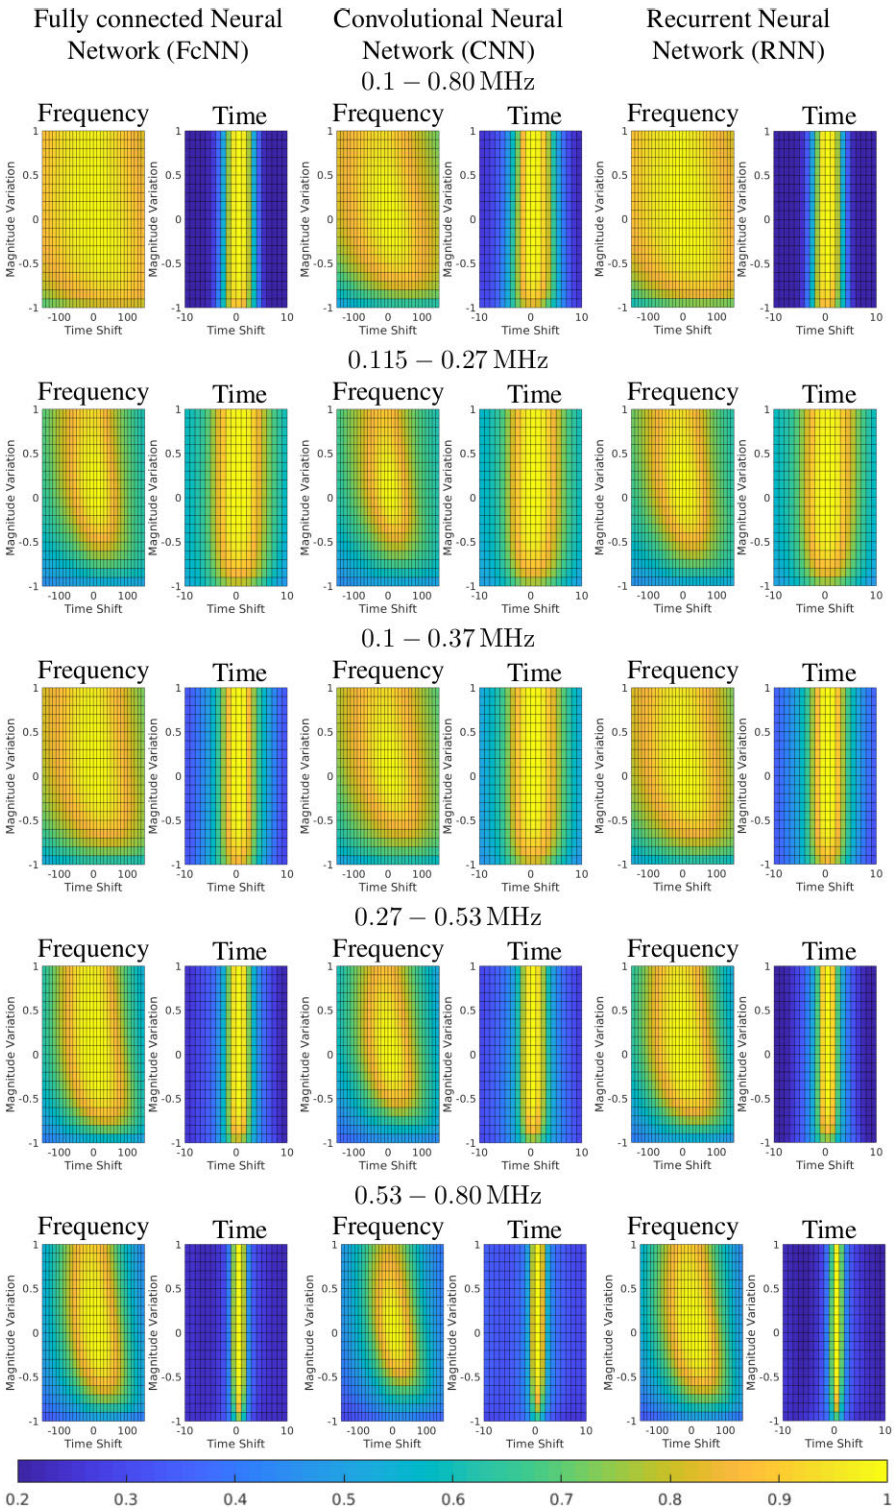

FIGURE 9. **The robustness of the proposed networks. Frequency- and time-dependent of FcNN (*left*), CNN *center*, and RNN *(right)*. The range of the time frameshift of the frequency-dependent data is ±150, and for the time-dependent data ±10. For the data augmentation, we used β in the range of −1 to 1.**

robustness at the frequency range of $0.115 - 0.27$ MHz. All of them have a similar robustness score, where FcNN has 75.8 %, the CNN has 75.3 %, and the RNN has 75.0 %. In all types of networks, those trained on the high-frequency range of $0.53 - 0.8$ MHz show the lowest robustness score.

These observations correspond well with our findings based on the activity map in Figure 7 where we showed that the low-frequency range was of high importance for the tissue classification. Comparing the similar lengths of frequency ranges 0.1−0.37 MHz, 0.27−0.53 MHz, and 0.53−0.8 MHz,

TABLE 4. The robustness expressed as the mean value of the classification accuracy achieved by the NNs with data augmentation at test time.

[MHz]	0.1-0.8	0.115-0.27	0.1-0.37	0.25-0.53	0.53-0.8
Freq.					
FcNN	92.5%	75.4%	84.9%	77.6%	70.1%
CNN	85.2%	73.2%	81.6%	71.2%	62.4%
RNN	91.4%	76.8%	85.0%	76.9%	70.8%
Time					
FcNN	41.9%	75.8%	59.2%	46.8%	34.8%
CNN	54.2%	75.3%	73.5%	50.5%	41.9%
RNN	41.3%	75.0%	60.6%	42.2%	33.1%

also shows that the highest robustness score is achieved by the low-frequency range. In fact, with the time-dependent data, a shorter frequency range of $0.115 - 0.27$ MHz exceeds all other tested frequency ranges.

V. CONCLUSION

Our NNs use the acoustic waves emitted during tissue ablation with a microsecond pulsed Er:YAG laser to classify the ablated tissue. Even though we used fewer training data compared to [25], we substantially improved the classification accuracy of the tissues (hard bone, soft bone, muscle, fat, and skin tissue), again compared to [25], where their machine learning approached achieved a mean accuracy of under 91 %. Our network managed to classify all tissue types with a classification accuracy of 100 % for both approaches: the NNs with time-dependent data, and the NNs with frequency-dependent data. We believe that the methods of [25] are limited to a lower accuracy because of hand-drafted features, i.e., the input data was projected to the scores of the first three principal components, and therefore, important information of the acoustic wave was not utilized.

We used an activation map to find the essential frequency ranges of the acoustic waves and conclude that the low-range frequencies have the highest impact on our network's classification, which coincides with the claim in [25].

Since all of our NNs achieved a classification accuracy of 100 %, we used a robustness test to further analyze their performance. The results imply that the frequency-dependent networks were more robust than the time-dependent networks. Although the FcNN and the RNN had similar robustness scores, the most robust network was the FcNN that used the frequency-dependent data on the whole available frequency range of $0.1 - 0.8$ MHz, unlike the NNs that used the time-dependent data, where the low-frequency range was the most robust method. However, the CNN has a higher robustness score for most time-dependent data in comparison to the FcNN and the RNN. Although the FcNNs, the CNNs, and the RNNs achieve similar results, the FcNNs performed slightly better. We conclude, that after a certain complexity of the network, no significant improvements are gained.

With a time window of 0.1 ms the classification approximately takes 1 ms and, therefore, the method could be used as a real-time classifier at the current laser repetition rate of 2 Hz. In fact, the real-time classifier could still work, even when the laser's repetition rate is increased.

A. FUTURE WORK

We solely used specimens consisting of one tissue (hard bone, soft bone, muscle, fat, and skin tissue). This setup, however, is not feasible during surgery. Therefore, we plan on investigating the more challenging case where specimens consist of multiple layers of tissues. We plan to develop a depth approximation for the Er:YAG laser ablation. This approximation can provide further crucial information to assist tissue differentiation during laser ablation.

ACKNOWLEDGMENT

The authors are grateful to Robin Sandkühler and Florentin Bieder for the helpful discussions. The authors would like to thank the authors of [25], for providing them with their experimental data. *(Antal Huck and Hervé Nguendon Kenhagho are co-second authors.) (Azhar Zam and Philippe C. Cattin are co-last authors.)*

REFERENCES

[1] J. C. Hu, X. Gu, S. R. Lipsitz, M. J. Barry, A. V. D'Amico, A. C. Weinberg, and N. L. Keating, "Comparative effectiveness of minimally invasive vs open radical prostatectomy," *JAMA*, vol. 302, no. 14, pp. 1557–1564, 2009.

[2] J. D. Luketich, M. Alvelo-Rivera, P. O. Buenaventura, N. A. Christie, J. S. McCaughan, V. R. Litle, P. R. Schauer, J. M. Close, and H. C. Fernando, "Minimally invasive esophagectomy: Outcomes in 222 patients," *Ann. Surg.*, vol. 238, no. 4, p. 486, 2003.

[3] D. D. Lo, M. A. Mackanos, M. T. Chung, J. S. Hyun, D. T. Montoro, M. Grova, C. Liu, J. Wang, D. Palanker, A. J. Connolly, M. T. Longaker, C. H. Contag, and D. C. Wan, "Femtosecond plasma mediated laser ablation has advantages over mechanical osteotomy of cranial bone," *Lasers Surg. Med.*, vol. 44, no. 10, pp. 805–814, Dec. 2012.

[4] B. Kyung-Won, M. Dard, H.-F. Zeilhofer, P. Cattin, and P. Jürgens, "Comparing the bone healing after cold ablation robot-guided Er:YAG laser osteotomy and piezoelectric osteotomy—A pilot study in minipig mandible," *Lasers Surg. Med.*, vol. 53, no. 3, pp. 291–299, 2021.

[5] M. Augello, W. Deibel, K. Nuss, P. Cattin, and P. Jürgens, "Comparative microstructural analysis of bone osteotomies after cutting by computer-assisted robot-guided laser osteotome and piezoelectric osteotome: An *in vivo* animal study," *Lasers Med. Sci.*, vol. 33, no. 7, pp. 1471–1478, Sep. 2018.

[6] S. R. Visuri, J. T. Walsh, and H. A. Wigdor, "Erbium laser ablation of dental hard tissue: Effect of water cooling," *Lasers Surg. Med.*, vol. 18, no. 3, pp. 294–300, 1996.

[7] K.-W. Baek, W. Deibel, D. Marinov, M. Griessen, M. Dard, A. Bruno, H.-F. Zeilhofer, P. Cattin, and P. Juergens, "A comparative investigation of bone surface after cutting with mechanical tools and Er:YAG laser," *Lasers Surg. Med.*, vol. 47, no. 5, pp. 426–432, Jul. 2015.

[8] H. Kang, I. Rizoiu, and A. Welch, "Hard tissue ablation with a spray-assisted mid-ir laser," *Phys. Med. Biol.*, vol. 52, no. 24, p. 7243, 2007.

[9] H. K. N. Kenhagho, G. Rauter, R. Guzman, P. C. Cattin, and A. Zam, "Comparison of acoustic shock waves generated by micro and nanosecond lasers for a smart laser surgery system," *Proc. SPIE*, vol. 10484, Feb. 2018, Art. no. 104840P.

[10] L. M. B. Bernal, G. Shayeganrad, G. Kosa, M. Zelechowski, G. Rauter, N. Friederich, P. C. Cattin, and A. Zam, "Performance of Er:YAG laser ablation of hard bone under different irrigation water cooling conditions," *Proc. SPIE*, vol. 10492, Feb. 2018, Art. no. 104920B.

[11] H. K. Nguendon, N. Faivre, B. Meylan, S. Shevchik, G. Rauter, R. Guzman, P. C. Cattin, K. Wasmer, and A. Zam, "Characterization of ablated porcine bone and muscle using laser-induced acoustic wave method for tissue differentiation," in *Proc. Eur. Conf. Biomed. Opt.*, 2017, Art. no. 104170N.

[12] H. Abbasi, L. Beltrán, G. Rauter, R. Guzman, P. C. Cattin, and A. Zam, "Effect of cooling water on ablation in Er:YAG laserosteotome of hard bone," *Proc. SPIE*, vol. 10453, Aug. 2017, Art. no. 104531I.

[13] G. Rauter, "The miracle," in *Lasers Oral Maxillofacial Surgery*. Cham, Switzerland: Springer, 2020, pp. 247–253.

[14] M. Eugster, P. Weber, P. Cattin, A. Zam, G. Kosa, and G. Rauter, "Positioning and stabilisation of a minimally invasive laser osteotome," in *Proc. 10th Hamlyn Symp. Med. Robot.*, Jun. 2017, p. 21.
[15] M. Eugster, P. C. Cattin, A. Zam, and G. Rauter, "A parallel robotic mechanism for the stabilization and guidance of an endoscope tip in laser osteotomy," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2018, pp. 1306–1311.
[16] H. Abbasi, G. Rauter, R. Guzman, P. C. Cattin, and A. Zam, "Differentiation of femur bone from surrounding soft tissue using laser-induced breakdown spectroscopy as a feedback system for smart laserosteotomy," *Proc. SPIE*, vol. 10685, May 2018, Art. no. 1068519.
[17] A. Zam, "Optical tissue differentiation for sensor-controlled tissue-specific laser surgery," Friedrich-Alexander-Universität Erlangen-Nürnberg, Technische Fakultät, Erlangen, Germany, Tech. Rep., 2011, doi: 10.25593/978-3-87525-318-4.
[18] F. Stelzle, A. Zam, W. Adler, K. Tangermann-Gerk, A. Douplik, E. Nkenke, and M. Schmidt, "Optical nerve detection by diffuse reflectance spectroscopy for feedback controlled oral and maxillofacial laser surgery," *J. Transl. Med.*, vol. 9, no. 1, p. 20, Dec. 2011.
[19] Y. A. Bayhaqi, A. Navarini, G. Rauter, P. C. Cattin, and A. Zam, "Neural network in tissue characterization of optical coherence tomography (OCT) image for smart laser surgery: Preliminary study," *Proc. SPIE*, vol. 11044, Apr. 2019, Art. no. 1104402.
[20] Y. A. Bayhaqi, G. Rauter, A. Navarini, P. C. Cattin, and A. Zam, "Fast optical coherence tomography image enhancement using deep learning for smart laser surgery: Preliminary study in bone tissue," *Proc. SPIE*, vol. 11207, Oct. 2019, Art. no. 112070Z.
[21] L. Grad, J. I. Mozina, D. Sustercic, N. Funduk, U. Skaleric, M. Lukac, S. Cencic, and K. N. Nemes, "Optoacoustic studies of Er:YAG laser ablation in hard dental tissue," *Proc. SPIE*, vol. 2128, pp. 456–465, Sep. 1994.
[22] K. Tangermann, S. Roth, D. Müller, H. Tragler, J. Uller, and S. Rupprecht, "Sensor-controlled laser processes for medical applications," *Proc. SPIE*, vol. 5287, pp. 24–34, Dec. 2003.
[23] S. Rupprecht, K. Tangermann-Gerk, J. Wiltfang, F. W. Neukam, and A. Schlegel, "Sensor-based laser ablation for tissue specific cutting: An experimental study," *Lasers Med. Sci.*, vol. 19, no. 2, pp. 81–88, Oct. 2004.
[24] H. N. Kenhagho, G. Rauter, R. Guzman, P. Cattin, and A. Zam, "Optoacoustic tissue differentiation using a Mach–Zehnder interferometer: Preliminary results," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Oct. 2018, pp. 1–9.
[25] H. N. Kenhagho, F. Canbaz, T. E. G. Alvarez-Arenas, R. Guzman, P. Cattin, and A. Zam, "Machine learning-based optoacoustic tissue classification method for laser osteotomes using an air-coupled transducer," *Lasers Surg. Med.*, vol. 53, no. 3, pp. 377–389, Mar. 2021.
[26] U. Nahum, A. Zam, and P. C. Cattin, "Bone reconstruction and depth control during laser ablation," in *Proc. Int. Workshop Comput. Methods Clin. Appl. Musculoskeletal Imag.* Springer, 2018, pp. 126–135.
[27] U. Nahum, C. Seppi, and P. C. Cattin, "Joint inverse medium and optimal control problem for acoustic waves," in *Proc. Platform Adv. Sci. Comput. Conf.*, Jun. 2019, p. 8.
[28] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
[29] S. N. Deepa, "A survey on artificial intelligence approaches for medical image classification," *Indian J. Sci. Technol.*, vol. 4, no. 11, pp. 1583–1595, Nov. 2011.
[30] Q. Li, W. Cai, X. Wang, Y. Zhou, D. D. Feng, and M. Chen, "Medical image classification with convolutional neural network," in *Proc. 13th Int. Conf. Control Autom. Robot. Vis. (ICARCV)*, Dec. 2014, pp. 844–848.
[31] A. Kumar, J. Kim, D. Lyndon, M. Fulham, and D. Feng, "An ensemble of fine-tuned convolutional neural networks for medical image classification," *IEEE J. Biomed. Health Inform.*, vol. 21, no. 1, pp. 31–40, Jan. 2017.
[32] A. S. Miller, B. H. Blott, and T. K. Hames, "Review of neural network applications in medical imaging and signal processing," *Med. Biol. Eng. Comput.*, vol. 30, no. 5, pp. 449–464, Sep. 1992.
[33] S. Kiranyaz, T. Ince, and M. Gabbouj, "Real-time patient-specific ecg classification by 1-d convolutional neural networks," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 3, pp. 664–675, 2015.
[34] Y. Han and K. Lee, "Convolutional neural network with multiple-width frequency-delta data augmentation for acoustic scene classification," in *Proc. IEEE AASP Challenge Detection Classification Acoustic Scenes Events*, Sep. 2016. [Online]. Available: http://dcase.community/documents/challenge2016/technical_reports/DCASE2016_Lee_1034.pdf
[35] H. Kamper, W. Wang, and K. Livescu, "Deep convolutional acoustic word embeddings using word-pair side information," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2016, pp. 4950–4954.
[36] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
[37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: http://arxiv.org/abs/1412.6980
[38] A. B. Downey, *Think DSP: Digital Signal Processing in Python*. Sebastopol, CA, USA: O'Reilly Media, 2016.
[39] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in Pytorch," in *Proc. 31st Conf. Neural Inf. Process. Syst. (NIPS)*, Long Beach, CA, USA, 2017.

CARLO SEPPI started to study mathematics at the Mathematical Institute, University of Basel, Allschwil, Switzerland, in 2009. He received the B.Sc. and M.Sc. degrees in mathematics, in 2013 and 2015, respectively. He is currently pursuing the Ph.D. degree with the Planning and Navigation Group, Department of Biomedical Engineering, University of Basel.

During his studies, he worked as a Teaching Assistant with the Mathematical Institute. He started his Swiss Civil Service at the Department of Biomedical Engineering, University of Basel, in 2016. He helped to develop novel visualization methods for nuclear medicine to detect sentinel lymph nodes. His project involves solving inverse problems in medical applications: e.g., finding the position of sentinel lymph nodes using a pinhole collimator (reconstructing a ^{99m}Tc source using a single image of the detector), reconstruction of the medium velocity and the acoustic source with the help of the Helmholtz equation, and extracting information of the acoustic wave during the ablation of tissue.

ANTAL HUCK was born in Basel, Switzerland. He received the B.Sc. and M.Sc. degrees in mathematics from the University of Basel, Switzerland, with a focus on numerical analysis, and the Ph.D. degree from the Faculty of Medicine, University of Basel, in 2019, working on the research project "Segmentation of the Gray and White Matter in the Human Spinal Cord." Since then, he has been a Postdoctoral Researcher with the Group of Prof. Philippe Cattin, Department of Biomedical, University of Basel.

HERVÉ NGUENDON KENHAGHO received the B.E. degree (Hons.) in communication and electronic engineering and the M.Sc. degree in microelectronic and communication engineering from Northumbria University, U.K., in 2009 and 2010, respectively, the professional certificate as a Laser Safety Officer from Universität des Saarlandes, Germany, in 2016, and the Ph.D. degree in biomedical engineering from the University of Basel, Switzerland, in 2020, with the focus on the development of opto-acoustical feedback sensors and advanced signal processing for smart laser surgery. After a Calibration Engineer position at Repair & Cal Ltd., U.K., a Research Engineer positions at the University of Lunebürg, and the University Carlos Tercero de Madrid, and a Researcher Assistant and a Collaborator positions at the Biomedical Laser and Optics Group (BLOG), Department of Biomedical Engineering, University of Basel, he joined the Sensor Application Innovation Laboratory (SAIL), Huawei Technologies, Grenoble, France, as a Senior Research and Development Engineer in (electro)-optical systems, in 2021. His main research interests include the development of smart sensor devices and miniaturized systems for mobile, automotive, and medical applications.

EVA SCHNIDER received the M.Sc. degree in mathematics from the University of Basel, Switzerland, in 2018. She is currently pursuing the Ph.D. degree with the MIRACLE Planning and Navigation Group of Prof. Philippe Cattin, Department of Biomedical Engineering, University of Basel.

Her thesis focused on spectral decomposition as a regularization method for inverse problems involving the acoustic wave equation in breast tumor detection. Her project focuses on the development of machine learning algorithms for distinct bone segmentation in CT images. The results of which will then be used in the group's virtual reality surgical planning systems. She interned at Google Research in Zürich and London working on machine perception and on fairness aspects of machine learning for time series predictions. Her research interests include medical image segmentation, neural network architectures for three-dimensional imaging modalities, and data augmentation schemes.

GEORG RAUTER (Member, IEEE) received the Diploma degree in mechanical modeling from MATMECA, Bordeaux, France, in 2006, the second Diploma degree in mechanical engineering from TU Graz, Austria, in 2008, and the Ph.D. degree in robotics from ETH Zurich, Switzerland, in 2014. From 2014 to 2016, he was a Postdoctoral Researcher in data analysis and rehabilitation robotics at ETH Zurich, the University of Southern California, and the University of Zurich. In 2016, he commercialized the gait rehabilitation robot the FLOAT in collaboration with the company Lutz Medical Engineering and the Spinal Cord Injury Center at Balgrist University Hospital, Zürich. Since 2016, he has been heading the BIROMED-Lab as an Assistant Professor of medical robotics and mechatronics with the Department of Biomedical Engineering, University of Basel, Switzerland. His research interests include the development of robotic and mechatronic systems, with a special focus on enabling minimally invasive surgical procedures and providing new and efficient rehabilitation methods. He has been the Chair of IFToMMm, Switzerland, since 2019.

AZHAR ZAM (Member, IEEE) received the B.Sc. degree in medical physics from the University of Indonesia, Depok City, Indonesia, in 2004, the M.Sc. degree in biomedical engineering from the University of Lübeck, Lübeck, Germany, in 2007, and the Ph.D. degree in advanced optical technologies from Friedrich-Alexander University Erlangen-Nuremberg, Germany, with the focus on optical feedback for tissue-specific laser surgery, in 2011. He held research positions at the University of Waterloo, Waterloo, ON, Canada, the National University of Ireland Galway, Galway, Ireland, and the University of California at Davis, Davis, CA, USA. He joined the Department of Biomedical Engineering, University of Basel, Allschwil, Switzerland, in 2016, as an Assistant Professor, where he founded and leads the Biomedical Laser and Optics Group (BLOG). He has authored over 80 peer-reviewed articles, book chapters, books, and patents. His main research interests include the development of smart devices for medical therapy, diagnostics, and monitoring using novel optical technologies, which include smart laser surgery, optical coherence tomography (OCT), photoacoustics, biomedical spectroscopy, AI-aided optical diagnostics and imaging, optical-based smart wearable sensors, and miniaturized systems.

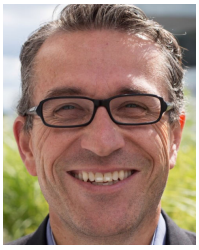

PHILIPPE C. CATTIN (Member, IEEE) was born in Baden, Switzerland, in 1967. He received the B.Sc. degree from the University of Applied Sciences and Arts Northwestern Switzerland, Windisch, Switzerland, in 1991, and the M.Sc. degree in computer science and the Ph.D. degree in robotics from ETH Zurich, Zürich, Switzerland, in 1995 and 2003, respectively.

From 2003 to 2007, he was a Postdoctoral Fellow with the Computer Vision Laboratory, ETH Zurich. In 2007, he became an Assistant Professor with the University of Basel, Switzerland, and was promoted to an Associate Professor, in 2015, and a Full Professor, in 2019. He is the Founder of the Center for medical Image Analysis & Navigation (CIAN), Medical Faculty, University of Basel, where he is the Founding Head and still heading the Department of Biomedical Engineering. In 2017, he joined Brigham and Women's Hospital, Boston, MA, USA, as a Research Fellow. He is also the Founder of two spin-off companies and licensed his patents and software to medical device companies. His research interests include medical image analysis, image-guided therapy, robotics-guided laser osteotomy, and virtual reality. As the Principal Investigator, he has finished many projects in these areas and published over 250 papers, patents, and book chapters.