

A Deep Learning Pipeline for Identification of Motor Units in Musculoskeletal Ultrasound

HAZRAT ALI¹, JOHANNES UMANDER, ROBIN ROHLÉN¹,
AND CHRISTER GRÖNLUND¹

Department of Radiation Sciences, Umea University, 901 87 Umea, Sweden

Corresponding author: Christer Grönlund (christer.gronlund@umu.se)

This work was supported in part by the Swedish Research Council under Grant dnr 2015-04461, and in part by the Kempe Foundations under Grant dnr JCK-1115 and Grant SMK-1868.

ABSTRACT Skeletal muscles are functionally regulated by populations of so-called motor units (MUs). An MU comprises a bundle of muscle fibers controlled by a neuron from the spinal cord. Current methods to diagnose neuromuscular diseases and monitor rehabilitation, and study sports sciences rely on recording and analyzing the bio-electric activity of the MUs. However, these methods provide information from a limited part of a muscle. Ultrasound imaging provides information from a large part of the muscle. It has recently been shown that ultrafast ultrasound imaging can be used to record and analyze the mechanical response of individual MUs using blind source separation. In this work, we present an alternative method - a deep learning pipeline - to identify active MUs in ultrasound image sequences, including segmentation of their territories and signal estimation of their mechanical responses (twitch train). We train and evaluate the model using simulated data mimicking the complex activation pattern of tens of activated MUs with overlapping territories and partially synchronized activation patterns. Using a slow fusion approach (based on 3D CNNs), we transform the spatiotemporal image sequence data to 2D representations and apply a deep neural network architecture for segmentation. Next, we employ a second deep neural network architecture for signal estimation. The results show that the proposed pipeline can effectively identify individual MUs, estimate their territories, and estimate their twitch train signal at low contraction forces. The framework can retain spatio-temporal consistencies and information of the mechanical response of MU activity even when the ultrasound image sequences are transformed into a 2D representation for compatibility with more traditional computer vision and image processing techniques. The proposed pipeline is potentially useful to identify simultaneously active MUs in whole muscles in ultrasound image sequences of voluntary skeletal muscle contractions at low force levels.

INDEX TERMS Motor unit, decomposition, ultrafast ultrasound, medical imaging, deep learning, mechanical response, neural networks, recurrent neural networks.

I. INTRODUCTION

The motor unit (MU) is the smallest voluntarily activatable unit in the skeletal muscles. Its function (and control) is important to study for the diagnosis of neuromuscular diseases and understanding of skeletal muscle physiology in sports sciences and rehabilitation [1], [2]. An MU is defined as a motor neuron connected to a bundle of muscle fibers located within a given territory (Fig. 1A) [3]. The control of an MU is encoded in a firing pattern (Fig. 1A) originating in the spinal cord mediated by the motor neuron. The corresponding output is characterized

by repeated electrical depolarizations of the fibers and subsequent repeated shortening and thickening of the fibers (mechanical twitch train) [4] (Fig. 1A and B).

Electromyography (EMG) is the gold standard technique to study MUs where electrodes are used to record the fibers' electrical activity either invasively or from the surface of the skin [5]–[7]. This technique provides high-quality data, but due to a low pass filtering effect of the tissue, there is a restricted field of view [8]. Ultrasound imaging is a non-invasive technique allowing mechanical information from a large field of view [9].

Recently our group presented a method [10] to study MU activity based on the mechanical response of individual MUs using ultrafast ultrasound imaging (>2000 images per

The associate editor coordinating the review of this manuscript and approving it for publication was Jun Wang¹.

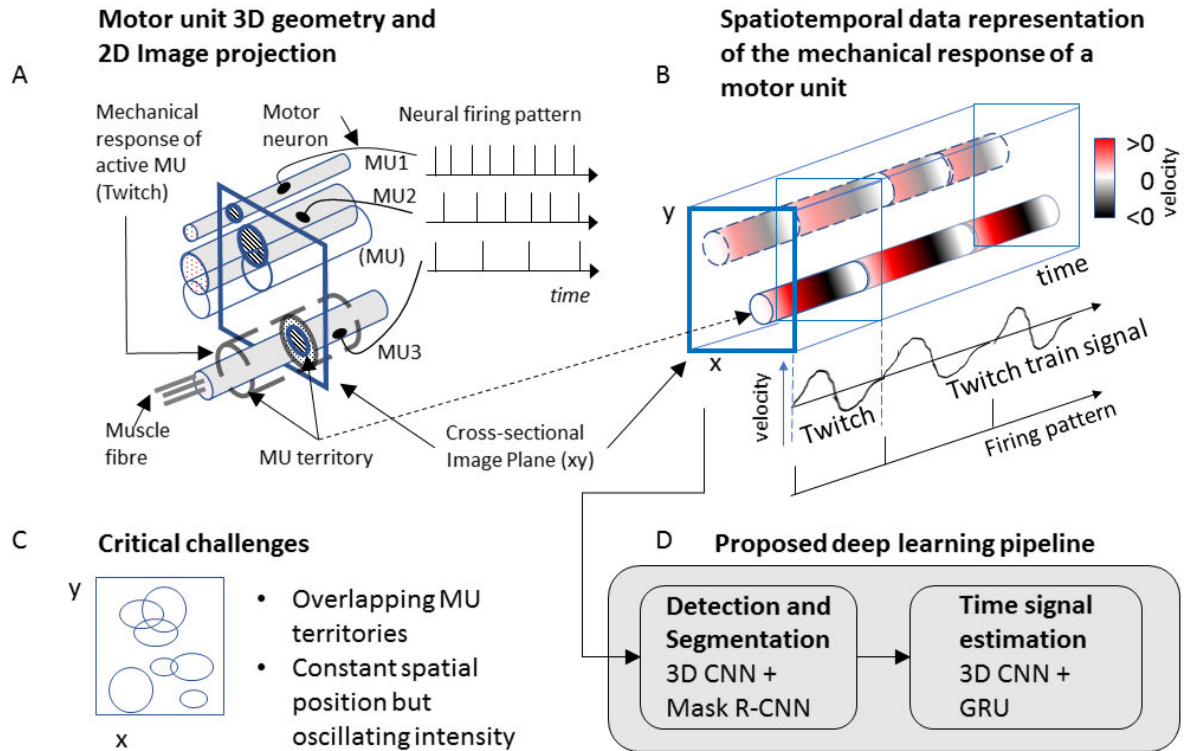


FIGURE 1. Illustration of principal skeletal muscle anatomy and recorded ultrasound image sequence. A) Shows three motor units, aligned in parallel, and they are activated with unique neural firing patterns. An activated unit’s mechanical response results in a twitch - a thickening and shortening of the unit’s fibers. B) An illustration of the spatio-temporal features of the recorded image sequence from the cross-sectional plane. C) The main challenges of the data, including the overlap of motor unit territories. D) An overview of the proposed deep learning approach comprising two modules of detection and segmentation and time signal detection.

second [11]). This method was based on a blind source separation framework and decomposition of spatio-temporal components. However, the performance was found to decrease with an increasing number of active MUs in the contractions. This problem’s challenge is that muscle activation is a highly complex physiological process, where tens to hundreds of MUs with overlapping MU territories can be active simultaneously with individual firing patterns (Fig. 1A and C).

One interesting approach that has shown tremendous potential to learn complex patterns is using deep learning models comprising neural networks with several layers’ architecture. In particular, in medical imaging, there is a vast literature on deep learning applications for detection and segmentation [12]. *In this work, we hypothesize that a deep learning methodology can improve the performance of identification of individual MUs, by learning the underlying complex interaction patterns using the full image sequence information.*

This work aims to develop and evaluate a deep learning pipeline to 1) detect individual active MUs, 2) segment their territories, and 3) estimate their activation twitch signal, using ultrafast image sequences of voluntary skeletal muscle contractions. The model is trained and evaluated using simulated ultrasound data [10].

This work presents a deep learning-based method to identify individual MUs in spatio-temporal data of contracting muscles. To the best of our knowledge, it is the first deep learning approach to identify a varying number of fixed-position objects with unique individual temporal patterns of intensity changes.

The rest of the paper is organized as follows: In Section II, we review related work. In Section III, we present the proposed deep learning pipeline in detail and also present the evaluation metrics. Section IV gives an overview of the simulation model and data sets generated. Section V presents the results, and Section VI gives a discussion of the findings. Finally, Section VII concludes the paper.

II. RELATED WORK

Deep learning [13]–[15] has greatly revolutionized many different domains involving analysis of a large image, audio, text, video, or tabular data. Of particular relevance to the work reported in this article are the advancement made in image and video processing using deep learning methods for segmentation, identification, recognition [16], deep learning for time-series data (e.g., speech) [17], and deep learning in medical imaging [12], [18]. Hence, we present the relevant details on deep learning for segmentation and signal detection in the subsequent text.

A. OBJECT DETECTION AND SEGMENTATION

The advancements made on instance segmentation tasks in the computer vision field have paved the way for many progress applications for deep learning applications in the medical imaging domain. Traditionally, the classification task was performed to categorize an image into a distinct class (e.g., cat vs. dog classification). For a more realistic task, we are usually interested in the object's position within an image, referred to as the localization task. When there are multiple instances of the same object within an image, we perform object detection, which localizes the object. For more practical use, we do pixel level localization of the object referred to as the segmentation.

In semantic segmentation, each class in the image is masked with a different color. For example, all the pixels containing dogs might be colored blue, and all the pixels containing cats may be colored red. A problem with this approach is when objects of the same class overlap, they are merged under the same mask, and it is difficult to differentiate between them. To solve this problem, one can use instance segmentation where every object (instance) gets its own mask [19]–[22].

One of the simplest methods to perform object detection is to crop out multiple locations of an image and run a Convolutional Neural Network (CNN) to classify the cropped region. The problem with this approach is that it is extremely slow. R-CNN (Region-based CNN) [19] tried to solve this by first applying a non-learning-based algorithm called a selective search that returned 2,000 likely region proposals that a CNN then classified and predicted a more refined bounding box around the object. This approach is considered slow because CNN has to classify 2,000 regions. Fast R-CNN [20] solved this problem by using a CNN to process the entire image into a convolutional features map. Then, to classify one of the region proposals, one can crop out a region of this features map corresponding to the region in the image using RoI pooling and then classify that data. This approach is faster than R-CNN, but it still relies on the non-learning based selective search to find interesting region proposals. Faster R-CNN [21] replaced the selective search algorithm with a new network called region proposal network (RPN) that used the information from the convolutional features map to generate region proposals.

Mask R-CNN [22] is an extension of the Faster R-CNN architecture to introduce instance segmentation. To achieve this, mask R-CNN modified some parts of the network. The convolutional features map is replaced by a feature pyramid network (FPN), which contains feature maps at multiple scales of the image. RoI pooling is replaced with a new method called RoIAlign, which works better when pixel-level accuracy is required. A pixel map containing the object is generated by adding a new head to the classifier and bounding box predictor heads for predicting the mask.

B. TIME SIGNAL ESTIMATION

Typically, recurrent neural networks (RNN) have been popular with time-series data processing. Successful use cases have been reported for sequential data, in particular in natural language processing. Comparative studies have shown the benefit of RNN for time dependencies modeling and signal tracking [17], [23]–[26]. For example, the work in [17] achieved a state-of-the-art error reduction on the popular TIMIT dataset for speech recognition. The Gate Recurrent Unit (GRU) is a type of RNN that helps overcome the vanishing gradient issue in training an RNN. This issue is typically achieved by updating and reset gates controlling the inward and outward flow of information from the neural network's memory states. This process effectively helps in removing information that is redundant for the prediction task. For additional details on RNN and GRU, we refer the reader to [17], [27].

III. PROPOSED MODEL

A. PREREQUISITES AND OVERVIEW

Three key features characterize the mechanical response of individual MUs that we want to identify:

- 1) Spatio-temporal characteristics of units: The mechanical response of an MU is characterized by a fixed location of a spatial territory with varying intensity (Fig 1B). The time signal intensity variation of a unit is unique due to an MU's unique neural firing pattern (Fig 1A) [3].
- 2) Unknown number of units: The number of active (and thus visible) MUs in an image sequence is unknown because the recruitment of units is highly complex coordination by the central nervous system [3], [10].
- 3) Overlapping territories of the units: The territories of two or more MUs may be overlapping [3], causing spatial and temporal interference of their activity (Fig 1A and C).

NB: In the text, we sometimes write *object* for MU or unit, and *signal estimation* for extraction of the mechanical twitch train signal, to harmonize with the terminology in machine learning and signal/image processing.

Given the spatio-temporal nature of the mechanical response of an MU (Fig 1B), we split the problem into two main modules (Fig 1D).

The first module is the **detection and segmentation model**, which detects and segments the MU territories within the image sequence. The second module, called the **time signal estimation model**, determines the mechanical activation signal (the twitch train) caused by a specified MU. As typical, the best parameters for the deep networks used here are determined empirically and through a grid-search on a finite set of parameters.

B. DETECTION AND SEGMENTATION MODEL

This module processes an image sequence to perform the detection and segmentation of the MUs. This process is

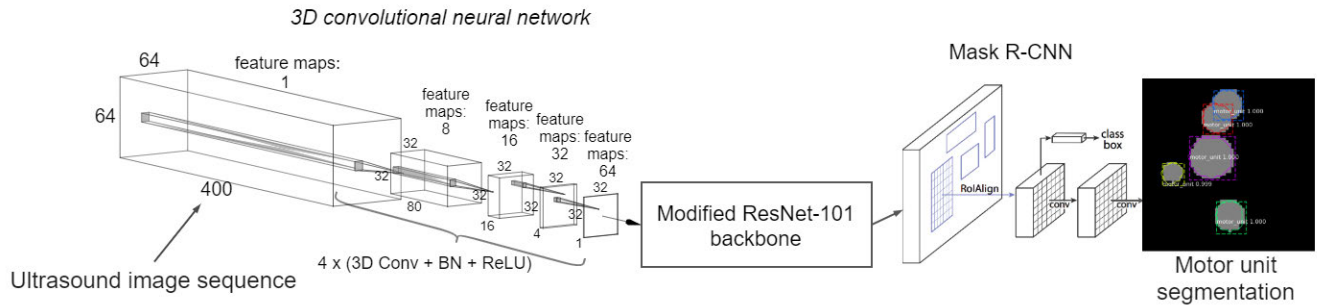


FIGURE 2. Detection and segmentation model architecture. Input data is the spatio-temporal image sequence of the complex activity of activated motor units in a skeletal muscle ($64 \times 64 \times 400$ is *Depth* \times *Width* \times *Frames* and corresponds to 40×40 mm and 1 second). The output is the segmented spatial territories of identified motor units (right).

particularly challenging due to the spatio-temporal nature of the mechanical response of MUs (as previously pointed out). In short, we use a 3D CNN, which helps to retain the temporal information while generating a 2D representation. The transformed 2D representation is used for instance segmentation using Mask R-CNN approach.

1) ARCHITECTURE

The first thing that takes place in the model is to convert the data into a 2D representation using a series of 3D convolutions. The slow fusion approach [28] inspires this way of processing image sequence data. The network preserves the temporal information on the action potential. More global information is made accessible to the higher layers in the network, retaining both the ultrasound sequence's spatial and temporal aspects. During this transformation, we keep the spatial resolution while reducing the temporal dimension. This architecture is visualized on the left part of Figure 2.

The first layer receives the standardized image sequence of size $64 \times 64 \times 400 \times 1$, and it has to reduce the size of the data to make the computational problem feasible. It uses a strided convolution of dimensions $2 \times 2 \times 5$ (*height* \times *width* \times *time*) to reduce the number of computations required and also reduces the spatial dimensions by a factor of two and the temporal dimension by a factor of 5. It uses a kernel size of $(7 \times 7 \times 7)$ and 8 feature maps followed by batch normalization (BN) and the ReLU activation function. The strided convolution reduces the size of the data to $32 \times 32 \times 80 \times 8$. A second convolution is performed on this data, which uses a kernel size of $(3 \times 3 \times 7)$ and 16 feature maps followed by BN and ReLU. Since performance is not as critical at this point compared to the first layer due to the reduced data size, max pooling is used instead of strided convolutions to reduce the temporal dimension. A max pooling layer reduces the temporal dimension by a factor of 5 to produce a data size of $32 \times 32 \times 16 \times 16$. The third convolution uses a kernel size of $(3 \times 3 \times 5)$ and 32 feature maps followed by BN, ReLU, and a max pooling layer, reducing the temporal dimension a factor of 4. This results in a data size of $32 \times 32 \times 4 \times 32$. At this stage, the final convolution uses a kernel of size $3 \times 3 \times 5$,

64 feature maps, BN, ReLU, and a max pooling layer that reduces the temporal dimension a factor of 4 to create a data size of $32 \times 32 \times 1 \times 64$.

At this stage, the data gets a 2D representation with 32×32 pixels and 64 channels. We now employ a mask R-CNN model [22], with implementation from [29] with a ResNet-101 model used for features extractions [30] as the feature extractor¹ (See Fig. 2). An issue that arises is that ResNet is normally used for high-resolution images. The ResNet network's first stage reduces the spatial dimensions by a factor of 4, which is unwanted when we already have a low-resolution image and would result in a very poor segmentation. In our work, this issue is addressed by removing the first stage of the ResNet network.

The actual object detection and segmentation is then performed by Mask R-CNN, which receives the features generated by ResNet-101 through a feature pyramid network (FPN). The Mask R-CNN network in this model retains the default parameters used in the implementation.

2) PERFORMANCE METRICS

We evaluate the detection and segmentation performances through precision and recall measures.

An MU is considered correctly detected if the intersection over union (IoU) measure for the MU mask is greater than 0.5. The detection step drives the segmentation performance, as we would like to consider MUs that as classified as true positives during the detection step.

3) TRAINING PROCESS

Ten thousand (10,000) simulated images sequences were generated for the training (see section IV. B Datasets). To further increase the training data, we perform data augmentation (random flipping) and get a total training set of 40,000 simulated image sequences, containing 699,788 MUs.

To speed up the training process and achieve better convergence, we used transfer learning. We use ResNet-101 [30] and Mask R-CNN models [22], pre-trained on the Microsoft

¹The implementation is in Python3 and uses the tensorflow and keras frameworks.

COCO dataset [31]. The following changes are introduced. The COCO dataset has 80 different categories. So, the classifier layer is modified to suit our task of binary classification. The 3D CNN layers are trained from scratch with weights initialized using the Xavier initialization method [32].

The model is trained using stochastic gradient descent (SGD) with an initial learning rate of 0.001 and momentum set to 0.9.²

For the first 20,000 mini-batches, the ResNet backbone weights are kept fixed so that to preserve the knowledge from the previous application and the other layers have to adjust to it. Afterward, the learning rate is reduced by a factor of 10 for training up to 200,000 mini-batches. Finally, the learning rate is further reduced by a factor of 10 as the training continues to the 500,000 mini-batch. The training takes a total of 67 hours as we get the optimal weights. The training and validation loss graphs can be seen in Figure 9 (in Appendix A). The loss function used for training this model is the sum of five different loss functions: the region proposal networks loss, the bounding box loss, the Mask R-CNNs loss, bounding box loss, and the mask loss.

In addition, to improve the model's tolerance to noise, the model is also further trained for an additional 100,000 mini-batches with a random noise level value between 30 to 10 dB SNR.

C. TIME SIGNAL ESTIMATION MODEL

The time signal estimation model is designed to learn to estimate the MU twitch train signal from a detected and segmented MU from the first model. Thus it removes signals that do not originate from the MU of interest. In short, with motivation from the slow fusion approach by Karpathy *et al.* [28], we first use a 3D CNN to transform the image sequences into a time-series data and then train a GRU network to estimate the twitch train.

1) ARCHITECTURE

The input to the model is a cropped version of the spatio-temporal image sequence. It has a spatial resolution of 16×16 pixels and 400 timesteps (corresponding to a one-second sequence within a 10×10 mm ROI), which is sufficient to encompass an MU size in the biceps muscles [33]).

Unlike the segmentation module, here we want to reduce the spatial dimensions and retain the temporal dimension. Similar to the segmentation module, a 3D CNN, inspired by the slow fusion concept [28], is adapted to transform the image sequences for subsequent time signal estimation. It is processed through five layers of 3D convolutions to extract spatio-temporal features.

The first convolution has a stride of 2 in two dimensions (say, x and y) to reduce the spatial resolution to 8×8 and thus reduce the computations.

²The model is trained on an Nvidia RTX 2070 with 8GB of memory which allowed for 8 examples per mini-batch.

Each convolution layer creates 16 filters and uses a $3 \times 3 \times 15$ kernel except the first layer, which has a $5 \times 5 \times 15$ kernel. Each layer is followed by batch normalization and then the ReLU activation function. The resulting data from these 3D convolutions is of dimension $8 \times 8 \times 400 \times 16$. The 3D CNN is visualized in the top left of Figure 3. This data is then sliced by the time steps and flattened so that each time step can be processed individually. The idea behind processing each time step individually is to create higher-level features that are spatially invariant, e.g., the width of the MU at that timestep. This processing is done in a fully connected neural network (FCNN). First, a layer of 1,024 neurons is applied, followed by a layer with 512 neurons, and finally, a layer with 256 neurons. The fully connected layers use the ReLU activation function, and all the time steps share the weights for the FCNN.

The data now consists of a feature vector of 256 numbers for each time step passed to two RNNs, each going in opposite directions, forming a bidirectional RNN. The motivation of using RNN here is to create a temporal signal based on the information from all the timesteps. For example, suppose the signal from an MU is unintelligible during an interval. In that case, RNNs could be used to look at the previous and preceding time steps to estimate the signal in the missing interval. Each RNN consists of 512 gated recurrent units (GRU). The RNNs are configured to return the entire sequence and not just the last output. The output of the RNNs is then concatenated, resulting in 1,024 values for each timestep.

The data is run through a final FCNN at each time step to produce the final result (the amplitude of the twitch train signal of the MU at each time step). First, a layer of 1,024 neurons is applied, followed by a layer with 512 neurons, and finally, a layer produces a single value (the amplitude). All layers except the final layer use ReLU, and all time steps share the weights. The data now consists of 400 values representing the mechanical response (tissue velocity) at each time step.

2) PERFORMANCE METRICS

We evaluate the extracted signals' performance by comparing the firing patterns of the estimated and simulated signals. The performance metric selected for this task is the rate of agreement (RoA), as it provides easy interpretation, captures the information we are interested in [34], and was used in previous work on similar signals [10]. The RoA measure ranges between 0 and 1 where 1 is a perfect score, and it measures the agreement between the firings of two signals. The formula for RoA is:

$$RoA = \frac{c}{c + A + B}, \quad \text{such that } 0 \leq RoA \leq 1 \quad (1)$$

Here, c is the number of times when both the true firings and the estimated firings have matched. A is the number of times signal A had a firing but not signal B, and B is the number of times signal B had a firing but not signal A. Two firings are considered matched if they occur within 15 ms




FIGURE 3. Time signal estimation model architecture. The input to this model is cropped image sequence data centered around the detected motor unit territory as segmented in model 1 ($16 \times 16 \times 400$ is *Depth* \times *Width* \times *Frames* corresponding to 10×10 mm, and 1 second sequence). The output is the corresponding motor unit twitch train signal (bottom).

(6 timesteps at 400 Hz) of each other. An example of how RoA is calculated for a given MU firing can be seen in Figure 11 in Appendix C.

The MU firings are extracted as the local maxima of the estimated twitch train signals (Fig. 11). Before this, a low-pass filter is applied to the signals using a running average with a window of 11 timesteps (27.5 ms). Local maxima with an amplitude (velocity) of less than 0 m/s were excluded.

3) TRAINING PROCESS

The twitch train estimation model uses the same training set as the segmentation model. The model is trained for 100,000 mini-batches using Adam optimizer with a learning rate of 0.0001 and clipping the gradient norm to a maximum of 1.³

This model tends to quickly overfit the training data, as shown in Figure 10. No regularization technique seems to fix this problem without also regressing the validation performance. Therefore, early stopping is used to choose the best performing model (on the validation set) instead of the last

³8GB of video memory on the RTX 2070 allows for 32 simulations per mini-batch.

model. The training process took around 17 hours (although, as per Figure 10 shown in Appendix B, only about 10 hours were necessary to find the best performing model).

The loss function used for this model’s training is the mean squared error between the predicted signal and the ground truth signal. However, one issue found with this loss function is that ground truth signals with large amplitudes tend to get much greater losses than ground truth signals with smaller amplitudes. This issue happens even though the predicted signals for the larger amplitude ground truth visually followed the signal much better. This issue results in that the optimizer primarily focusing on improving signals with larger amplitudes and ignoring MUs with weaker signals. This problem was solved by normalizing all the ground truth signals to an amplitude of 1 when calculating the loss and scaling the predicted signals using the same coefficient.

IV. SIMULATION OF MUSCLE ACTIVATION

A simulator [10] generated the data used for training and evaluating the models in this work. Since the simulator knows all the latent variables used to generate the data, it can also

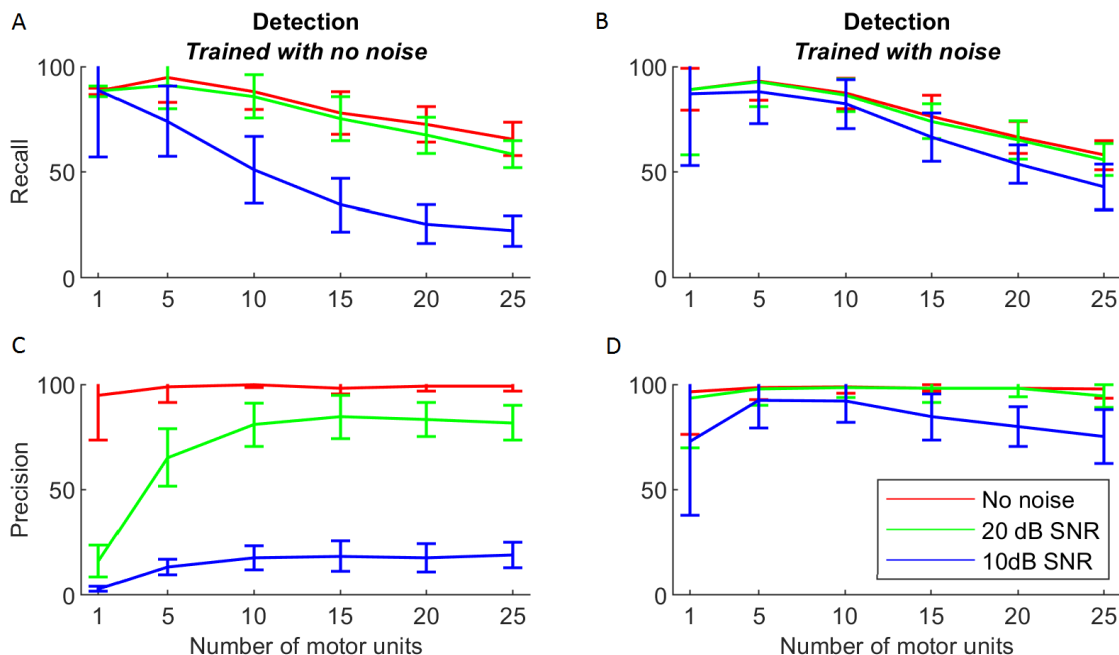


FIGURE 4. Object detection recall (A) and (B) and precision (C) and (D). The model's result was evaluated on the test set with a varying number of motor units and noise. Each data point shows the mean and standard deviation of 100 simulations. Figure 4(A) and (C) correspond to the model trained with noise-free data. Figure 4 (B) and (D) correspond to the model trained on noisy data.

provide the labels for the example in the form of masks of the cross-sectional territory encompassed by each MU and the mechanical twitch train signal for each MU.

A. SIMULATION MODEL

The simulation model used in this work was previously described in Rohlén *et al.* [10], and here we give a brief description. The model generates the tissue velocity image sequences of a contracting muscle based on a modified EMG simulation model [35], where the electrical action potential responses are replaced by mechanical spatio-temporal twitch responses. The mechanical response, in the plane perpendicular to the fiber direction (cross-sectional), is modeled (Fig 1A). An MU territory is modeled as a circular region, and the corresponding mechanical twitch response is modeled using in vivo empirical MU tissue velocity waveform from electro-stimulation experiments [36], [37]. It is assumed that the force is transmitted along the fiber direction only and that there is no mechanical connectivity between the fibers of different MUs. Parameters were set to simulate a biceps brachii muscle at weak isometric contraction levels. The firing patterns of the MUs had a firing rate (FR) in the range 8 and 13 Hz (randomly distributed) with an inter-pulse-interval variation of $\mathcal{N}(0, 0.2/FR)$ [3], [38]. Synchronization of MU firings was simulated in the range of 0-10 % and was computed as the percentage of MU firings synchronized with (firings of) other MUs [39], [40]. The territories of the MUs were randomly located within the simulated muscle cross-section and had a diameter in the range 2.5mm to 10 mm (randomly distributed) [33].

B. DATASETS

We generated three datasets - training, validation, and test sets - consisting of 10,000, 1,000, and 600 simulated image sequences.

The simulated (tissue velocity) image sequences were $64 \times 64 \times 400$ pixels, corresponding to $40 \text{ mm} \times 40 \text{ mm} \times 1 \text{ s}$ with spatial resolution of 0.625 mm/pixel, and 400 Hz frame rate. Each simulated image sequence in the training and validation sets contains between 5 and 30 motor units. In contrast, the test set contains 100 simulated image sequences of each of the following categories: 1, 5, 10, 15, 20, and 25 MUs.

Gaussian white noise was added to the simulated signals at 10, 20, and ∞ dB.

V. RESULTS

Figure 12 and 13 of the Appendix present examples of the image sequences of simulated mechanical response of skeletal muscle activity with 5 and 15 active MUs. Figure 7 and 8 show the true and estimated territories and twitch train signals for these corresponding datasets. It can be seen that the complexity of the mechanical response pattern increases with an increasing number of active MUs. For example, when 5 units were active, all MUs were detected, but when 15 units were active, two units failed to be detected.

A. MOTOR UNIT DETECTION AND SEGMENTATION PERFORMANCE

The object detection results for the segmentation model trained, as described in the method section, can be seen

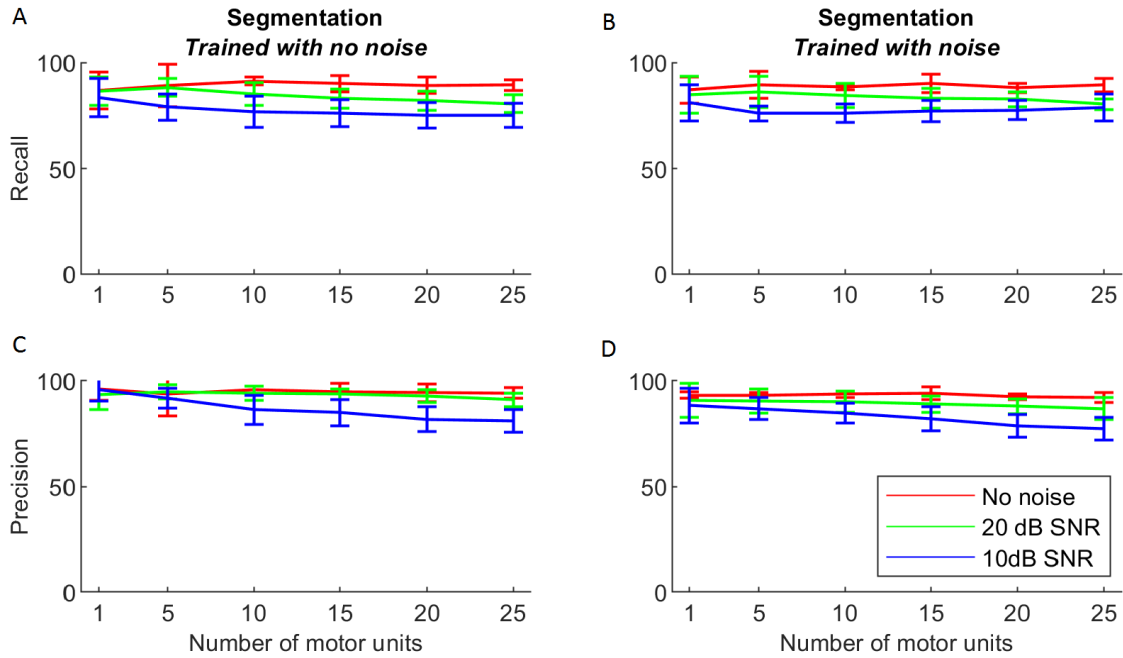


FIGURE 5. Object segmentation recall (A) and (B) and precision (C) and (D). The segmentation model results when evaluated on the test set with varying numbers of motor units and noise. Each data point is computed as the mean of 100 examples. Figure 5(A) and (C) correspond to the model trained with noise-free data. Figure 5 (B) and (D) correspond to the model trained on noisy data.

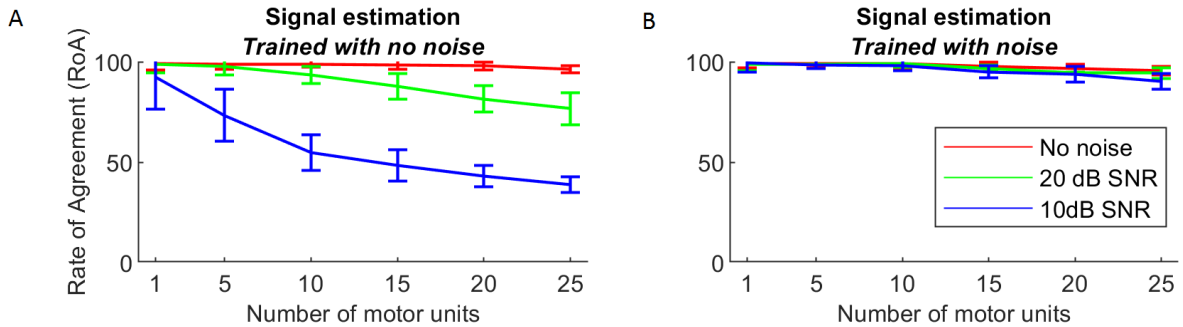


FIGURE 6. Rate of agreement of the twitch train estimation model for evaluation on different noise levels and a different number of motor units. (A) The model is trained with no noise data. (B) The model is trained with noisy data.

in Figure 4. When the model is trained on signals without noise, it achieves high precision in the case when the data has no noise, which means that the model rarely makes miss predictions (i.e., lower false positives). When noise is applied at 20 dB SNR, the model still correctly detects most of the MUs, but it starts to fail to detect some MUs and detect some false MUs as can be seen in the reduced recall and precision, respectively. At 10dB SNR, the model performance is greatly reduced with respect to recall and precision.

In general, training the model with noisy data significantly improved the model (Fig. 4 B and D). The improvement was modest for noise-free or high SNR data but was large for low SNR data and precision. The recall decreased from 90% to 60% with an increasing number of MUs. Also, the recall was approximately stable at >80% for all noise levels.

The segmentation results are shown in Figure 5 (A) and (C) for training with noise-free data and Figure 5 (B) and (D) training with noisy data respectively. The number of MUs did not significantly influence the segmentation performance. Also, training with noise did not impact the segmentation performance compared to when it was trained on noise-free data.

B. TWITCH TRAIN ESTIMATION PERFORMANCE

The twitch train (signal) estimation model’s performance can be seen in Figure 6. When trained without noise, the model performs almost perfectly for low noise case, but performance drops when noise is present, mainly when more MUs are present (Fig 6 A). When the model was trained with noise, the performance was over 90% independently on the number of active MUs or the data’s noise level (Fig 6 B).

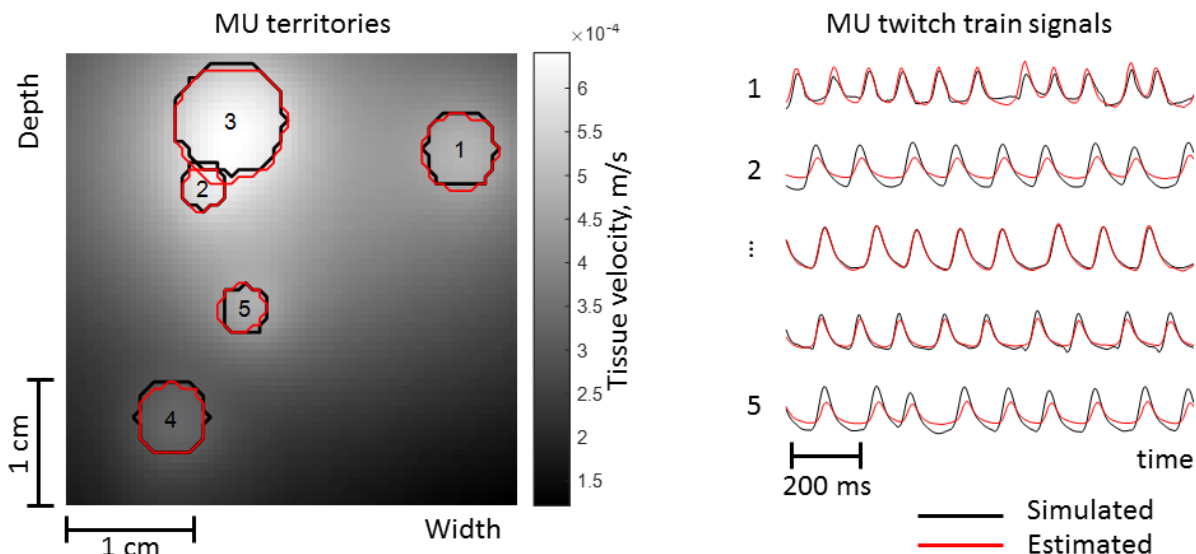


FIGURE 7. Example of results from a simulated tissue velocity image sequence with five active MUs. The left panel shows the simulated and segmented MU territories and an underlying map of the tissue velocity distribution within the image. The right panel shows the simulated and estimated twitch train signals of the corresponding MUs.

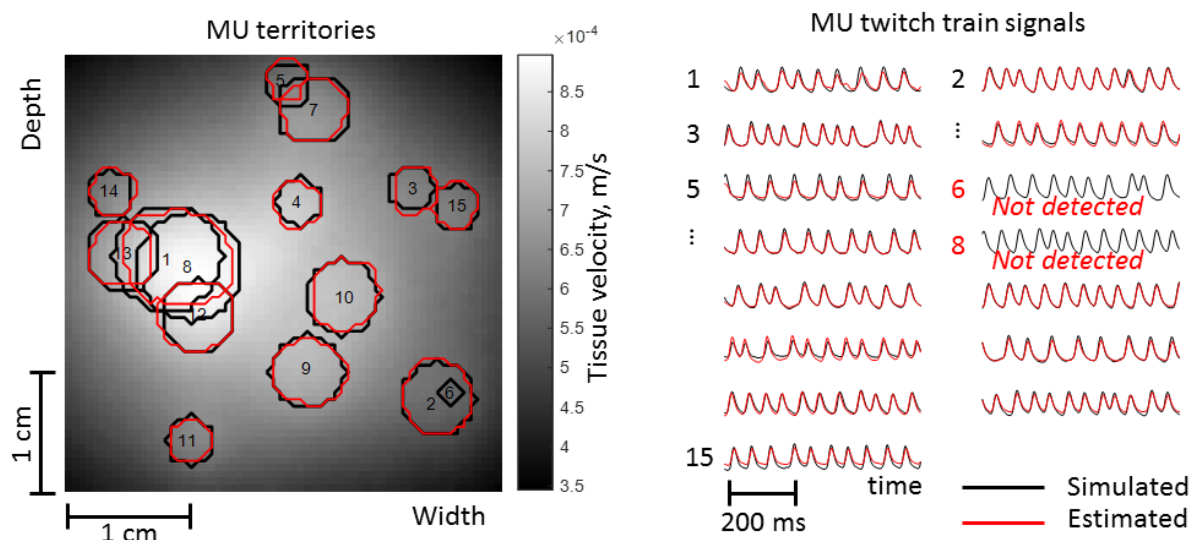


FIGURE 8. Example of results from a simulated tissue velocity image sequence with 15 active MUs. The left panel shows the simulated and segmented MU territories and an underlying map of the tissue velocity magnitude distribution within the image. The right panel shows the simulated and estimated twitch train signals of the corresponding MUs.


VI. DISCUSSION

In this work, a deep learning pipeline is suggested to identify MUs, segmentation of their territories, and estimate their twitch train activation based on ultrasound image sequence data of skeletal muscle contractions. The proposed model’s performance is evaluated using simulated ultrasound data mimicking the complex activation pattern of tens of activated MUs with overlapping territories and partially synchronized activation patterns. Performance evaluation shows that the proposed pipeline can effectively identify individual MUs,


and estimate their territories and twitch train signal at low contraction forces.

A. EVALUATION OF PERFORMANCE

First, the influence of including noisy signals in training was large. When models were trained on noise-free data, the performance was significantly lower for noisy signals as compared to noise-free signals (e.g., Figure 4 and Figure 6). When training the models with noisy signals, the performance was high independently of the noise level. These observations are



(a) Training loss.

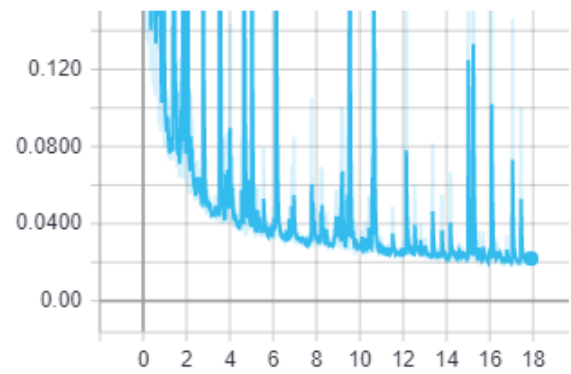


(b) Test loss.


FIGURE 9. Losses plotted over time (hours) when training the segmentation model. Ten mini-batches of validation data are computed for every 100 mini-batches of training data.

consistent with the results in the literature on the robustness of CNNs on noisy images. The addition of noise during the deep learning architecture training has been argued to have a regularization effect and, thus, gives robustness to the model [41], [42].

Second, the results showed that the detection recall decreased with the activation level. For $N=25$ units active, about 60% of the MUs could be detected. The recall and precision of the territories' segmentation were greater than 80%, and RoA was greater than 90% for all activation levels ($N=1 - 25$ MUs) when trained with noisy signals. Compared to the previous work by Rohlén *et al.* [10], who used the same simulation method and evaluation metrics, the detection performance also decreased with increasing activation level. However, at 25 MUs, their method achieved a higher recall of 75%. The proposed deep learning pipeline consistently outperformed Rohlén *et al.* [10] regarding the territory segmentation and firing pattern estimation, which had a recall in the range of 50 to 60% and a declining pattern for RoA vs. activation level, achieving 60% for 25 active MUs. In summary, we observe that the deep learning pipeline has better performance in terms of the segmentation and twitch-train estimation. However, the detection task needs further improvement.



(a) Training loss.



(b) Validation loss.

FIGURE 10. Losses plotted over time (hours) when training the signal estimation model. Ten mini-batches of validation data is computed for every 100 mini-batches of training data.

B. LIMITATIONS

The deep learning pipeline was trained using simulated muscle contraction data [10]. In this context, its performance should be interpreted as a proof-of-concept and demonstration of the principle. A limitation of this approach is that the method has learned the simulation data features compared to the previously suggested approach using blind source separation [10], which does not rely on learning the data is a more generalized approach. Consequently, we do not expect that the present trained network should have necessarily high performance on experimental data. There are three key simplifications/differences of the simulated data compared to experimental data: 1) superposition of mechanical responses of multiple MUs, 2) no extracellular matrix (fascia) connecting the fibers is included, and 3) non-physiological noise is additive and white. The first assumption has been shown valid at low force contractions where only a limited number of MUs are active as in our case [10], [43], [44]. The second assumption is a simplification of the true anatomy. Still, it has been indicated that in the cross-sectional view of the muscle, the spatio-temporal pattern is highly similar, comparing image sequences of simulated and experimental muscle contractions [10]. The impact on the third assumption's training was clear in the results and shows that a relevant noise model will be influential on performance.




FIGURE 11. Illustration of an RoA computation. The red marks correspond to the predicted signal, and the black marks correspond to the simulated ground truth signal. The dashed curves are the original signals, and the solid curves are smoothed versions of those. Vertical lines are the estimated firings, and horizontal lines are the intervals within which firings are considered to match. In this example $c = 7, A = 3$ and $B = 2$ resulting in $RoA = 0.583$.

Tissue velocity image sequence, 5 MUs




FIGURE 12. Example of a simulated tissue velocity image sequence with five active MUs. Approximately one cycle of the contraction of the units can be seen. White color means contraction and dark color relaxation.

A potential solution to translate the pipeline to an experimental application is to use the trained model from this work and transfer learning. This translation can be done by fine-tuning the model with labeled experimental data from gold-standard measurements of invasive EMG methods (needle-EMG) that can record the activation of individual MUs [45].

From the computational perspective, implementing the two key modules is based on a number of deep learning architectures, essentially requiring GPUs to perform training and inference. So, it would be useful to build methods with lower computational complexity while attaining a comparable performance. We can certainly hope that given the rising popularity of deep learning methods for medical imaging in




FIGURE 13. Example of a simulated tissue velocity image sequence with 15 active MUs. Approximately two cycles of the contraction of the units can be seen. White color means contraction and dark color relaxation.

general, developing more computationally efficient methods for motor unit detection and twitch train estimation is only a matter of time.

C. APPLICATIONS

The ability to identify the activity of MUs in the whole muscle (large field of view) would allow larger accessibility than current EMG methods that suffer from a small field of view. The proposed technique has many interesting and important applications, given a successful translation training the pipeline on experimental data. For example, for recording the neural firing patterns of MUs to control prostheses [46], studying strategies of the central nervous system on MU recruitment in endurance/fatiguing tasks [47], or clinical diagnostics when territories and/or firing pattern are altered due to pathological processes e.g., [48]. Altogether, the proposed method could allow the study of various questions that previously were difficult or not possible to address.

VII. CONCLUSION

In this work, a deep learning pipeline is suggested to identify the mechanical response of individual MUs, segmentation of their MU territories, and estimate their twitch train activations based on ultrasound image sequence data in voluntary skeletal muscle contractions. The results show that the proposed pipeline can effectively identify individual MUs, and estimate their territories and twitch train signal at low contraction forces. The proposed method is potentially useful to progress with experimental data. The ability of an ultrasound imaging based non-invasive large field of view of the active MUs would make it possible to address a variety of questions that were difficult to address before.

ACKNOWLEDGMENT

The authors thank M. D. Lars-Johan Liedholm at Dept of Neurophysiology, Västerbotten County Council, Umeå, Sweden, for useful discussions in the initial phase of the work.

APPENDIX A LOSSES FOR SEGMENTATION MODEL

Losses for the segmentation model on the training and validation sets are shown in Figure 9.

APPENDIX B LOSSES FOR TIME SIGNAL ESTIMATION MODEL

Losses for the twitch train estimation model on the training and validation sets are shown in Figure 10.

APPENDIX C RoA COMPUTATION

The computation of rate of agreement is illustrated in Figure 11.

APPENDIX D EXAMPLES OF SIMULATED IMAGE SEQUENCES

Examples of simulated tissue velocity image sequences for 5 MUs and 15 MUs are shown in Figure 12 and Figure 13, respectively.

REFERENCES

- [1] D. G. Sale, "Influence of exercise and training on motor unit activation," *Exercise sport Sci. Rev.*, vol. 15, pp. 95–151, 1987.
- [2] R. D. Henderson and P. A. McCombe, "Assessment of motor units in neuromuscular disease," *Neurotherapeutics*, vol. 14, no. 1, pp. 69–77, Jan. 2017.
- [3] J. V. Basmajian and C. J. De Luca, *Muscles Alive: Their Functions Revealed By Electromyography*, vol. 5. Baltimore, MA, USA: Williams & Wilkins, 1985.
- [4] M. Lindkvist, G. Granåsen, and C. Grönlund, "Precontractile optical response during excitation-contraction in human muscle revealed by non-invasive high-speed spatiotemporal NIR measurement," *Sci. Rep.*, vol. 8, no. 1, pp. 1–9, Dec. 2018.
- [5] J. C. King, D. Dumitru, and S. Nandedkar, "Concentric and single fiber electrode spatial recording characteristics," *Muscle Nerve, Off. J. Amer. Assoc. Electrodiagnostic Med.*, vol. 20, no. 12, pp. 1525–1533, Dec. 1997.
- [6] E. Martinez-Valdes, C. M. Laine, D. Falla, F. Mayer, and D. Farina, "High-density surface electromyography provides reliable estimates of motor unit behavior," *Clin. Neurophysiol.*, vol. 127, no. 6, pp. 2534–2541, Jun. 2016.
- [7] D. Farina and A. Holobar, "Characterization of human motor units from surface EMG decomposition," *Proc. IEEE*, vol. 104, no. 2, pp. 353–373, Feb. 2016.

- [8] A. J. Fuglevand, D. A. Winter, A. E. Patla, and D. Stashuk, "Detection of motor unit action potentials with surface electrodes: Influence of electrode size and spacing," *Biol. Cybern.*, vol. 67, no. 2, pp. 143–153, Jun. 1992.
- [9] T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*. New York, NY, USA: Academic, 2004.
- [10] R. Rohlén, E. Stålberg, K.-H. Stöverud, J. Yu, and C. Grönlund, "A method for identification of mechanical response of motor units in skeletal muscle voluntary contractions using ultrafast ultrasound imaging—Simulations and experimental tests," *IEEE Access*, vol. 8, pp. 50299–50311, 2020.
- [11] J. Bercoff, "Ultrafast ultrasound imaging," in *Ultrasound Imaging—Medical Applications*. Rijeka, Croatia: InTech Open, 2011, pp. 3–24, doi: 10.5772/19729.
- [12] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghahfarooian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.
- [13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [14] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Netw.*, vol. 61, pp. 85–117, Jan. 2015.
- [15] I. G. Y. Bengio and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2017, ch. 6, pp. 163–187.
- [16] A. Garcia-Garcia, S. Orts-Escobedo, S. Oprea, V. Villena-Martinez, P. Martinez-Gonzalez, and J. Garcia-Rodriguez, "A survey on deep learning techniques for image and video semantic segmentation," *Appl. Soft Comput.*, vol. 70, pp. 41–65, Sep. 2018.
- [17] A. Graves, A.-R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2013, pp. 6645–6649.
- [18] T. Iqbal and H. Ali, "Generative adversarial network for medical images (MI-GAN)," *J. Med. Syst.*, vol. 42, no. 11, p. 231, Nov. 2018.
- [19] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 580–587.
- [20] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1440–1448.
- [21] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, Jun. 2017.
- [22] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
- [23] W. Yin, K. Kann, M. Yu, and H. Schütze, "Comparative study of CNN and RNN for natural language processing," 2017, *arXiv:1702.01923*. [Online]. Available: <http://arxiv.org/abs/1702.01923>
- [24] U. Ugurlu, I. Oksuz, and O. Tas, "Electricity price forecasting using recurrent neural networks," *Energies*, vol. 11, no. 5, p. 1255, May 2018.
- [25] B. Wang, "Disconnected recurrent neural networks for text categorization," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics, Long Papers*, vol. 1, 2018, pp. 2311–2320.
- [26] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A search space odyssey," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 10, pp. 2222–2232, Oct. 2017.
- [27] Y. Yu, X. Si, C. Hu, and J. Zhang, "A review of recurrent neural networks: LSTM cells and network architectures," *Neural Comput.*, vol. 31, no. 7, pp. 1235–1270, Jul. 2019.
- [28] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, "Large-scale video classification with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1725–1732.
- [29] W. Abdulla. (2017). *Mask R-CNN for Object Detection and Instance Segmentation on Keras and TensorFlow*. [Online]. Available: https://github.com/matterport/Mask_RCNN
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [31] T. Y. Lin et al., "Microsoft COCO: Common objects in context," in *Computer Vision—ECCV (Lecture Notes in Computer Science)*, vol. 8693, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham, Switzerland: Springer, 2014, doi: 10.1007/978-3-319-10602-1_48.
- [32] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, in Proceedings of Machine Learning Research, vol. 9, Y. W. Teh and M. Titterton, Eds. Sardinia, Italy: Chia Laguna Resort, May 2010, pp. 249–256.
- [33] E. Stålberg and P. Dioszeghy, "Scanning EMG in normal muscle and in neuromuscular disorders," *Electroencephalogr. Clin. Neurophysiol./Evoked Potentials Sect.*, vol. 81, no. 6, pp. 403–416, 1991.
- [34] F. Negro, S. Muceli, A. M. Castronovo, A. Holobar, and D. Farina, "Multi-channel intramuscular and surface EMG decomposition by convolutive blind source separation," *J. Neural Eng.*, vol. 13, no. 2, Feb. 2016, Art. no. 026027.
- [35] D. Farina and R. Merletti, "A novel approach for precise simulation of the EMG signal detected by surface electrodes," *IEEE Trans. Biomed. Eng.*, vol. 48, no. 6, pp. 637–646, Jun. 2001.
- [36] T. Deffieux, J.-L. Gennisson, M. Tanter, and M. Fink, "Assessment of the mechanical properties of the musculoskeletal system using 2-D and 3-D very high frame rate ultrasound," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 55, no. 10, pp. 2177–2190, Oct. 2008.
- [37] C. Grönlund, K. Claesson, and A. Holtermann, "Imaging two-dimensional mechanical waves of skeletal muscle contraction," *Ultrasound Med. Biol.*, vol. 39, no. 2, pp. 360–369, Feb. 2013.
- [38] M. S. Stock and B. J. Thompson, "Motor unit interpulse intervals during high force contractions," *Motor Control*, vol. 20, no. 1, pp. 70–86, Jan. 2016.
- [39] W. Yao, R. J. Fuglevand, and R. M. Enoka, "Motor-unit synchronization increases EMG amplitude and decreases force steadiness of simulated contractions," *J. Neurophysiol.*, vol. 83, no. 1, pp. 441–452, Jan. 2000.
- [40] P. A. Kirkwood and T. A. Sears, "Cross-correlation analyses of motoneuron inputs in a coordinated motor act," in *Neuronal Cooperativity (Springer Series in Synergetics)*, vol. 49. Berlin, Germany: Springer, 1991, doi: 10.1007/978-3-642-84301-3_10.
- [41] G. An, "The effects of adding noise during backpropagation training on a generalization performance," *Neural Comput.*, vol. 8, no. 3, pp. 643–674, Apr. 1996.
- [42] C. M. Bishop, "Training with noise is equivalent to Tikhonov regularization," *Neural Comput.*, vol. 7, no. 1, pp. 108–116, Jan. 1995.
- [43] Y. Yoshitake, M. Shinohara, H. Ue, and T. Moritani, "Characteristics of surface mechanomyogram are dependent on development of fusion of motor units in humans," *J. Appl. Physiol.*, vol. 93, no. 5, pp. 1744–1752, Nov. 2002.
- [44] C. Cescon, P. Madeleine, and D. Farina, "Longitudinal and transverse propagation of surface mechanomyographic waves generated by single motor unit activity," *Med. Biol. Eng. Comput.*, vol. 46, no. 9, p. 871, 2008.
- [45] D. C. Preston and B. E. Shapiro, *Electromyography and Neuromuscular Disorders E-Book: Clinical-Electrophysiologic Correlations (Expert Consult-Online)*. Amsterdam, The Netherlands: Elsevier, 2012.
- [46] D. Farina, I. Vujaklija, M. Sartori, T. Kapelner, F. Negro, N. Jiang, K. Bergmeister, A. Andalib, J. Principe, and O. C. Aszmann, "Man/machine interface based on the discharge timings of spinal motor neurons after targeted muscle reinnervation," *Nature Biomed. Eng.*, vol. 1, no. 2, pp. 1–12, Feb. 2017.
- [47] A. Holtermann, C. Grönlund, J. S. Karlsson, and K. Roeleveld, "Differential activation of regions within the biceps Brachii muscle during fatigue," *Acta Physiologica*, vol. 192, no. 4, pp. 559–567, Apr. 2008.
- [48] E. Stålberg, "Use of single fiber EMG and macro EMG in study of reinnervation," *Muscle Nerve, Off. J. Amer. Assoc. Electrodiagnostic Med.*, vol. 13, no. 9, pp. 804–813, Sep. 1990.



HAZRAT ALI received the B.Sc. and M.Sc. degrees in electrical engineering in 2009 and 2012, respectively, and the Ph.D. degree from the University of Science and Technology Beijing, China, in 2015. He is currently a Researcher with Umea University, Sweden. He has authored/coauthored more than 50 journal/conference papers and book chapters. His research interests include unsupervised learning, generative and discriminative approaches, and speech and medical image processing. He is selected as a Young Researcher at the 5th Heidelberg Laureate Forum, Heidelberg, Germany. He was a recipient of the HEC Scholarship, the Top10 Research Pitch Award, the Kempe Foundation Sweden grant, the IEEE Student Travel Award, the IBRO grant, the TERENA/CISCO Travel grant, the QCRI/Boeing Travel grant, and the Erasmus Mundus STRoNGTiES research grant. He is an Associate Editor of IEEE and serves as a Reviewer for the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, the IEEE TRANSACTIONS ON MEDICAL IMAGING, the IEEE TRANSACTIONS ON EMERGING TOPICS IN COMPUTATIONAL INTELLIGENCE, *IET Signal Processing*, and many other international publishers.



JOHANNES UMANDER was born in Umeå, Sweden, in 1995. He received the master's degree in computer science from Umeå University, in 2020. His research interests include programming, machine learning, medical imaging, and generally reading about new technologies.



ROBIN ROHLÉN was born in Umeå, Sweden, in 1990. He received the master's degree in mathematical statistics from Umeå University, in 2016, and the master's degree in finance from the University of Gothenburg, in 2018. He is currently pursuing the Ph.D. degree in biomedical engineering with the Department of Radiation Sciences, Umeå University, under the supervision of Christer Grönlund. His current research interests include ultrasound imaging of muscle tissue and statistical modeling in space and time.



CHRISTER GRÖNLUND received the Ph.D. degree from Umea University, Sweden, in 2007. Since 2012, he has been an Associate Professor in biomedical engineering with the Department of Radiation Sciences, Umea University. He is also the Principal Investigator of a Research Group working on method development for electrophysiological, bio-optical, and ultrasound imaging. His medical research interests include vascular assessment for cardiovascular disease risk prediction, and neuromuscular diseases and disorders for diagnostics and sports medicine.

...