

Type of the Paper (Article)

Estimating Full Regional Skeletal Muscle Fibre Curvature from b-Mode Ultrasound Images Using Convolutional-Deconvolutional Neural Networks

Ryan Cunningham ^{1,*}, María B Sánchez ¹, Greg May ¹ and Ian Loram ^{1,*}

¹ School of Healthcare Science, Manchester Metropolitan University, Manchester, England, UK

* Correspondence: ryan.cunningham@mmu.ac.uk, i.loram@mmu.ac.uk

Abstract: Direct measurement of strain within muscle is important for understanding muscle function in health and disease. Current technology (kinematics, dynamometry, electromyography) provides limited ability to measure strain within muscle. Regional fiber orientation and length are related with active/passive strain within muscle. Currently, ultrasound imaging provides the only non-invasive means of observing regional fiber orientation within muscle during dynamic tasks. Previous attempts to automatically estimate fiber orientation from ultrasound are not adequate, often requiring manual region selection, feature engineering, providing low-resolution estimations (one angle per muscle), and deep muscles are often not attempted. Here, we propose deconvolutional neural networks (DCNN) for estimating fiber orientation at the pixel-level. Dynamic ultrasound images sequences of the calf muscles were acquired (25 Hz) from 8 healthy volunteers (4 male, ages: 25–36, median 30). A combination of expert annotation and interpolation/extrapolation provided labels of regional fiber orientation for each image. We then trained DCNNs both with and without dropout using leave one out cross-validation. Our results demonstrated robust estimation of regional fiber orientation with approximately 3° error, which was an improvement on previous methods. The methods presented here provide new potential to study muscle in disease and health.

Keywords: ultrasound; b-mode; skeletal muscle; fascicle orientation; pennation angle; fiber orientation; fiber tract; fascicle tract; convolutional neural network; deconvolutional neural network

1. Introduction

This paper presents a novel application of deconvolutional neural networks for estimating full regional skeletal muscle fiber orientation directly from standard frame rate (25Hz) b-mode ultrasound images. In recent years, ultrasound has become a valuable and ubiquitous clinical and research tool for understanding the changes which take place within muscle in ageing, disease, atrophy, and exercise. Ultrasound has been proposed [1] as a non-invasive alternative to intramuscular electromyography (iEMG) for measuring twitch frequency, useful for early the diagnosis of motor neuron disease (MND). Ultrasound has also recently demonstrated application to rehabilitative biofeedback [2]. Other computational techniques have been developed for muscle-ultrasound analysis which would allow estimation of changes in muscle length during contraction[3], and changes in fiber orientation and length [4]–[8]. Regional muscle fiber orientation (curvature) and length change when muscle is under active (generated internally through contraction) and/or passive strain (generated externally through joint movements or external pressure) [9], [10]. Muscle fiber orientation is one of the main identifying features of muscle state [11]. Below, we review the previous computational methods which attempt to automatically estimate fascicle orientation from ultrasound

images of skeletal muscle. We do not consider feature tracking based approaches [8], [12], since within ultrasound feature tracking is an ill-posed problem, leading to tracking drift and feature dropout due low signal to noise ration and out of plane feature motion [13].

Previous attempts to automatically estimate fiber orientation directly from b-mode ultrasound images are inadequate, typically requiring manual region identification, *a priori* feature engineering, presumptuous problem approach (e.g. fascicles and connective tissues appear as straight lines), and/or typically providing low resolution estimates (i.e. one angle for an entire muscle region). Further, to our knowledge, nobody has attempted the deep (more challenging) muscles. All previous methods depend on some parameterization of features such as Gabor wavelets [4], [5], and/or edge detectors and vessel enhancement filters [4]–[8], [14]–[16]. In all cases, parameters are empirically chosen and/or are based on assumptions about how the descriptive features will present within an image. In the studies we are aware of, parameters are chosen presumptuously or empirically rather than tuned with a multiple-fold cross-validation set, and since the authors unanimously write that results are sensitive to parameter changes, the assumption is that real world performance may not be as optimistic as reported.

Zhou and colleagues [17] developed a method based on the Revoting Hough Transform (RVHT) which provides an estimate of the overall fiber orientation in a single muscle, the gastrocnemius in the calf. Based upon the incorrect assumption that the fiber paths are straight lines, they use the RVHT to detect an empirically predetermined number of lines within a manually defined region of interest and then take the median orientation of all detected lines as the overall orientation of the fibers. The approach of Zhou and colleagues is fundamentally limited to detecting straight lines, whereas muscle fibers do not always present that way. When observing muscle fibers using ultrasound, there are many other features which can appear as straight lines (blood vessels, noise/dropout, artefacts/reflections, skin/fat layers, connective tissues, etc...). These fundamental facts are a strong limiting factor to the potential of methods such as the RVHT and the Radon Transform [7].

Rana and colleagues [5] introduced a method which potentially allowed estimation of local orientation, although after suggesting this, the mean of local orientations is used to produce another method which provides an estimate of the overall fiber orientation. Local orientations were identified by convolving a bank of Gabor Wavelet filters with a processed version of the image. They compare their proposed method with manual annotations from 10 experts and the Radon Transform, which like the Hough Transform (or RVHT) is a method which can be used for detecting straight lines in images. Both the Radon Transform and Wavelet methods require a preprocessing step described as vessel enhancement filtering, which enhances anisotropic features within the image. The vessel enhancement method is parametric, where parameters were chosen empirically with no cross-validation or evaluation of results over varying parameters. Their results showed that the Radon Transform was not significantly different to the expert annotations, whereas the Wavelet method was significantly different. Other than significance values, the only accessible results they report are mean differences (1.41°) between the Wavelet method and the expert annotations of real images, and mean differences between both the Radon and Wavelet method applied to synthetic images with known orientations ($< 0.06^\circ$). Although the Wavelet method performed comparatively poorly, the authors rightly suggest that this approach has potential to allow tracking of fiber paths throughout the muscle, whereas Radon and Hough Transform methods do not.

Namburete and colleagues [6] expanded upon the methods of Rana and colleagues [5], developing the hypothesis that regional fiber orientation can be tracked using local orientated Gabor Wavelets and vessel enhancement filters, for the first time going beyond linear approximations to the overall orientation. The proposed extension is to convolve the local orientations with a median filter, which effectively smooths the local orientations. The fiber trajectories are then tracked between the muscle boundaries on a continuous coordinate system with linear (following the dominant orientation in a local 15×15 -pixel region centered on the current fiber track) steps of 15 pixels. Finally, curvature is quantified using the Frenet-Serret curvature formula [18]. Namburete's method differs significantly from all preceding methods because it provides an estimation of local fiber orientation, rather than a more global estimate. However, the authors do not provide an estimate of the error for example when comparing to expert annotations; instead they evaluate errors on synthetic images with known orientations, which are trivial. Further, they do not apply this method to the deeper muscles, which are even more challenging.

Several important problems have been identified as unaddressed from a review of previous methods for estimating full regional fiber orientation in multiple layers of muscle; most commonly, the lack of error evaluation on real data, estimating overall fiber orientation and not local orientation, limited application to superficial muscles only, and presumptuous justification for chosen parameters. We propose to address these problems by introducing advanced machine learning methods and cross-validation against expert annotation of real data. In recent years the popularity of machine learning (in particular neural networks) has surged since a number of successive algorithmic, methodological, and computational hardware developments were introduced [19]–[23]. The application of machine learning to estimating local fiber orientations has thus far not been considered. In recent years, the development of deconvolutional neural networks (DCNN) allows robust regression (heatmaps [24]–[26]) or classification (semantic segmentation [27]) of pixel level labels in full resolution images, without any feature engineering or preprocessing (i.e. all parameters are learned from data, and cross-validated against held-out validation and test sets). Therefore, we expect DCNNs to perform well at predicting fiber orientation heat maps from ultrasound images of muscle. For comparison we implement a version of the wavelet method described by Namburete, and colleagues [6], since no data are presented on the accuracy of their method on real ultrasound images.

2. Materials and Methods

2.1. Deep Learning Software

All neural network software was developed from scratch by the authors using C/C++ and CUDA-C (Nvidia Corporation, Santa Clara, California). No libraries other than the standard CUDA libraries (runtime version 8.0 `cuda.h`, `cuda_runtime.h`, `curand.h`, `curand_kernel.h`, `cuda_occupancy.h`, and `device_functions.h`) and the C++ 11 standard library were used. All DCNNs were trained on an Intel Xeon CPU E5-2697 v3 (2.60GHz), 64GB (2133 MHz), Nvidia with a TitanX (Maxwell) GPU.

2.2. Data Acquisition

Ultrasound image data were recorded at 25Hz (AlokaSSD-5000 PHD, 7.5MHz) from the calf muscles (medial gastrocnemius and soleus) of 8 healthy volunteers (4 male, ages: 25-36, median 30)

during dynamic maximum isometric voluntary contractions. Images represented a $6\text{cm} \times 5\text{cm}$ (horizontal \times vertical axes) cross-section of the calf muscles. Volunteers lay prone on a physio bed with their right ankle strapped to a immobilised pedal. Volunteers were asked to push their toes down against the pedal as hard as they could. A dynamometer (Cybex) recorded the torque at 100Hz at the ankle during the contraction. Matlab (Matlab, R2013a, The MathWorks Inc., Natick, MA) was used to acquire ultrasound frames and a hardware trigger was used to initiate recording at the start of each trial.

2.3. Generating Ground Truth

Following data collection, we extracted all frames beginning one frame before contraction started, and ending with the frame when contraction peaked, where initial and peak contraction were identified manually. This resulted in a total of 504 images containing spatiotemporal variations in both fiber orientation and muscle thickness. Two experts were asked to manually identify fiber paths in all muscles/compartments in all images by marking a series of multiple-point lines/curves (typically 20-30 per muscle/compartment). The same experts were asked to identify the boundaries of the medial gastrocnemius (superficial muscle) and the boundaries and internal compartments (where visible) of the Soleus (deep muscle) by marking multiple-points lines/curves around the boundaries and visible intramuscular compartments.

To create labels for each pixel, first a blank image (matrix of zeros with equal dimensions to the ultrasound image) was created. Then, the angles along each of the annotated fiber lines were calculated, and for each pixel under the lines the calculated angle was stored. The result was an image of mostly zeroes and a few angled lines, where the value of the non-zero pixels represents the local angle of the line intersecting each pixel. Then, within each muscle/compartment (defined by the experts), nearest neighbor interpolation was used to fill the gaps between lines, followed by nearest neighbor extrapolation to fill the muscle/compartment. Between muscles/compartments, nearest neighbor interpolation was also used to fill the gaps. All other pixels (outside muscle, e.g. skin) were set to zero. Following, additional data were created to introduce variation in the data/labels and help prevent over-fitting, by randomly (uniform) rotating each image and corresponding labels/angles between -5° and 5° , thus doubling the size dataset to 1008.




Figure 1. Image annotation. (a) shows the raw ultrasound image, (b) shows expert annotated line traces of visible fibre paths, (c) shows nearest neighbor region filling of annotated fiber paths within annotated muscles/compartments, (d) shows line traces over the filled regions. The colors in (b)-(d) represent the local fibre orientation in degrees (see colour bar in (c)).

2.4. Wavelet Filter Bank Method

The wavelet method was implemented faithful to the original paper [6]. A bank of 180 Gabor wavelet filters was generated spanning 0° - 179° in 1° increments using,

$$g(x, y) = -e^{\left(\frac{(x-k-1)^2(y-k-1)^2}{-dk}\right)} \cos\left(\frac{2\pi((x-k-1)\cos(\alpha)-(y-k-1)\sin(\alpha))}{f}\right),$$

where k is the kernel size ($k = 20$), d is a damping term ($d = 51.243$), α is the orientation, and f is the spatial frequency ($f = 7$). Initially, all parameters were set as originally described in [6], and after a first analysis, some parameters (f and k) were varied ($f = 9, f = 11, f = 13, k = f*3-1$) to give some scope to the potential of this method on real data. Vessel enhancement was performed using a Matlab version of the Frangi multiscale vessel enhancement filter [28] which can be found here [29]. Following [6], filtering was performed at 3 resolutions, 2, 3, and 4. Following a first analysis, a second analysis was done at resolutions, 4, 6, and 8, for all variants of the Gabor wavelet filters described above.

To analyse an image, the image was filtered with the vessel enhancement filter and then convolved with the entire wavelet filter bank. At each spatial location in the image, we take the α corresponding to the filter with the highest convolution in that location. The result is a map of local orientations. Following [6], the image was then convolved with a 35×35 median filter resulting in the final map of local fiber orientations.

2.5. Deconvolution Neural Network Method




Figure 2. DCNN schematic. On the far left the grayscale input image is connected to a series of convolutional (Conv.) and max-pooling (Pool.) layers, followed by a series of fully connected (FC.) layers, followed by a series of max-un-pooling (UnPool.) and deconvolutional (DeConv.). The green, blue and red arrows and numbers represent the vertical and horizontal axes, and the number of filters (or depth/color channels with respect to the input/output layers) respectively in each layer. This network has 1,942,528 units (neurons), and 9,187,904 trainable weights/biases (synapses).

Our primary concern when deciding on a DCNN architecture (number of layer and units per layer) was having a large enough model to learn the training set. Our secondary concern was training time, since we planned to execute an 8-fold cross-validation. Addressing the latter concern, both input images and labels were down-sampled (bilinear interpolation) to 128×128 pixels. Addressing the former concern, the DCNN implemented 32 convolutional filters in the input layer, with 4 additional convolutional layers, the first of which had 32 filters and each of the following had double the number of filters in the preceding convolutional layer. Between convolutional layers, the spatial dimensions were down-sampled (max-pooling) by a factor of 2. A dense fully connected layer of 1024 nodes was fully connected to an initial up-sampling (2×2 max-un-pooling) layer, followed by a deconvolution layer. An additional 4 up-sampling with deconvolutional layers following, completes

the network architecture, with the final layer being a full 128×128 resolution regression map. Every layer consisted of rectified linear units (ReLU), with the exception of the final (output) layer, which consisted of linear units. The DCNN was trained by minimizing mean square error (MSE) between the labels and the output layer using online stochastic gradient descent (batch size of 1 sample).


2.6. DCNN Cross-Validation and Regularization

To optimize and test generalization (estimated real-world performance), leave one person out cross-validation was performed on 3 DCNNs with different dropout parameter settings (0%, 30%, 50%), where dropout was enabled in the fully connected layers only. A small amount of L2 weight decay was used, $5e^{-4}$, to prevent the weights from growing too large and to help convergence. The weights of the DCNNs were initialized using Xavier initialization [30]. To train the DCNNs we used a learning rate of $5e^{-2}$ and weight gradient momentum of $9.5e^{-1}$, for a maximum of 300 full batch iterations (approx. 300,000 weight updates) using stochastic online gradient descent (minibatch=1). All learning parameters were chosen empirically on a small fraction of the training data. The DCNNs were trained 8 distinct times (one for each person), each time reinitializing the networks and leaving a different person out for testing, and randomly splitting the remaining data (7 people) into validation and training sets (10% and 90% respectively). Prior to training, each pixel of the input images was normalized over the training set to zero mean and unit variance, and each pixel of the output images was divided by a constant (45) so that most of the pixels (angles) fell between -1 and +1. Validation and test errors were observed periodically, and the test error observed at the lowest validation error (early stopping method) gave results for the validation set and the test set (estimated real-world generalization error) for the held-out person. Test results for all held-out participants were combined to reveal the performance of the optimal DCNN.

3. Results

Following model optimization of the wavelet and DCNN methods, both methods were compared against the ground truth labels using a range of error measures, thus making the work accessible and comparable to others. In sections 3.1 and 3.2 results are presented respectively for the wavelet method and the DCNN method. With respect to data ranges, within the GM muscle, the maximum range of angles within a single muscle region was 34.84° , whereas the maximum range of angles over all participants was 76.58° . Within the Sol muscle (compartment 1), the maximum range of angles within a single muscle region was 35.02° , whereas the maximum range of angles over all participants was 64.92° . Within the Sol muscle (compartment 2), the maximum range of angles within a single muscle region was 65.15° , whereas the maximum range of angles over all participants was 79.05° .

3.1. Wavelet Method



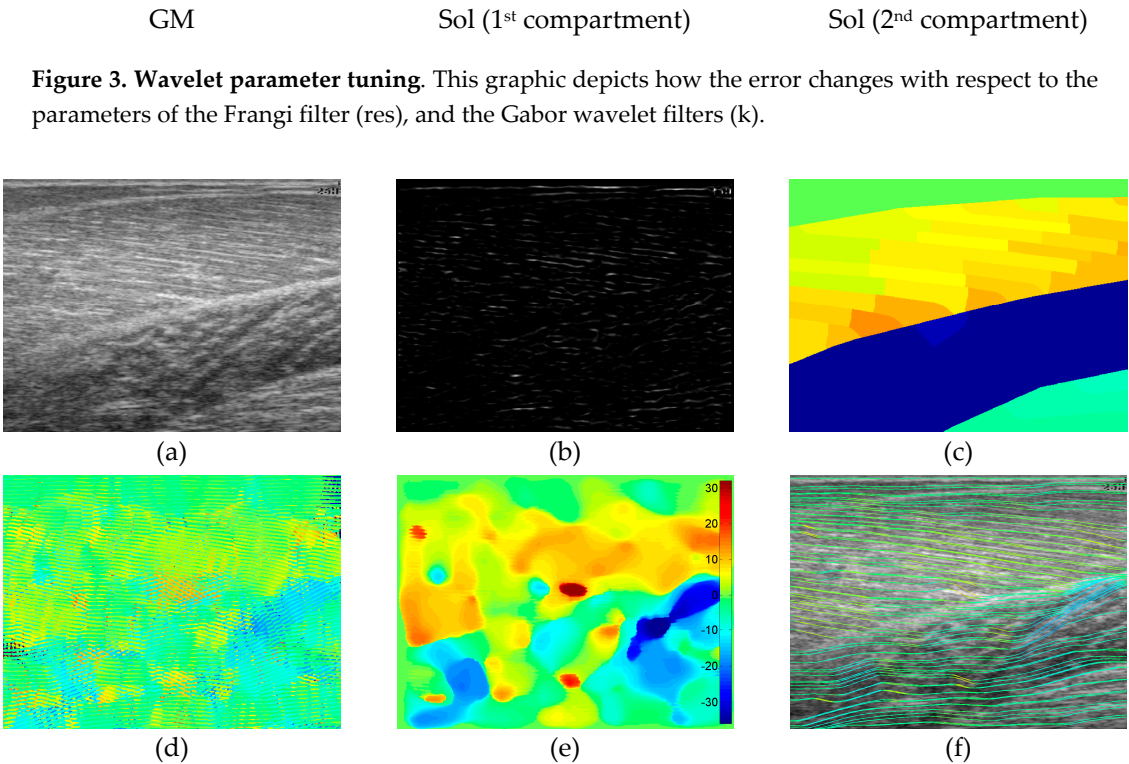


Figure 4. Representative wavelet visual result. (a) shows the raw image, (b) shows the vessel-enhanced image, (c) shows the ground truth, (d) shows the result of Gabor wavelet convolution with (b), (e) shows the result of a (35 × 35) median filter applied to (d), (f) shows line traces of (e). It is clear, when comparing (e) to (c), that the heatmaps partial match within the GM region, but almost not at all in the Sol region. The line trace visually confirms that the fibre traces are roughly aligned with the visible fibres in the raw image in the GM muscle, and the second compartment of the Sol muscle only.

The wavelet method was evaluated using the optimal parameters (within the range of parameters investigated) for each muscle/compartment, separately. To find the optimal parameters within each muscle, the local orientation over each pixel identified by the wavelet method was compared to each pixel of the ground truth by root mean square error (RMSE) (see fig. 3). Within the parameters we investigated, for the GM muscle, smaller wavelet filters resulted in smaller RMSE for the higher-resolution Frangi filter (res= [4, 6, 8]), conversely larger wavelet filters results in smaller RMSE for the lower-resolution Frangi filter. Within the parameters we investigated, the first compartment of the Sol muscle yields best results with the higher-resolution Frangi filter and larger wavelet filters. Conversely, within the parameters we investigated, the second compartment of the Sol muscle yields best results for smaller wavelet filters and lower-resolution Frangi filter. The models with the lowest RMSE were compared to ground truth using mean difference (MD) and mean absolute error (MAE), to be compared with the DCNN method (section 3.2).

For the optimal models within each muscle/compartment, the deviations of predicted local fibre orientations from ground truth were much too high to be useful in any real context, particularly in the Sol (compartment 1) muscle, which reports an optimal RMSE of over 26° (see fig. 3). Discrepancies were much lower in the GM and Sol (compartment 2) muscles at 11.3° and 7.1°, respectively. Visual comparisons were also made by comparing the wavelet analysis, presented as a heat-map, to the ground truth, also presented as a heat-map (see fig. 4). Visual comparison revealed quite sporadic agreement between wavelet predictions and ground truth. Visual comparison also revealed that any agreement between wavelet predictions and ground truth was not only heavily dependent on the

presence of well-defined muscle fibre tracts, but also on the ability of the Frangi filter to extract those features.

Application of the wavelet method was computationally expensive using the Matlab framework, with full regional analysis (whole image including all muscles) taking between one and two hours to process all 504 images for a single set of parameters. We note that this can be significantly improved upon with the use of graphics processing units (GPUs) and/or a full compiled language (such as C/C++).

3.2. DCNN Method




Figure 5. DCNN regularization. This graphic depicts how the error changes with respect to the percentage of dropout enabled in the fully connected layers of the DCNN. These results show that dropout improved lowered RMSE, with 50% dropout being most effective, and therefore revealing the optimal model.

The DCNN method was evaluated using the optimal models as identified by cross-validation (see sec. 2.6) and regularization using dropout (see fig. 5). Following, the output heat-map from the optimal DCNN for all 504 images were compared to the ground truth using mean difference (MD) and mean absolute error (MAE), to be compared with the wavelet method (see sec. 3.1). Visual comparisons were also made by comparing the DCNN predictions, presented as a heat-map, to the ground truth, also presented as a heat-map (see fig. 5). The most immediate result of visual comparisons was the ability of the DCNN to ‘color-code’ the correct regions; the GM and Sol (compartment 1) muscles typically present with opposing fibre orientations, with negative angles in the Sol muscle and positive angles in the GM muscle. In almost all images the DCNN was able to identify the Sol and GM muscles simply by predicting negative and positive angles respectively for each compartment. Furthermore, this was broadly true for entire muscle regions, even where there were no visually identifiable fibre tracts (see fig. 5).




Figure 5. Representative DCNN visual result. (a) shows the raw ultrasound image, (b) shows the ground truth, (c) shows the raw output from the DCNN, (d) shows line traces of (c). Careful observation of (c) reveals that the DCNN has predicted a bifurcation in the deep part of the Sol muscle

(lower right), which disagrees with the expert in (b). Now look at the lower right portion of (a) and see that this prediction is entirely plausible, as confirmed by (d) which shows line traces following the visible fiber paths.

Table 1. Table of results per muscle/compartiment. This table shows results computed by comparing predictions of each method, over all pixels ($128 \times 128 = 16384$), to the ground truth.

Muscle/ Compartment	Error Measure	Namburete, <i>et al</i> [6]	Proposed (DCNN)
GM	MD ¹	$7.57^\circ \pm 8.37^\circ$	$3.09^\circ \pm 8.74^\circ$
	MAE ²	$8.23^\circ \pm 7.72^\circ$	$6.24^\circ \pm 6.85^\circ$
	RMSE ³	11.29°	9.27°
Sol (1st Compartment)	MD ¹	$-22.84^\circ \pm 14.01^\circ$	$-4.99^\circ \pm 12.09^\circ$
	MAE ²	$23.09^\circ \pm 13.61^\circ$	$10.64^\circ \pm 7.61^\circ$
	RMSE ³	26.80°	13.08°
Sol (2nd Compartment)	MD ¹	$1.61^\circ \pm 6.92^\circ$	$10.64^\circ \pm 9.93^\circ$
	MAE ²	$5.03^\circ \pm 5.01^\circ$	$12.25^\circ \pm 7.86^\circ$
	RMSE ³	7.11°	14.56°

¹ mean difference. ² mean absolute error. ³ root mean square error.

4. Discussion

In this paper we reviewed the state of the art computational methods for estimating regional fibre orientation from ultrasound images of skeletal muscle. We found that the previous approaches were fundamentally lacking in their comparison to real ultrasound data, thus rendering them incomparable unless implemented. We also found that previous methods were targeted at estimating global/overall muscle fibre orientation by using some line detection technique such as the Hough or Radon transform [5], [7], [14], [17]. Fibres are known to curve within a muscle, and here we have also provided data in support of this, based on expert identification of fibre tracts within 3 muscles/compartments over 8 people, which show how fibre orientation is different within a muscle/compartiment (see paragraph 1 of Results section).

In review of previous methods we found an approach which demonstrated potential to provide estimates of local orientation [5], [6]. The authors did not present any comparable data to assess the performance of this method on real ultrasound data, only reporting that there were significant differences between their method and expert annotations of fibre orientation in real ultrasound images. The authors also highlight weaknesses in their approach such as the need for manual region identification, where the region selected should contain visibly well-defined fibre tracts. We recognized the potential for this method to fail in the presence of noise, artefacts, other quasi-line structures (blood vessels, connective tissues, etc.), and other muscles where fibres do not present as clearly, such as the deep muscles.

We proposed the application of a deep learning method known as deconvolutional neural networks (DCNN). We identified that such a method is capable of robust regression (heatmaps [24]–[26]) or classification (semantic segmentation [27]) of pixel level labels in full resolution images. The advantage of a deep learning approach is that the model can learn to identify fibres in well-defined regions, interpolate between well-defined regions and extrapolate from regions to fill the muscle/compartiment with predictions of local fibre orientation. Deep learning models can also learn to ignore regions with spurious (non-fibrous) features. After training 3 DCNNs with different levels of dropout enabled, using 8-fold cross-validation and early stopping using a validation set, we compared the results to the wavelet method applied to the same data. Results were compared within specific muscle regions (identified by two experts) since large variations were expected due to the challenging textural appearance of the deep muscles (Sol). Results revealed that both methods showed large deviations from the ground truth, however the DCNN method gave much more accurate predictions for the superficial and deep muscles (first compartment) than the wavelet

method. Those predictions were also much more consistent as identified by standard deviations of the error metrics (see table. 1). Our results showed that dropout applied to fully connected layers, improved overall performance of the DCNN, and we suggest that further improvement is possible with additional tuning of the dropout parameter, or even the configuration of the network architecture, and certainly with additional data. In contrast, although figure 3 indicates that there is scope to tune the wavelet method, it is much harder to do so in such a way which generalizes and deals with all of the problems we identified as well as the original authors (image artefacts, spurious fibre features, noise, reflections, etc.).

5. Conclusions

In this paper we have provided the first comprehensive analysis of an existing and a novel computational method for estimating full regional fibre orientation from ultrasound images of human skeletal muscle. We have proposed a novel application of deep learning to a long-standing and challenging problem, and demonstrated state of the art results. We present analyses in a form which is comparable to any future developments, and we also publish our ultrasound and ground truth data to support this end. The application of DCNNs to this problem has opened up new potential to hi-resolution analysis of skeletal muscle, from prediction of motion maps to segmentation of muscles and other structures of interest. This paper provides further evidence that deep learning methods can surpass state of the art performance, even when there is not an abundance of labeled data available. With additional data we propose that this project could easily be extended successfully, and this preliminary muscle analysis step could very likely form part of a skeletal muscle analysis system which accurately predicts the passive and active muscle forces non-invasively directly from single ultrasound images and sequences. Such a contribution could enable early diagnoses of diseases such as MND, and would enable personalized musculoskeletal medical diagnosis, monitoring, treatment targeting, and care.

Acknowledgments: This research was funded by the Dalton Research Institute and the School of Healthcare Science, Manchester Metropolitan University.

Author Contributions: Loram and Cunningham conceived the work. Loram, Cunningham and May designed and performed the data acquisition experiment. Cunningham and Sánchez annotated muscle boundaries and muscle fibre traces, providing ground truth for all images. Cunningham wrote the paper. All authors revised versions of the paper. Cunningham wrote the wavelet analysis software in Matlab and performed the wavelet analysis. Cunningham wrote the deep learning software in C++/Cuda-C and performed the deep learning analysis. Loram was the supervisor of this work.

Conflicts of Interest: The authors declare no conflict of interest.

References

- [1] P. J. Harding, I. D. Loram, N. Combes, and E. F. Hodson-Tole, "Ultrasound-Based Detection of Fasciculations in Healthy and Diseased Muscles," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 3, pp. 512–518, 2016.
- [2] I. D. Loram, B. Bate, P. Harding, R. Cunningham, and A. Loram, "Proactive Selective Inhibition Targeted at the Neck Muscles: This Proximal Constraint Facilitates Learning and Regulates Global Control," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 4, pp. 357–369, 2017.
- [3] I. Loram, C. Maganaris, and M. Lakie, "Use of ultrasound to make noninvasive in vivo measurement of continuous changes in human muscle contractile length," *J. Appl. Physiol.*, pp. 1311–1323, 2006.
- [4] G. Q. Zhou, P. Chan, and Y. P. Zheng, "Automatic measurement of pennation angle and fascicle length of gastrocnemius muscles using real-time ultrasound imaging," *Ultrasonics*, vol. 57, no. C, pp. 72–83, 2015.
- [5] M. Rana, G. Hamarneh, and J. M. Wakeling, "Automated tracking of muscle fascicle orientation in B-mode ultrasound images," *J. Biomech.*, vol. 42, no. 13, pp. 2068–2073, 2009.

- [6] A. I. L. Namburete, M. Rana, and J. M. Wakeling, "Computational methods for quantifying in vivo muscle fascicle curvature from ultrasound images," *J. Biomech.*, vol. 44, no. 14, pp. 2538–2543, 2011.
- [7] X. Chen, Q. Li, S. Qi, H. Zhang, S. Chen, and T. Wang, "Continuous fascicle orientation measurement of medial gastrocnemius muscle in ultrasonography using frequency domain Radon transform," *Biomed. Signal Process. Control*, vol. 20, pp. 117–124, Jul. 2015.
- [8] J. Darby, B. Li, N. Costen, I. Loram, and E. Hodson-Tole, "Estimating skeletal muscle fascicle curvature from B-mode ultrasound image sequences," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 7, pp. 1935–1945, 2013.
- [9] R. D. Herbert, S. C. Gandevia, and R. D. Herbert, "Changes in pennation with joint angle and muscle torque: in vivo measurements in human brachialis muscle," *J. Physiol.*, vol. 484, no. Pt2, pp. 523–532, 1995.
- [10] M. V Narici, T. Binzoni, E. Hiltbrand, J. Fasel, F. Terrier, and P. Cerretelli, "In vivo human gastrocnemius architecture with changing joint angle at rest and during graded isometric contraction," *J. Physiol.*, vol. 496, no. 1, pp. 287–297, 1996.
- [11] R. L. Lieber, J. Fridé N, and J. Fridén, "Functional and Clinical Significance of Skeletal Muscle Architecture," *Muscle Nerve*, vol. 23, no. November, pp. 1647–1666, 2000.
- [12] J. Darby, E. F. Hodson-Tole, N. Costen, and I. D. Loram, "Automated regional analysis of B-mode ultrasound images of skeletal muscle movement," *J. Appl. Physiol.*, vol. 112, no. 2, pp. 313–327, 2012.
- [13] F. Yeung, S. F. Levinson, and K. J. Parker, "Multilevel and Motion Model-Based Ultrasonic Speckle Tracking Algorithms," *Ultrasound Med. Biol.*, vol. 24, no. 3, pp. 427–441, Mar. 1998.
- [14] Y. Zhou, J.-Z. Li, G. Zhou, and Y.-P. Zheng, "Dynamic measurement of pennation angle of gastrocnemius muscles during contractions based on ultrasound imaging," *Biomed. Eng. Online*, vol. 11, no. 1, p. 63, 2012.
- [15] G. Q. Zhou and Y. P. Zheng, "Automatic Fascicle Length Estimation on Muscle Ultrasound Images With an Orientation-Sensitive Segmentation," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 12, pp. 2828–2836, 2015.
- [16] F. Chen, "Hand gesture recognition using a real-time tracking method and hidden Markov models," *Image Vis. Comput.*, vol. 21, no. 8, pp. 745–758, 2003.
- [17] Y. Zhou and Y. P. Zheng, "Estimation of Muscle Fiber Orientation in Ultrasound Images Using Revoting Hough Transform (RVHT)," *Ultrasound Med. Biol.*, vol. 34, no. 9, pp. 1474–1481, Sep. 2008.
- [18] H. Stark and N. Schilling, "A novel method of studying fascicle architecture in relaxed and contracted muscles," *J. Biomech.*, vol. 43, no. 15, pp. 2897–2903, 2010.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Adv. Neural Inf. Process. Syst.* 25, pp. 1–9, 2012.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Arxiv.Org*, vol. 7, no. 3, pp. 171–180, 2015.
- [21] G. Hinton, "Dropout : A Simple Way to Prevent Neural Networks from Overfitting," vol. 15, pp. 1929–1958, 2014.
- [22] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," *Proc. 27th Int. Conf. Mach. Learn.*, no. 3, pp. 807–814, 2010.
- [23] L. D. Le Cun Jackel, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, B. Le Cun, J. Denker, and D. Henderson, "Handwritten Digit Recognition with a Back-Propagation Network," *Adv. Neural Inf. Process. Syst.*, pp. 396–404, 1990.
- [24] T. Pfister, J. Charles, and A. Zisserman, "Flowing convnets for human pose estimation in videos," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, vol. 2015 Inter, pp. 1913–1921.

- [25] P. Fischer, A. Dosovitskiy, E. Ilg, ..., and T. Brox, "FlowNet: Learning Optical Flow with Convolutional Networks," *Iccv*, p. 8, 2015.
- [26] C. Payer, D. Štern, H. Bischof, and M. Urschler, "Regressing heatmaps for multiple landmark localization using CNNs," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016, vol. 9901 LNCS, pp. 230–238.
- [27] H. Noh, S. Hong, and B. Han, "Learning deconvolution network for semantic segmentation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2016, vol. 11–18–Dece, pp. 1520–1528.
- [28] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. a Viergever, "Multiscale vessel enhancement filtering," *Medial Image Comput. Comput. Intervention - MICCAI'98. Lect. Notes Comput. Sci. vol 1496*, vol. 1496, pp. 130–137, 1998.
- [29] D.-J. Kroon, "Hessian based Frangi Vesselness filter," *MATLAB Central/File Exchange*, 2010. [Online]. Available: <https://uk.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-filter>. [Accessed: 01-Nov-2017].
- [30] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," *Proc. 13th Int. Conf. Artif. Intell. Stat.*, vol. 9, pp. 249–256, 2010.