

1 *Type of the Paper (Article)*

2 **Estimating Full Regional Skeletal Muscle Fibre** 3 **Curvature from b-Mode Ultrasound Images Using** 4 **Convolutional-Deconvolutional Neural Networks**

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

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

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

8 **Abstract:** Direct measurement of strain within muscle is important for understanding muscle
9 function in health and disease. Current technology (kinematics, dynamometry, electromyography)
10 provides limited ability to measure strain within muscle. Regional fiber orientation and length are
11 related with active/passive strain within muscle. Currently, ultrasound imaging provides the only
12 non-invasive means of observing regional fiber orientation within muscle during dynamic tasks.
13 Previous attempts to automatically estimate fiber orientation from ultrasound are not adequate,
14 often requiring manual region selection, feature engineering, providing low-resolution estimations
15 (one angle per muscle), and deep muscles are often not attempted. Here, we propose
16 deconvolutional neural networks (DCNN) for estimating fiber orientation at the pixel-level.
17 Dynamic ultrasound images sequences of the calf muscles were acquired (25Hz) from 8 healthy
18 volunteers (4 male, ages: 25–36, median 30). A combination of expert annotation and
19 interpolation/extrapolation provided labels of regional fiber orientation for each image. We then
20 trained DCNNs both with and without dropout using leave one out cross-validation. Our results
21 demonstrated robust estimation of regional fiber orientation to within 5° error, which was
22 comparable to previous methods. The methods presented here provide new potential to study
23 muscle in disease and health.

24 **Keywords:** ultrasound; b-mode; skeletal muscle; fascicle orientation; pennation angle; fiber
25 orientation; fiber tract; fascicle tract; convolutional neural network; deconvolutional neural network
26

27 **1. Introduction**

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

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

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

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

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

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

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

114 **2. Materials and Methods**

115 *2.1. Deep Learning Software*

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

122
123
124


125 2.2. Data Acquisition

126 Ultrasound image data were recorded at 25Hz (AlokaSSD-5000 PHD, 7.5MHz) from the calf
 127 muscles (medial gastrocnemius and soleus) of 8 healthy volunteers (4 male, ages: 25-36, median 30)
 128 during dynamic maximum isometric voluntary contractions. Images represented a 6cm × 5cm
 129 (horizontal × vertical axes) cross-section of the calf muscles. Volunteers lay prone on a physio bed
 130 with their right ankle strapped to a immobilised pedal. Volunteers were asked to push their toes
 131 down against the pedal as hard as they could. A dynamometer (Cybex) recorded the torque at 100Hz
 132 at the ankle during the contraction. Matlab (Matlab, R2013a, The MathWorks Inc., Natick, MA) was
 133 used to acquire ultrasound frames and a hardware trigger was used to initiate recording at the start
 134 of each trial.

135 2.3. Generating Ground Truth

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

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



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

160

161 2.4. Wavelet Filter Bank Method

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

$$164 \quad g(x, y) = -e^{\left(\frac{(x-k-1)^2(y-k-1)^2}{-dk}\right)} \cos\left(\frac{2\pi((x-k-1)\cos(\alpha)-(y-k-1)\sin(\alpha))}{f}\right),$$


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

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

178 2.5. Deconvolution Neural Network Method

179 Our primary concern when deciding on a DCNN architecture (number of layer and units per
180 layer) was having a large enough model to learn the training set. Our secondary concern was training
181 time, since we planned to execute an 8-fold cross-validation. Addressing the latter concern, both
182 input images and labels were down-sampled (bilinear interpolation) to 128×128 pixels. Addressing
183 the former concern, the DCNN implemented 16 convolutional filters in the input layer, with 4
184 additional convolutional layers, each with double the number of filters in the preceding
185 convolutional layer. Between convolutional layers, the spatial dimensions were down-sampled (max-
186 pooling) by a factor of 2. A dense fully connected layer of 512 nodes was fully connected to an initial
187 up-sampling (2×2 max-un-pooling) layer, followed by a deconvolution layer. An additional 4 up-
188 sampling with deconvolutional layers following, completes the network architecture, with the final
189 layer being a full 128×128 resolution regression map. Every layer consisted of rectified linear units
190 (ReLU), with the exception of the final (output) layer, which consisted of linear units. The DCNN was
191 trained by minimizing mean square error (MSE) between the labels and the output layer using online
192 stochastic gradient descent (batch size of 1 sample).

193



194 **Figure 2. DCNN schematic.** On the far left the grayscale input image is connected to a series of
 195 convolutional (Conv.) and max-pooling (Pool.) layers, followed by a series of fully connected (FC.)
 196 layers, followed by a series of max-un-pooling (UnPool.) and deconvolutional (DeConv.). The green,
 197 blue and red arrows and numbers represent the vertical and horizontal axes, and the number of filters
 198 (or depth/color channels with respect to the input/output layers) respectively in each layer. This
 199 network has 1,286,656 units (neurons), and 4,983,552 trainable weights/biases (synapses).

200 2.6. DCNN Cross-Validation and Regularization


201 To optimize and test generalization (estimated real-world performance), leave one person out
 202 cross-validation was performed on a DCNN with a small amount of L2 weight decay, $5e^{-4}$, to prevent
 203 the weights from growing too large. The weights of the DCNN were initialized using Xavier
 204 initialization [30]. To train the DCNN we used a learning rate of $2e^{-2}$ and weight gradient momentum
 205 of $9.5e^{-1}$, for a maximum of 200 full batch iterations (approx. 200,000 weight updates). All learning
 206 parameters were chosen empirically on a small fraction of the training data. The DCNN was trained
 207 8 distinct times (one for each person), each time reinitializing the network and leaving a different
 208 person out for testing, and randomly splitting the remaining data (7 people) into validation and
 209 training sets (10% and 90% respectively). Prior to training, each pixel of the input images were
 210 normalized over the training set to zero mean and unit variance, and each pixel of the output images
 211 were simply divided by a constant (45) so that each pixel fell between -1 and +1. Validation and test
 212 errors were observed periodically, and the test error recorded at the lowest validation error (early
 213 stopping method) gave results for the validation set and the test set (estimated real-world
 214 generalization error) for the held-out person. Test results for all held-out participants were combined
 215 to reveal the performance of the DCNN.

216 3. Results

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

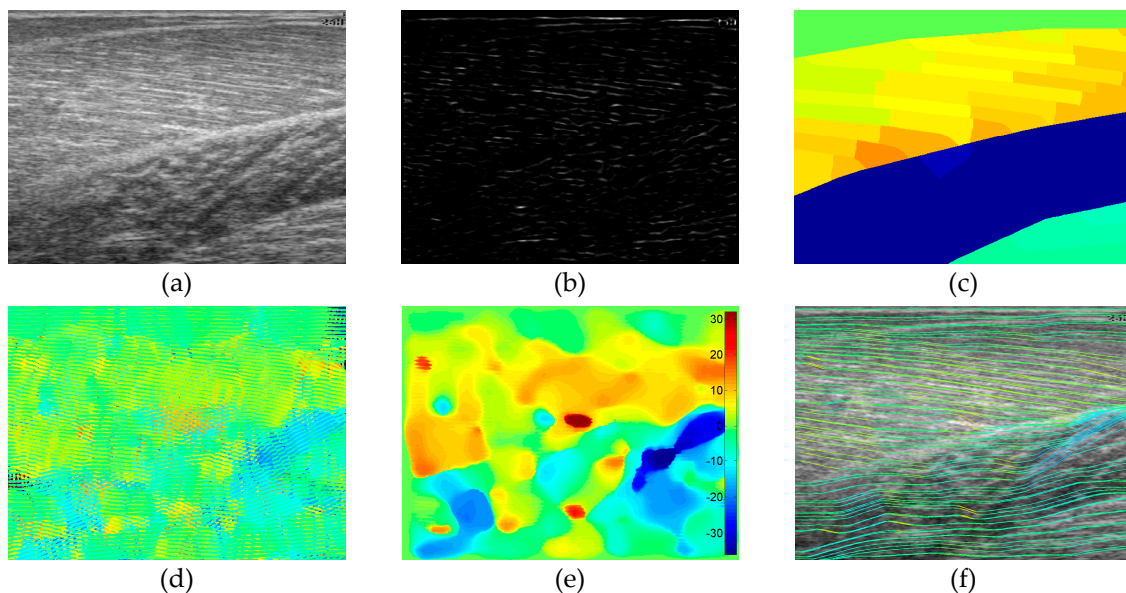
225 a single muscle region was 65.15° , whereas the maximum range of angles over all participants was
 226 79.05° .

227 3.1. Wavelet Method



228 **Figure 3. Wavelet parameter tuning.** This graphic depicts how the error changes with respect to the
 229 parameters of the Frangi filter (res), and the Gabor wavelet filters (k).

230



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

237

238 The wavelet method was evaluated using the optimal parameters (within the range of
 239 parameters investigated) for each muscle/compartment, separately. To find the optimal parameters
 240 within each muscle, the local orientation over each pixel identified by the wavelet method was
 241 compared to each pixel of the ground truth by root mean square error (RMSE) (see fig. 3). Within the
 242 parameters we investigated, for the GM muscle, smaller wavelet filters resulted in smaller RMSE for
 243 the higher-resolution Frangi filter ($res=[4, 6, 8]$), conversely larger wavelet filters results in smaller
 244 RMSE for the lower-resolution Frangi filter. Within the parameters we investigated, the first
 245 compartment of the Sol muscle yields best results with the higher-resolution Frangi filter and larger
 246 wavelet filters. Conversely, within the parameters we investigated, the second compartment of the

247 Sol muscle yields best results for smaller wavelet filters and lower-resolution Frangi filter. The
 248 models with the lowest RMSE were compared to ground truth using mean difference (MD) and mean
 249 absolute error (MAE), to be compared with the DCNN method (section 3.2).


250 For the optimal models within each muscle/compartiment, the deviations of predicted local fibre
 251 orientations from ground truth were much too high to be useful in any real context, particularly in
 252 the Sol (compartment 1) muscle, which reports an optimal RMSE of over 26° (see fig. 3). Discrepancies
 253 were much lower in the GM and Sol (compartment 2) muscles at 11.3° and 7.1° , respectively. Visual
 254 comparisons were also made by comparing the wavelet analysis, presented as a heat-map, to the
 255 ground truth, also presented as a heat-map (see fig. 4). Visual comparison revealed quite sporadic
 256 agreement between wavelet predictions and ground truth. Visual comparison also revealed that any
 257 agreement between wavelet predictions and ground truth was not only heavily dependent on the
 258 presence of well-defined muscle fibre tracts, but also on the ability of the Frangi filter to extract those
 259 features.

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

265 3.2. DCNN Method

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

277



278 **Figure 5. Representative DCNN visual result.** (a) shows the ground truth, (b) shows the output from
 279 the DCNN, (c) shows line traces of (b). Compare (b) to (a), and notice that the heatmaps differ around
 280 the deep aponeurosis (boundary between Sol and GM). The line traces confirm this, by depicting
 281 visually accurate fibre traces throughout each muscle/compartiment until the deep aponeurosis,
 282 where curvature starts to increase rapidly.

283
284**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	$7.57^\circ \pm 8.37^\circ$	$4.46^\circ \pm 10.50^\circ$
	MAE ²	$8.23^\circ \pm 7.72^\circ$	$8.32^\circ \pm 7.81^\circ$
	RMSE ³	11.29°	11.41°
Sol (1st Compartment)	MD ¹	$-22.84^\circ \pm 14.01^\circ$	$-7.62^\circ \pm 12.88^\circ$
	MAE ²	$23.09^\circ \pm 13.61^\circ$	$12.14^\circ \pm 8.75^\circ$
	RMSE ³	26.80°	14.97°
Sol (2nd Compartment)	MD ¹	$1.61^\circ \pm 6.92^\circ$	$11.79^\circ \pm 10.86^\circ$
	MAE ²	$5.03^\circ \pm 5.01^\circ$	$13.33^\circ \pm 8.90^\circ$
	RMSE ³	7.11°	16.03°

285

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

286

4. Discussion

287

288

289

290

291

292

293

294

295

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).

296

297

298

299

300

301

302

303

304

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.

305

306

307

308

309

310

311

312

313

314

315

316

317

318

319

320

321

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 a DCNN 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 deep muscles than the wavelet method. Those predictions were also much more consistent as identified by standard deviations of the error metrics (see table. 1). Regardless of the similarity in performance, we suggest that the DCNN method could be improved with more sophisticated regularization (e.g. dropout and residual connections), or with additional labelled 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

322 and deals with all of the problems we identified as well as the original authors (image artefacts,
323 spurious fibre features, noise, reflections, etc.).

324 5. Conclusions

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

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

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

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

349 References

- 350 [1] P. J. Harding, I. D. Loram, N. Combes, and E. F. Hodson-Tole, "Ultrasound-Based Detection of
351 Fasciculations in Healthy and Diseased Muscles," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 3, pp. 512–518,
352 2016.
- 353 [2] I. D. Loram, B. Bate, P. Harding, R. Cunningham, and A. Loram, "Proactive Selective Inhibition Targeted
354 at the Neck Muscles: This Proximal Constraint Facilitates Learning and Regulates Global Control," *IEEE
355 Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 4, pp. 357–369, 2017.
- 356 [3] I. Loram, C. Maganaris, and M. Lakie, "Use of ultrasound to make noninvasive in vivo measurement of
357 continuous changes in human muscle contractile length," *J. Appl. Physiol.*, pp. 1311–1323, 2006.
- 358 [4] G. Q. Zhou, P. Chan, and Y. P. Zheng, "Automatic measurement of pennation angle and fascicle length
359 of gastrocnemius muscles using real-time ultrasound imaging," *Ultrasonics*, vol. 57, no. C, pp. 72–83,
360 2015.
- 361 [5] M. Rana, G. Hamarneh, and J. M. Wakeling, "Automated tracking of muscle fascicle orientation in B-
362 mode ultrasound images," *J. Biomech.*, vol. 42, no. 13, pp. 2068–2073, 2009.
- 363 [6] A. I. L. Namburete, M. Rana, and J. M. Wakeling, "Computational methods for quantifying in vivo
364 muscle fascicle curvature from ultrasound images," *J. Biomech.*, vol. 44, no. 14, pp. 2538–2543, 2011.
- 365 [7] X. Chen, Q. Li, S. Qi, H. Zhang, S. Chen, and T. Wang, "Continuous fascicle orientation measurement of
366 medial gastrocnemius muscle in ultrasonography using frequency domain Radon transform," *Biomed.
367 Signal Process. Control*, vol. 20, pp. 117–124, Jul. 2015.

- 368 [8] J. Darby, B. Li, N. Costen, I. Loram, and E. Hodson-Tole, "Estimating skeletal muscle fascicle curvature
369 from B-mode ultrasound image sequences," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 7, pp. 1935–1945, 2013.
- 370 [9] R. D. Herbert, S. C. Gandevia, and R. D. Herbert, "Changes in pennation with joint angle and muscle
371 torque: in vivo measurements in human brachialis muscle," *J. Physiol.*, vol. 484, no. Pt2, pp. 523–532,
372 1995.
- 373 [10] M. V Narici, T. Binzoni, E. Hiltbrand, J. Fasel, F. Terrier, and P. Cerretelli, "In vivo human gastrocnemius
374 architecture with changing joint angle at rest and during graded isometric contraction," *J. Physiol.*, vol.
375 496, no. 1, pp. 287–297, 1996.
- 376 [11] R. L. Lieber, J. Fridé N, and J. Fridén, "Functional and Clinical Significance of Skeletal Muscle
377 Architecture," *Muscle Nerve*, vol. 23, no. November, pp. 1647–1666, 2000.
- 378 [12] J. Darby, E. F. Hodson-Tole, N. Costen, and I. D. Loram, "Automated regional analysis of B-mode
379 ultrasound images of skeletal muscle movement," *J. Appl. Physiol.*, vol. 112, no. 2, pp. 313–327, 2012.
- 380 [13] F. Yeung, S. F. Levinson, and K. J. Parker, "Multilevel and Motion Model-Based Ultrasonic Speckle
381 Tracking Algorithms," *Ultrasound Med. Biol.*, vol. 24, no. 3, pp. 427–441, Mar. 1998.
- 382 [14] Y. Zhou, J.-Z. Li, G. Zhou, and Y.-P. Zheng, "Dynamic measurement of pennation angle of
383 gastrocnemius muscles during contractions based on ultrasound imaging," *Biomed. Eng. Online*, vol. 11,
384 no. 1, p. 63, 2012.
- 385 [15] G. Q. Zhou and Y. P. Zheng, "Automatic Fascicle Length Estimation on Muscle Ultrasound Images With
386 an Orientation-Sensitive Segmentation," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 12, pp. 2828–2836, 2015.
- 387 [16] F. Chen, "Hand gesture recognition using a real-time tracking method and hidden Markov models,"
388 *Image Vis. Comput.*, vol. 21, no. 8, pp. 745–758, 2003.
- 389 [17] Y. Zhou and Y. P. Zheng, "Estimation of Muscle Fiber Orientation in Ultrasound Images Using Revoting
390 Hough Transform (RVHT)," *Ultrasound Med. Biol.*, vol. 34, no. 9, pp. 1474–1481, Sep. 2008.
- 391 [18] H. Stark and N. Schilling, "A novel method of studying fascicle architecture in relaxed and contracted
392 muscles," *J. Biomech.*, vol. 43, no. 15, pp. 2897–2903, 2010.
- 393 [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional
394 Neural Networks," *Adv. Neural Inf. Process. Syst.* 25, pp. 1–9, 2012.
- 395 [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Arxiv.Org*, vol. 7,
396 no. 3, pp. 171–180, 2015.
- 397 [21] G. Hinton, "Dropout : A Simple Way to Prevent Neural Networks from Overfitting," vol. 15, pp. 1929–
398 1958, 2014.
- 399 [22] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," *Proc. 27th*
400 *Int. Conf. Mach. Learn.*, no. 3, pp. 807–814, 2010.
- 401 [23] L. D. Le Cun Jackel, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, B. Le Cun, J.
402 Denker, and D. Henderson, "Handwritten Digit Recognition with a Back-Propagation Network," *Adv.*
403 *Neural Inf. Process. Syst.*, pp. 396–404, 1990.
- 404 [24] T. Pfister, J. Charles, and A. Zisserman, "Flowing convnets for human pose estimation in videos," in
405 *Proceedings of the IEEE International Conference on Computer Vision*, 2015, vol. 2015 Inter, pp. 1913–1921.
- 406 [25] P. Fischer, A. Dosovitskiy, E. Ilg, ..., and T. Brox, "FlowNet: Learning Optical Flow with Convolutional
407 Networks," *Iccv*, p. 8, 2015.
- 408 [26] C. Payer, D. Štern, H. Bischof, and M. Urschler, "Regressing heatmaps for multiple landmark localization
409 using CNNs," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence*
410 *and Lecture Notes in Bioinformatics)*, 2016, vol. 9901 LNCS, pp. 230–238.

- 411 [27] H. Noh, S. Hong, and B. Han, "Learning deconvolution network for semantic segmentation," in
412 *Proceedings of the IEEE International Conference on Computer Vision*, 2016, vol. 11–18–Dece, pp. 1520–1528.
- 413 [28] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. a Viergever, "Multiscale vessel enhancement filtering,"
414 *Medial Image Comput. Comput. Invervention - MICCAI'98. Lect. Notes Comput. Sci. vol 1496*, vol. 1496, pp.
415 130–137, 1998.
- 416 [29] D.-J. Kroon, "Hessian based Frangi Vesselness filter," *MATLAB Central/File Exchange*, 2010. [Online].
417 Available: [https://uk.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-](https://uk.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-filter)
418 [filter](https://uk.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-filter). [Accessed: 01-Nov-2017].
- 419 [30] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks,"
420 *Proc. 13th Int. Conf. Artif. Intell. Stat.*, vol. 9, pp. 249–256, 2010.