

# Deep Residual Networks for Quantification of Muscle Fiber Orientation and Curvature from Ultrasound Images

Ryan Cunningham<sup>(✉)</sup>, Peter Harding, and Ian Loram

Manchester Metropolitan University, Manchester, England, UK  
ryan.cunningham@mmu.ac.uk

**Abstract.** This paper concerns fully automatic and objective measurement of human skeletal muscle fiber orientation directly from standard b-mode ultrasound images using deep residual (ResNet) and convolutional neural networks (CNN). Fiber orientation and length is related with active and passive states of force production within muscle. There is currently no non-invasive way to measure force directly from muscle. Measurement of forces and other contractile parameters like muscle length change, thickness, and tendon length is not only important for understanding healthy muscle, but such information has contributed to understanding, diagnosis, monitoring, targeting and treatment of diseases ranging from myositis to stroke and motor neurone disease (MND). We applied well established deep learning methods to ultrasound data recorded from 19 healthy participants (5 female, ages:  $30 \pm 7.7$ ) and achieved state of the art accuracy in predicting fiber orientation directly from ultrasound images of the calf muscles. First we used a previously developed segmentation technique to extract a region of interest within the gastrocnemius muscle. Then we asked an expert to annotate the main line of fiber orientation in  $4 \times 4$  partitions of 400 normalized images. A linear model was then applied to the annotations to regulate and recover the orientation field for each image. Then we applied a CNN and a ResNet to predict the fiber orientation in each image. With leave one participant out cross-validation and dropout as a regulariser, we were able to demonstrate state of the art performance, recovering the fiber orientation with an average error of just  $2^\circ$ .

**Keywords:** Ultrasound · Muscle analysis · Residual networks · Convolutional networks · Wavelet · Fiber orientation

## 1 Introduction

This paper presents a novel fully automatic computational approach to non-invasively measuring human skeletal muscle architectural parameters directly from 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 (Harding et al. 2016) as a non-invasive alternative to intramuscular electromyography (iEMG) for measuring twitch frequency, useful for early the diagnosis of motor neurone disease (MND).

Ultrasound has also recently demonstrated application to rehabilitative biofeedback (Loram et al. 2017). Other computational techniques have been developed for muscle-ultrasound analysis which would allow estimation of changes in muscle length during contraction (Loram et al. 2006), and changes in fiber orientation and length (Darby et al. 2013; Namburete et al. 2011; Rana et al. 2009). Muscle fiber orientation is one of the main identifying features of muscle state (Lieber and Fridén 2000). Without highly invasive methods, the measurement of force from muscle is currently not possible (Barry and Ahmed 1986; Finni et al. 1998; Finni et al. 2000; Gregor et al. 1987; Holden et al. 1994; Komi 1996; Komi et al. 1987; Lewis et al. 1982). Muscle fiber orientation, curvature and length is known to change with changes in force within the muscle (Herbert et al. 1995; Narici et al. 1996).

Broadly there are two main approaches to extracting fiber orientation/curvature automatically from ultrasound. The first approach is feature tracking, proposed by Darby and others (Darby et al. 2013). Feature tracking can be robust for small movements (i.e. where the difference in texture appearance of features between frames is similar) (Loram et al. 2006), however drift due to noise is inevitable, and further, severe drift can occur where for large movements, which cause disparity between texture patches (Loram et al. 2006; Yeung et al. 1998). Drift can be lessened using methods such as the Kanade Lucas-Tomasi (Shi and Tomasi 1994; Tomasi 1991) pyramidal method of tracking (Darby et al. 2012), but not solved since features can completely disappear from view as they move outside the ultrasound image plane causing feature loss (Darby et al. 2012). Darby et al. (2013) address the drift problem using a Bayesian particle filtering framework which regularizes the tracking based on a Gaussian process (GP) model of fiber shape. Their approach was validated on synthetic data with known curvatures, and evaluation with real data in the form of comparisons to existing techniques. Although this is an important contribution, the authors conclude that their method is outperformed in some scenarios by existing methods on synthetic data, and produces physiologically unrealistic fiber length values on real data (one test person/trial).

The second approach can be described as feature engineering and extraction, as proposed by Rana et al. (2009). Feature engineering can be a powerful approach to image analysis if a good model of the features can be constructed. In the case of muscle fibers within an ultrasound image, they appear as dark parallel lines (with slight curvature) at particular angles (typically 5–30°) within the muscle belly. These dark lines are surrounded by connective tissues which appear bright. The difficulty is that there are other structures which exhibit these features such as blood vessels, image artifacts and muscle boundaries. Rana and others model these high-contrast tube-like/vessel structures using orientated anisotropic wavelets. Due to noise and contrast variations within the image, after manual region of interest (ROI) identification they process the image with a multi-scale vessel enhancement filtering technique (Frangi et al. 1998) which enhances tubular structures in the image. After filtering they convolve orientated wavelets (0–90°) with the image, and for each pixel the maximum convolution gives the orientation of the fiber at that pixel. Their technique was evaluated on synthetic images with known orientations, and real images with operator identification of true fiber orientation. For synthetic images they report accuracy less than 0.02° (mean absolute error) for synthetic images, and they do not quantify the error on real images,

instead they report statistically significant differences between the wavelet method and the operator identification of fiber orientation (one test participant/trial during cycling).

A third approach which has thus far not been considered is to learn features which predict fiber orientation directly from data. In recent years, with the advent of graphics processing unit (GPU) computing and advancements in machine learning algorithms, such as the introduction of the rectified linear unit, dropout, and residual networks for very deep learning (Krizhevsky et al. 2012b; He et al. 2015; Hinton, 2014; Krizhevsky et al. 2012b; LeCun et al. 1990; Nair and Hinton 2010), image understanding and analysis is more feasible for challenging images such as medical ultrasound images. As static analysis has a clear and demonstrable advantage over dynamic analysis (tracking), we propose to learn to extract fiber orientation directly from data with manual identifications of fiber orientations. We hypothesize a machine learning approach will provide an improvement over feature engineering (Rana et al. 2009) because a convolutional neural network (CNN) can learn more complex convolution filters, which are more descriptive and discriminative of fibers and other non-fibrous structures in the image, over a filter-bank technique. One problem with a machine learning approach, as highlighted by Rana et al. (2009), is that manual identification of fiber orientation is subjective and difficult, and since supervised machine learning methods typically require large volumes of accurately labeled data, the CNN may not have the required information to create an accurate generalizable solution. We address this problem by regulating the operator identification process using multiple annotations and linear models. We compare a standard 7-layer deep (not counting 4 max-pooling layers) CNN with a 27-layer *very* deep (not counting 4 max-pooling layers) residual neural network (ResNet), both with and without dropout for CNN and ResNet. Results are compared against the wavelet approach on real images undergoing joint rotations and contractions.

## 2 Methods

### 2.1 Data Collection

Ultrasound data were recorded from 19 participants (5 female, ages:  $30 \pm 7.7$ ) during dynamic standing tasks. Participants stood upright on a programmable/controllable foot pedal system during 3 tasks while strapped to a backboard. During the tasks we recorded calf muscle (*medial gastrocnemius*: MG) activation using electromyography (EMG), ankle joint angle and joint torque, all at 1000 Hz, and ultrasound of the MG at 25 Hz (AlokaSSD-5000 PHD, 7.5 MHz). The three tasks were designed to observe distinct muscle mechanical changes during isolated joint angle changes via sinusoidal modulation of the pedal system (resulting in passive changes in the length of the calf muscles), isolated contraction via pushing down of the toes onto the pedals while the pedals were fixed in a neutral position (resulting in active changes in length of the calf muscles), and combined joint angle changes with contractions (resulting in a mixture of passive and active changes in calf muscle length). Simulink (Matlab, R2013a, The-MathWorks Inc., Natick, MA) was used to interface with the lab equipment (pedal system and EMG), and for video synchronization a hardware trigger was used to initiate recording at the start of each trial.

## 2.2 Automatic Identification of Region of Interest (Muscle Segmentation)

For segmentation and ROI identification we used an already established technique previously developed by our group (Cunningham et al. 2015) for segmenting skeletal muscle within ultrasound. The technique involved manual annotation of a number of training images (100) and construction of a principal components model of shape. The shape model is then used to define a mean ultrasound image by warping each texture in the space to the mean texture (after shape alignment via Procrustes analysis). Then a dictionary of textures and associated shapes is created from the mean texture and shape using the same method by warping the mean texture to equidistant points in the component model orthogonal to the axes of the  $n$  largest components (representing 90% of the shape variance). To segment a new image the dictionary is used to define an initial segmentation, and then a modified Active Shape Model is iteratively used to refine and complete the segmentation. The segmentation of the gastrocnemius muscle allowed extraction of a rectangular ROI ( $496 \times 128$ ) equidistant from its boundaries and orthogonal to the main axis of the muscle (linear least squares fit over the upper and lower boundary). For the purposes of this paper, generalization performance is not of primary concern since we only need accurate segmentation to evaluate the fiber orientation methods we have developed here. However, accuracy was evaluated on a leave one participant out basis and the results are presented in Sect. 3.

## 2.3 Regulated Expert Identification of Fiber Orientation


After segmentation and region extraction, an expert was asked to identify regional fiber orientation in 400 normalized ( $128 \times 128$ ) training images and 50 validation and 50 test images. Since there is regional variation in orientation due to curvature and other factors, each image was divided into a  $4 \times 4$  grid of  $32 \times 32$  sub-images, and then the expert identified the main line of fiber orientation in each sub-image. In many cases the expert could not identify fibers. In these cases the expert was allowed to mark the sub-image as undefined. Following manual identification of fiber orientation, linear models were fit to the annotation of each individual image,

$$\mathbf{w} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y},$$

where  $\mathbf{X}$  is a matrix in which the columns are a constant 1 (bias term) and the horizontal and vertical center coordinates of each sub-image (not including sub-images marked as undefined by the expert) relative to the full image,  $\mathbf{y}$  is a row vector containing the corresponding angle of the expert-defined line of fiber orientation for each pair of coordinates, and  $\mathbf{w}$  therefore is a vector of coefficients which can multiply an arbitrary pair of coordinates in the space, recovering a linear estimate of local fiber orientation,

$$\theta = \mathbf{X}\mathbf{w}.$$

The linear model is used not only to reconstruct the orientations which the expert could not identify, but also to regulate the expert's inconsistencies and variations in challenging images since physiologically we expect fibers to follow similar trajectories (Fig. 1).



**Fig. 1.** Regulated expert identification of regional fiber orientation. Left: expert identification of main fiber orientation. Middle: reconstruction after linear model (blue dashed) of original identification (green solid). Right: Line trace of linear fit over entire image, revealing regional curvature over the muscle. (Color figure online)


## 2.4 Vessel Enhancement and Wavelet Method

After segmentation and ROI identification and before wavelet analysis all images were processed at full resolution ( $496 \times 128$ ) with the vessel enhancement filter in the same way as (Rana et al. 2009) to enhance the appearance of fibers (and other tubular structures) in the image. We used an existing implementation of (Frangi et al. 1998). Parameters were optimized empirically on the training set to obtain processed images similar to those presented by Rana et al. (2009).

Following vessel enhancement all processed images were convolved at full resolution ( $496 \times 128$ ) with a set of orientated ( $41 \times 41$ ) Gabor wavelets,

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

where  $k$  is the kernel size,  $d$  is a damping term,  $\alpha$  is the orientation, and  $f$  is the spatial frequency. All parameters were set as originally described in (Rana et al. 2009). Wavelets were constructed in the range of  $0^\circ$ – $90^\circ$  in  $1^\circ$  increments. Following convolution with the filter bank, the maximum wavelet convolution at each pixel was stored. From the list of maximum convolutions, we then removed any convolution values less than one standard deviation from the mean. At that point we applied a linear model to the location of the pixels and wavelet orientations given by the maximum convolutions. This was done as in Sect. 2.3 to resample the orientations at points comparable with the expert identifications and to regulate noise/error (Figs. 2 and 3).



**Fig. 2.** Vessel filter. This figure shows how a noisy ultrasound ROI image (left) can be enhanced (right) by the vesselness filter prior to being convolved with the wavelet filter bank.



**Fig. 3.** Wavelet filter bank. The figure shows 10 selected orientated wavelets ranging from  $17^\circ$  to  $80^\circ$  in increments of  $7^\circ$ .

## 2.5 Deep Learning Method


After segmentation and ROI identification, due to concern over model training time all images were resampled (bilinear interpolation) to a standard size of  $(128 \times 128)$ . We then trained a CNN and a ResNet, first without dropout, and then with dropout in the fully connected layer. The architecture we chose was a simple one due to concern over training time, yet we empirically confirmed that it was complex enough to learn the training set. The input to each net was the resampled images, and the output was 16 linear regression units, one for each fiber orientation of the subsampled images derived from the linear model of the expert identifications.

Each net had rectified linear units (ReLU) in the hidden layers, and 4 max pooling layers. In between pooling layers the number of convolutional filters in the CNN from input to output was  $25-p-36-p-49-p-100-p-196$  (where  $p$  denotes max pooling layer) and one fully connected layer of 1024 Leaky ReLUs ( $\alpha = 1e - 1$ ). The ResNet had the same architecture with the main difference being 5 times the number of convolutional layers in between max pooling layers, and there were identity shortcut connections used in alternating consecutive layers (2 shortcut connections per set of 5 convolutional layers). We did not use projection shortcuts since identity shortcuts were adequate for our purposes. Both CNN and ResNet used a convolution stride of  $1 \times 1$ , and the same max pooling dimensions (from input to output:  $2 \times 2 - 4 \times 4 - 2 \times 2 - 4 \times 4$ ). For the CNN the filter sizes were all  $3 \times 3$  with the exception of the input layer which was  $7 \times 7$ . For the ResNet the filter sizes were  $3 \times 3$  throughout. Initial weights were drawn from a normal distribution, normalized by the product of the size of fan-in and the dropout coefficient,  $w_i = N(0, 1) \sqrt{\frac{2}{n_{1-p}}}$ .

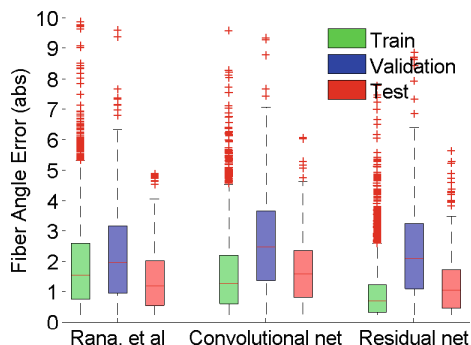
For training, both nets used a learning rate of  $5e - 4$  with a momentum of  $9e - 1$  and a small amount of L2 weight decay ( $1e - 3$ ), and no learning rate decay or adaptive learning rate was used. Models were trained with ( $p = 5e - 1$ ) and without dropout ( $p = 0$ ) for 250 full passes over the training set with a batch size of 1 (i.e. online stochastic gradient descent). For testing, the mean square error (MSE) was monitored on the validation and testing sets and the model with the lowest validation error was selected as the solution, at which point the testing error was used to estimate generalized performance on unseen data. We also present root MSE for performance evaluation in degree units. Finally, we quantify the linear predictive power ( $R^2$ ) of each approach for predicting the measured ankle joint torque during all three trials as described in Sect. 2.1.

### 3 Results

Images are presented along with predictions of the regional fiber orientation in the testing participant for the three approaches under investigation (see Fig. 4). All three methods clearly produce realistic estimates of the fiber direction; however the wavelet method demonstrates some fairly substantial deviations from the ground truth (expert annotations). RMSE measures report accuracy in general for all methods comfortable below  $2^\circ$ , with the ResNet approaching errors of  $1^\circ$  in the training and test cases (see Fig. 5). The most important result came from the  $R^2$  linear prediction of measured joint torque during the combined, active, and passive trials in the test participant. Both CNN and ResNet clearly demonstrate greater predictive power over the wavelet method, despite using a lower resolution image ( $128 \times 128$  vs  $496 \times 128$ ). Further, the ResNet method showed much less deviation from ground truth as confirmed by MSE and RMSE standard deviation (see Table 1) and visual assessment (see Fig. 4).



**Fig. 4.** Method comparison on held out test participant (left: wavelet method, middle: CNN, right: ResNet). The blue dashed line is the prediction of the fiber orientation from each of the methods, whereas the green solid line is the regulated expert label. Notice the clear improvement of the ResNet over the CNN and the wavelet-based method. (Color figure online)



**Fig. 5.** Absolute fiber angle prediction error. This figure summarizes the errors of the three approaches in real values (angle in degrees).

**Table 1.** Summary of results. This table summarises the results of all three methods for each of the labeled data sets. The predictive power is also given ( $R^2$ ) for the testing participant. In general the convolutional network and wavelet methods perform similarly. Conversely, the residual network reports a clear improvement with generally less deviation and lower error over all cases. The discriminating factor is in the predictive power, where both deep learning methods give a marked improvement over Rana et al.

Method	Data set	RMSE	MSE	$R^2$
Rana et al. (2009)	Training	$2.30^\circ \pm 3.46^\circ$	$0.84 \pm 1.22$	0.3837
	Validation	$2.76^\circ \pm 3.44^\circ$	$0.77 \pm 1.07$	
	Testing	$1.72^\circ \pm 1.98^\circ$	$0.56 \pm 0.99$	
Convolutional net	Training	$2.00^\circ \pm 3.50^\circ$	$0.60 \pm 0.95$	0.4222
	Validation	$3.08^\circ \pm 3.44^\circ$	$0.91 \pm 1.08$	
	Testing	$1.94^\circ \pm 2.08^\circ$	$0.70 \pm 1.35$	
Residual net*	Training	$1.25^\circ \pm 2.50^\circ *$	$0.24 \pm 0.49 *$	0.5897*
	Validation	$2.70^\circ \pm 3.03^\circ *$	$0.72 \pm 0.93 *$	
	Testing	$1.53^\circ \pm 1.88^\circ *$	$0.47 \pm 1.08 *$	

## 4 Discussion

The combination of segmentation (Cunningham et al. 2015) with deep learning provides a general framework for extracting physiological information directly from ultrasound images of human skeletal muscle. Comparison with an existing state of the art wavelet convolution method reveals a new benchmark set by deep residual networks (see Table 1). The convolutional network was comparable in accuracy to the wavelet method and in some cases proved more robust, yet it did not demonstrate the gain in performance that the residual network did. The benefits of deep residual networks are clear for classification problems (He et al. 2015), however the work presented here contributes evidence of success in the domain of multiple regression. All of the methods presented here provide the necessary information to quantify fiber orientation which consequently provides the information required to estimate fiber length and pennation angles, which are common features of interest within biomechanics and medical analysis.

Although labels were easy to acquire (c. 8 h of annotation), the sorts of volumes of data typically used to train deep networks was not feasible for this study. Although the results presented here are competitive with the state of the art – *if not a significant improvement* – we would argue strongly that those results could certainly be improved perhaps by an order of magnitude if the size of the labeled data set could be increased by hundreds or thousands. Due to concern over the said small data set, state of the art regularization techniques were used to improve model generalization: dropout, early stopping, weight sharing (convolutions), max-pooling, and identity connections (residual learning). The success of dropout as a regulariser is well established in classification problems (Krizhevsky et al. 2012a), however we our results evidence success in the domain of multiple regression. We show that dropout behaves as expected and continues to minimize error while the nets without dropout diverge before reaching the same level of performance as the nets with dropout.



The previous state of the art used engineered wavelets based on prior knowledge of fiber size and expected orientation, and these wavelets were convolved with the image (after vesselness filtering). This approach (feature engineering) is becoming an out-dated concept where there exists an abundance of data, since deep learning methods can learn exactly what sort of filters would extract only the relevant information and discard all of the irrelevant information by learning directly from data.

This preliminary work did not investigate many of the available free parameters available to deep learning methods. Preventing over-fitting was our primary concern after learning the training data, yet for this preliminary investigation we only used early stopping with and without dropout. Dropout demonstrably works well even without parameterization, although we must highlight that there is potential to use dropout in multiple layers and with different frequencies, all of which would need fully investigating on a validation set. We used some weight decay as is standard practice to help convergence; however we acknowledge that this parameter could also be tuned to improve generalization. We did investigate different architectures in terms of max pooling configuration and number of filters per layer, but we did not attempt to cross-validate all of the different models due to concern over training time – we chose a reasonably sized model with standard pooling configurations and then relied on dropout and early stopping. We note that investigations of different architectures would become more important when the project scales up in terms of increasing the size of the training set.

## 5 Conclusions

This paper has presented a robust and repeatable method for extracting fiber orientations directly from ultrasound images of human skeletal muscle. A clear improvement has been demonstrated over a widely used and well-established wavelet-based method (Rana et al. 2009). In this paper deep learning principals were applied to automatically learn discriminative features of muscle fibrous tissue from a relatively small data set of 400 labeled images. In order to provide robust and accurate labels a novel annotation method was introduced in which linear models are used to regulate annotations and recover a fiber orientation field. That development led to a successful extension of the wavelet-based method which facilitated comparison with the deep learning methods we have introduced here. The wavelet-based method previously was unable to recover robust estimations of regional fiber orientation, and the application of linear models allowed recovery of that information as evidenced (see Fig. 5 and Table 1). This paper provides further evidence that very deep networks (ResNets) can help performance and generalisation, even when there is not an abundance of labeled data available. With additional data we propose that this project could easily be extended more 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.

## References

- Barry, D., Ahmed, A.M.: Design and performance of a modified buckle transducer for the measurement of ligament tension. *J. Biomech. Eng.* **108**(2), 149–152 (1986). <https://doi.org/10.1115/1.3138594>
- Cunningham, R., Harding, P., Loram, I.: Real-time ultrasound segmentation, analysis and visualization of deep cervical muscle structure. *Trans. Med. Imaging* **36**(2), 653–665 (2015). <https://doi.org/10.1109/TMI.2016.2623819>
- Darby, J., Hodson-Tole, E.F., Costen, N., Loram, I.D.: Automated regional analysis of B-mode ultrasound images of skeletal muscle movement. *J. Appl. Physiol.* **112**(2), 313–327 (2012). <https://doi.org/10.1152/jappphysiol.00701.2011>
- Darby, J., Li, B., Costen, N., Loram, I., Hodson-Tole, E.: Estimating skeletal muscle fascicle curvature from B-mode ultrasound image sequences. *IEEE Trans. Biomed. Eng.* **60**(7), 1935–1945 (2013). <https://doi.org/10.1109/TBME.2013.2245328>
- Finni, T., Komi, P.V., Lukkariniemi, J.: Achilles tendon loading during walking: application of a novel optic fiber technique. *Eur. J. Appl. Physiol. Occup. Physiol.* **77**(3), 289–291 (1998). <https://doi.org/10.1007/s004210050335>
- Finni, T., Komi, P.V., Lepola, V.: In vivo human triceps surae and quadriceps femoris muscle function in a squat jump and counter. *Eur. J. Appl. Physiol.* **83**, 416–426 (2000). <https://doi.org/10.1007/s004210000289>
- Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A., Delp, S. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998). doi:10.1007/BFb0056195. <https://doi.org/10.1016/j.media.2004.08.001>
- Gregor, R.J., Komi, P.V., Järvinen, M.: Achilles tendon forces during cycling. *Int. J. Sports Med.* **8**(Suppl. 1), 9–14 (1987). <https://doi.org/10.1055/s-2008-1025698>
- Harding, P.J., Loram, I.D., Combes, N., Hodson-Tole, E.F.: Ultrasound-based detection of fasciculations in healthy and diseased muscles. *IEEE Trans. Biomed. Eng.* **63**(3), 512–518 (2016). <https://doi.org/10.1109/TBME.2015.2465168>
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *Arxiv.Org* **7**(3), 171–180 (2015). <https://doi.org/10.3389/fpsyg.2013.00124>
- Herbert, R.D., Gandevia, S.C., Herbert, R.D.: Changes in pennation with joint angle and muscle torque: in vivo measurements in human brachialis muscle. *J. Physiol.* **484**(Pt2), 523–532 (1995). <https://doi.org/10.1113/jphysiol.1995.sp020683>
- Hinton, G.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014)
- Holden, J.P., Grood, E.S., Korvick, D.L., Cummings, J.F., Butler, D.L., Bylski-Austrow, D.I.: In vivo forces in the anterior cruciate ligament: direct measurements during walking and trotting in a quadruped. *J. Biomech.* **27**(5), 517–526 (1994). [https://doi.org/10.1016/0021-9290\(94\)90063-9](https://doi.org/10.1016/0021-9290(94)90063-9)
- Komi, P.V.: Optic fibre as a transducer of tendomuscular forces. *Eur. J. Appl. Physiol. Occup. Physiol.* **72**(3), 278–280 (1996). <https://doi.org/10.1007/BF00838652>
- Komi, P.V., Salonen, M., Jarvinen, M., Kokko, O.: In vivo registration of Achilles tendon forces in man. I. Methodological development. *Int. J. Sports Med.* **8**(Suppl. 1), 3–8 (1987). <https://doi.org/10.1055/s-2008-1025697>
- Krizhevsky, A., Sutskever, I., Hinton, G. E.: ImageNet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1–9 (2012a). <https://doi.org/10.1117/12.2176558>

- Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (2012b)
- LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D.: Handwritten digit recognition with a back-propagation network. In: *Advances in Neural Information Processing Systems*, pp. 396–404 (1990). <https://doi.org/10.1111/dsu.12130>
- Lewis, J.L., Lew, W.D., Schmidt, J.: A note on the application and evaluation of the buckle transducer for the knee ligament force measurement. *J. Biomech. Eng.* **104**(2), 125–128 (1982). <http://www.ncbi.nlm.nih.gov/pubmed/7078126>
- Lieber, R.L., Fridén, J.: Functional and clinical significance of skeletal muscle architecture. *Muscle Nerve* **23**(11), 1647–1666 (2000). [https://doi.org/10.1002/1097-4598\(200011\)23:11<1647::AID-MUS1>3.0.CO;2-M](https://doi.org/10.1002/1097-4598(200011)23:11<1647::AID-MUS1>3.0.CO;2-M)
- Loram, I., Bate, B., Harding, P., Cunningham, R., Loram, A.: Proactive selective inhibition targeted at the neck muscles: this proximal constraint facilitates learning and regulates global control. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**(4), 357–369 (2017). <https://doi.org/10.1109/TNSRE.2016.2641024>
- Loram, I., Maganaris, C., Lakie, M.: Use of ultrasound to make noninvasive in vivo measurement of continuous changes in human muscle contractile length. *J. Appl. Physiol.* **100**, 1311–1323 (2006). <https://doi.org/10.1152/japplphysiol.01229.2005>
- Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. *Proc. Int. Conf. Mach. Learn.* **3**, 807–814 (2010). <https://doi.org/10.11.1.165.6419>
- Namburete, A.I.L., Rana, M., Wakeling, J.M.: Computational methods for quantifying in vivo muscle fascicle curvature from ultrasound images. *J. Biomech.* **44**(14), 2538–2543 (2011). <https://doi.org/10.1016/j.jbiomech.2011.07.017>
- Narici, M.V., Binzoni, T., Hiltbrand, E., Fasel, J., Terrier, F., Cerretelli, P.: In vivo human gastrocnemius architecture with changing joint angle at rest and during graded isometric contraction. *J. Physiol.* **496**(1), 287–297 (1996). <https://doi.org/10.1113/jphysiol.1996.sp021685>
- Rana, M., Hamarneh, G., Wakeling, J.M.: Automated tracking of muscle fascicle orientation in B-mode ultrasound images. *J. Biomech.* **42**(13), 2068–2073 (2009). <https://doi.org/10.1016/j.jbiomech.2009.06.003>
- Shi, J., Tomasi, C.: Good features to track. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 1994*, pp. 593–600 (1994). <https://doi.org/10.1109/cvpr.1994.323794>
- Tomasi, C.: Detection and tracking of point features. *School of Computer Science, Carnegie Mellon University*, **91**(April), 1–22 (1991). [https://doi.org/10.1016/S0031-3203\(03\)00234-6](https://doi.org/10.1016/S0031-3203(03)00234-6)
- Yeung, F., Levinson, S.F., Parker, K.J.: Multilevel and motion model-based ultrasonic speckle tracking algorithms. *Ultrasound Med. Biol.* **24**(3), 427–441 (1998). [https://doi.org/10.1016/S0301-5629\(97\)00281-0](https://doi.org/10.1016/S0301-5629(97)00281-0)