

This is the author-created version of the following work:

**Saleh, Alzayat, Laradji, Issam H., Lammie, Corey, Vazquez, David, Flavell, Carol Ann, and Rahimi Azghadi, Mostafa (2021) *A deep learning localization method for measuring abdominal muscle dimensions in ultrasound images*. IEEE Journal of Biomedical and Health Informatics, 25 pp. 3865-3873.**

Access to this file is available from:

<https://researchonline.jcu.edu.au/68402/>

(c) 2021 IEEE.

Please refer to the original source for the final version of this work:

<https://doi.org/10.1109/JBHI.2021.3085019>

# A Deep Learning Localization Method for Measuring Abdominal Muscle Dimensions in Ultrasound Images

Alzayat Saleh<sup>1</sup>, Issam H. Laradji<sup>2</sup>, Corey Lammie<sup>3</sup>, David Vazquez<sup>4</sup>, Carol A Flavell<sup>5</sup>, and Mostafa Rahimi Azghadi<sup>6</sup>

**Abstract**—Health professionals extensively use Two-Dimensional (2D) Ultrasound (US) videos and images to visualize and measure internal organs for various purposes including evaluation of muscle architectural changes. US images can be used to measure abdominal muscles dimensions for the diagnosis and creation of customized treatment plans for patients with Low Back Pain (LBP), however, they are difficult to interpret. Due to high variability, skilled professionals with specialized training are required to take measurements to avoid low intra-observer reliability. This variability stems from the challenging nature of accurately finding the correct spatial location of measurement endpoints in abdominal US images. In this paper, we use a Deep Learning (DL) approach to automate the measurement of the abdominal muscle thickness in 2D US images. By treating the problem as a localization task, we develop a modified Fully Convolutional Network (FCN) architecture to generate blobs of coordinate locations of measurement endpoints, similar to what a human operator does. We demonstrate that using the TrA400 US image dataset, our network achieves a Mean Absolute Error (MAE) of 0.3125 on the test set, which almost matches the performance of skilled ultrasound technicians. Our approach can facilitate next steps for automating the process of measurements in 2D US images, while reducing inter-observer as well as intra-observer variability for more effective clinical outcomes.

**Index Terms**—Ultrasound, Transversus Abdominis, Musculoskeletal, Deep Learning, Convolutional Neural Network

Manuscript received January 26, 2021.

Study approval was obtained from the Human Research Ethics Committee of the James Cook University (JCUH7878).

The author(s) received no financial support for the research, authorship, and/or publication of this article.

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

(Corresponding author: Alzayat Saleh and Mostafa Rahimi Azghadi.)

A. Saleh is with College of Science and Engineering, James Cook University, Townsville, Australia (e-mail: [alzayat.saleh@my.jcu.edu.au](mailto:alzayat.saleh@my.jcu.edu.au)).

I. H. Laradji is with McGill University, Montreal, Canada. and Element AI, Montreal, Canada (e-mail: [issam.laradji@gmail.com](mailto:issam.laradji@gmail.com)).

C. Lammie is with College of Science and Engineering, James Cook University, Townsville, Australia (e-mail: [corey.lammie@jcu.edu.au](mailto:corey.lammie@jcu.edu.au)).

D. Vazquez is with Element AI, Montreal, Canada (e-mail: [david.vazquez@servicenow.com](mailto:david.vazquez@servicenow.com)).

C. A. Flavell is with College of Healthcare Sciences, James Cook University, Townsville, Australia (e-mail: [carol.flavell@jcu.edu.au](mailto:carol.flavell@jcu.edu.au)).

M. Rahimi Azghadi is with College of Science and Engineering, James Cook University, Townsville, Australia (e-mail: [mostafa.rahimiazghadi@jcu.edu.au](mailto:mostafa.rahimiazghadi@jcu.edu.au)).

## I. INTRODUCTION

HEALTH professionals extensively use 2D Ultrasound (US) videos and images to visualize and measure internal organs for various purposes including evaluation of muscle architectural and functional changes for diagnostic purposes, and as an outcome measure to evaluate rehabilitation effects. Compared to Computerized Tomography (CT) scans, US imaging is low-cost [1], more sensitive [2], and does not expose patients to ionizing radiation [3].

Two-dimensional US imaging is a valid method of measuring abdominal muscle dimensions in patients undergoing physical rehabilitation [4]. Previous literature [5] has demonstrated that the measurement of the Transversus Abdominis (TrA), a broad paired muscular sheet found on the lateral sides of the abdominal wall, can be used to quantitatively express dysfunctional TrA activation. In terms of calculating TrA activation (TrA-C), US measurements of TrA thickness at rest (RTrA), and when fully contracted (CTrA) have been used as a preeminent measurement [5], [6]. Some examples of these measurements are shown in Fig. 1.

To evaluate the intra-rater and inter-rater reliability of measurements of TrA muscle thickness conducted by novice examiners, the authors of [7] conducted a study on military personnel examined by two ultrasound technicians and reported intra-rater Intraclass Correlation Coefficient (ICC) values from 0.90 to 0.98. The inter-rater ICC values ranged from 0.39 to 0.79 and concluded that the inter-rater reliability of two novice examiners was poor to fair. This low reliability between examiners challenges the efficacy of this method of muscle measurement in diagnostics and rehabilitation.

Reliability of US-based TrA measurement is deemed to be dependent on standardized methodological factors such as imaging technique, or issues associated with landmark detection on the 2D raw image [4]. An automated measurement tool that decreases variability at the image measurement stage may help health professionals make more accurate diagnoses and evaluations of rehabilitation progress. In addition, image measurement may be expedited by the tool, providing benefit in terms of saving clinical time and cost and improving patient outcomes. Deep Learning (DL) offers a great potential in developing such tool, and is therefore investigated in this paper.




Fig. 1: Images depicting TrA measurements, in contracted (left) and resting (right) states.

Previous DL works on US [8], [9] have framed the measurement task as a landmark detection problem, where the goal is to identify the key points on the image. In the general landmark detection approach (e.g. facial landmark detection), there is a defined local appearance of these landmarks. However, the measurements of TrA muscle thickness from US images is defined from local appearance and global structural information of the muscle. Consequently, we frame our task as localization of the caliper endpoints, similar to what a US technician does, to increase user-interpretability and use of global structural information.

For target localization in images and for many other computer vision tasks, Convolutional Neural Network (CNN) [10]–[12] has become the dominant DL solution. The standard CNN is not fully convolutional because it often contains dense layers [13], which have many parameters (i.e. computationally expensive). In contrast, FCN [14] is an architecture employing only convolution layers. Avoiding the use of dense layers means fewer parameters (i.e. faster to train the network). FCN consists of a downsampling path, used to extract and interpret the context, and an upsampling path, which allows for localization of each pixel. FCN also uses skip connections [15] to retrieve the fine-grained spatial information missed in the downsampling path.

In this paper, an object localization framework resembling the ultrasound technicians procedure is proposed to solve the problem of measurement of abdominal muscle thickness in 2D US images. Our specific contributions are as follows:

- 1) We propose a novel localization framework that uses FCN8 [14] with a CoordConv layer [16] to detect the points for muscle dimension measurement while leveraging Localization-based Counting Fully Convolutional Network (LCFCN) loss to maximize the performance of the network with point-level supervision;
- 2) We demonstrate that our novel approach of using LCFCN along with CoordConv yields significant improvement over conventional landmark detection and almost matches the performance of ultrasound technicians.

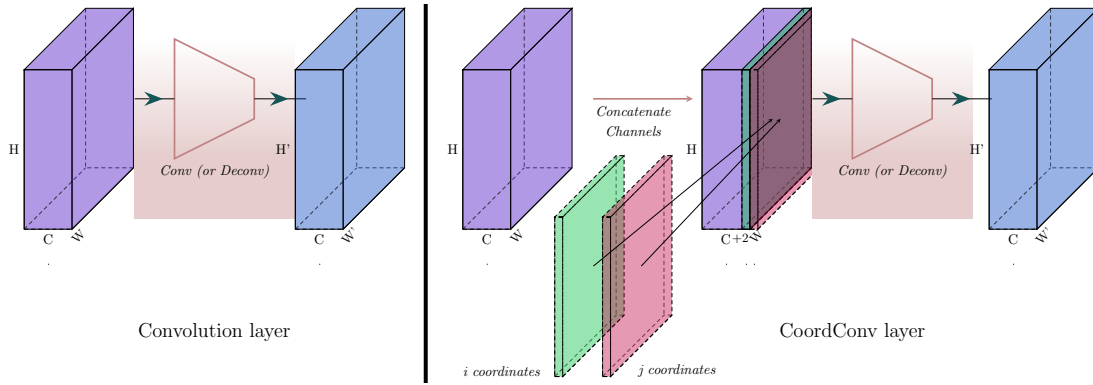
## II. RELATED WORK

DL has been extensively used in various medical applications including US image processing. For instance, it has been used for automated identification and segmentation of tumours in breast ultrasound images [17]–[19]. In addition, it has been employed for prenatal screening, which is a classic therapeutic application of US [20], [21]. The diagnosis of possible fetal problems or chromosomal abnormalities is done through monitoring fetal growth and development. This diagnosis can be complicated, requiring a highly trained doctor to promptly recognise the scan planes and structures of interest. Baumgartner *et al.* [22] used DL to significantly simplify this process by automated fetal standard scan plane for real-time localization and detection in these planes in freehand ultrasound. Cronin *et al.* [23] trained a deep neural networks (based on U-net) to detect muscle fascicles and aponeuroses using a set of labelled musculoskeletal ultrasound images. For ultrasound muscle segmentation, Cunningham *et al.* [24] have developed a method for segmenting five bilateral cervical muscles and the spine via ultrasound images. DL methods have also proven to outperform traditional methods in many other ultrasound imaging applications [25], [26].

Conventional landmark detection has so far been the dominant method in US localization approaches [8], [9], [27]–[29]. For instance, fusing local appearance of each landmark in an ultrasound image with the spatial configuration of the rest of the landmarks has achieved good landmark localization in a previous work [27]. The standard landmark detection approach has been used in [8] to detect measurement keypoint locations for left ventricular internal dimension measurements by computing their regression estimates using a centre of mass layer with a CNN. Gilbert *et al.* [9] extended the work of Sofka *et al.* [8] by targeting the intra-ventricular septum and the left ventricular posterior wall measurements. They introduced anatomically meaningful ground truth heatmaps, which follow the expected spatial distribution of keypoints.

The localization approach has been also used by Sloun *et al.* [28] to develop a CNN that automatically identifies clinically relevant B-line artifacts in lung from an ultrasound





**Fig. 3:** A visual representation of how CoordConv differs from the standard convolution layer. The left is a schematic diagram of standard convolutional layer maps from a representation block with shape  $h \times w \times c$  to a new representation of shape  $h' \times w' \times c'$ . The right schematic diagram of CoordConv layer shows that it differs from the standard convolutional layer by first concatenating extra hard-coded  $i$  and  $j$  coordinates channels to the incoming representation.

the activation output that determines the likelihood of a pixel belonging to the localization or segmentation target. In our localization task, we obtain per-pixel probabilities by applying the Softmax activation function to compute  $\hat{Y}$ , which contains the likelihood that a pixel either belongs to the background or muscle edge. The LCFCN loss  $\mathcal{L}_L$  is described using equation (1)

$$\mathcal{L}_L = \underbrace{\mathcal{L}_I(\hat{Y}, Y)}_{\text{Image-level loss}} + \underbrace{\mathcal{L}_P(\hat{Y}, Y)}_{\text{Point-level loss}} + \underbrace{\mathcal{L}_S(\hat{Y}, Y)}_{\text{Split-level loss}} + \underbrace{\mathcal{L}_F(\hat{Y}, Y)}_{\text{False positive loss}}, \quad (1)$$

where  $Y$  represents the point annotation ground-truth.

The **image-level loss** ( $\mathcal{L}_I$ ) trains the model to predict whether there is an object in the image, and it is described using equation (2).

$$\mathcal{L}_I(\hat{Y}, Y) = -\frac{1}{|C_e|} \sum_{c \in C_e} \log(\hat{Y}_{t_c c}) - \frac{1}{|C_{-e}|} \sum_{c \in C_{-e}} \log(1 - \hat{Y}_{t_c c}) \quad (2)$$

where  $t_c = \text{argmax}_{i \in \mathcal{I}} \hat{Y}_{i c}$ . Here, each class  $c \in C_e$  which shows the set of classes present in the image, while  $C_{-e}$  is the set of classes not present in the image.

The **point-level loss** ( $\mathcal{L}_P$ ) encourages the model to predict a pixel for each object instance, and it is described using equation (3).

$$\mathcal{L}_P(\hat{Y}, Y) = -\sum_{i \in \mathcal{I}_s} \log(\hat{Y}_{i Y_i}), \quad (3)$$

where  $Y_i$  represents the true label of pixel  $i$  and  $\mathcal{I}_s$  represents the locations of the object instances.

The **split-level loss** ( $\mathcal{L}_S$ ) enforces the model to predict a single blob per instance [43], and it is described using equation (4).

$$\mathcal{L}_S(\hat{Y}, Y) = -\sum_{i \in T_b} \alpha_i \log(\hat{Y}_{i 0}), \quad (4)$$

where  $\hat{Y}_{i 0}$  is the probability that pixel  $i$  belongs to the background class,  $\alpha_i$  is the number of point-annotations in

the blob in which pixel  $i$  lies, and  $T_b$  is the set of pixels representing the boundaries determined by the local and global segmentation.

The **false-positive loss** ( $\mathcal{L}_F$ ) discourages the model from predicting a blob with no point annotations, in order to reduce the number of false positive predictions, and is described using equation (5).


$$\mathcal{L}_F(\hat{Y}, Y) = -\sum_{i \in B_{fp}} \log(\hat{Y}_{i 0}), \quad (5)$$

where  $B_{fp}$  is the set of pixels constituting the blobs predicted for each class (except the background class) that contain no ground-truth point annotations (note that  $S_{i 0}$  is the probability that pixel  $i$  belongs to the background class).


Applying LCFCN loss on the original activation map generates a small blob around the centre of the object, in our case the caliper endpoints. To measure the TrA thickness, we take the centroid of each blob to get coordinate locations and measure the distances between the coordinate pairs that our model predicted.

#### D. TrA400 Dataset

Ethics approval for the use of the ultrasound image set in this study was granted by the Human Research Ethics Committee of James Cook University (JCU) (JCUH7878). The TrA400 dataset consists of 400 2D US images used to measure the thickness and activation changes in the TrA. The US images were captured from abdominal US videos using methods described previously in Kennedy *et al.* [5]. The 400 2D US images are from 56 participants and at least 4 images per participant were captured as detailed below. For the first 38 participants, each one was examined by two examiners each taking 2 types of measurements, “freehand” and “probe force device” for two different states, i.e. CTrA and RTrA. Therefore,  $38 \times 2 \times 2 \times 2 = 304$  images were captured. For some of these 38 participants, 20 additional random measurements were also performed, bringing the total number of measurements to 324. For the remaining 18 participants, 4 measurements per participant were captured in CTrA and



**Fig. 4:** Point annotations in a sample US image ( $X$ ) (left). The points in the training mask image ( $Y$ ) (middle) are extracted from a human-annotated image (right) and are inflated and highlighted for visibility, but only the center pixel and its class label are collected and used.



**Fig. 5:** A sample full resolution US image (left) and its corresponding (right) cropped image ( $895 \times 745$  pixels).

RTrA and using “freehand” and “probe force device”. For 10 of these 18 participants 2 extra measurements were also randomly conducted. These resulted in 82 images from the 18-subject cohort, bringing the total number of images in our collected dataset to 406 images. Two of these images were unreadable by humans so were removed. In addition, an additional 4 images were removed randomly for the built dataset to reach 400 still images to carry out cross-validation across 10 folds.

US videos were imported into VideoPad® (NCH Software, 2016). The examiner extracted one still image from each imported video. Still resting thickness of the TrA (RTrA) images from all videos were captured at the end-expiratory phase and contracted thickness of the TrA (CTrA) images were taken during Abdominal Drawing-In Maneuver (ADIM) [5]. The end-expiratory phase and contracted images were extracted from different videos. The examiner selected still images based on the highest quality visualization of TrA available. Several videos were captured per participant, during the US session. Participants remained on the treatment plinth for the duration of video capture. This gave the examiner a good opportunity to select a proper still image from each of the videos.


Full details regarding the study design and inclusion and exclusion criteria of participants are available in previous studies [5], [6].

## IV. EXPERIMENTS

In this section we describe the experiments performed, and provide details of the methods and evaluation metrics used in our experiments.

### A. Dataset Preparation

The extracted US images were annotated by an examiner to measure the muscle thickness. The caliper endpoints were defined for abdominal muscle dimension measurements as the start and the endpoints of the line in the middle of the TrA, which is drawn by a human examiner between the superior aspect of the inferior hyperechoic fascial border of the TrA and the inferior border of its superior hyperechoic fascial line, to measure the muscle thickness, as depicted in Fig. 1. In order to train our deep learning network, point-level annotated binary masks were obtained by converting the start and the end of the line in the middle of the TrA to  $(x_1, y_1)$  and  $(x_2, y_2)$  coordinates, respectively. Here,  $(x_i, y_i)$  coordinates have been given a value one as foreground, while the rest of the image pixels are turned to zeros as the background. The whole TrA400 dataset consists of precise image-mask pairs ( $X$ ,  $Y$ ) (see Fig. 4). The extracted US images original size are between



**Fig. 6:** Cross-validation distribution for TrA400 dataset. Each rectangle represents a group of 40 images. Green rectangles are used to denote test splits, yellow rectangles are used to show validation splits, and blue rectangles are used to demonstrate train splits. Performance is reported across all folds.

1010 × 1040 pixels to 898 × 945 pixels. However, the region of interest for the TrA measurements is within a 895 × 745 pixels range. Therefore, we cropped the images to 895 × 745 pixels region as shown in Fig. 5. We manually defined the region of interest to exclude the unwanted parts of the images that included markers and text, which were positioned outside the region of interest, i.e. the TrA muscle in the middle of the image.

All 400 images of the TrA400 dataset were split into 10 folds of 3 subsets including training set ( $n = 320$ ), validation set ( $n = 40$ ), and test set ( $n = 40$ ), as depicted in Fig. 6. Using this cross-validation method ensures that each validation or test set has been used only once during training in each iteration and never used before in any other iterations. In addition, this method tests the proposed algorithm on all of the 400 images and produces the average performance as a more reliable evaluation metric.

### B. Implementation Details

Our method uses an Imagenet [33] pre-trained VGG16 FCN8 network [14]. Images were first converted to unsigned 8-bit integers, then resized to a fixed dimension 512 × 512 pixels and normalized using ImageNet statistics. The predicted measurements are based on these resized images. For each image in the training phase, we applied the following: random contrast, random brightness and random gamma transformations. The training was carried out on an Ubuntu 20.04.1 LTS machine with an NVIDIA GeForce RTX 2080 TI Graphics Card. The model was trained using PyTorch 1.5 with Python 3.7, a batch size of 1 for 100 epochs. The inference time per image is 47 milliseconds excluding pre/post-processing steps, which take 12 milliseconds. The overall processing time of 59 milliseconds is well within the requirements of clinical use for automatic muscle measurements. ADAM [44] was used with an initial learning rate of  $10^{-4}$ , and a learning rate decay of 0.1 after every 30 epochs. To overcome the small size problem of the TrA400 dataset, cross-validation was used to report performance across 10 folds, where all images in the dataset are measured using the proposed network and one measurement prediction is generated for each of the 400 images.

## V. RESULTS

To fully evaluate our model and compare it with other methods, we provide its experimental results along with a baseline landmark detection approach and a LCFCN approach. Our shown results in Table I are obtained from four fully-supervised networks using point-level annotations. These include

- **Landmark:** As a baseline, the standard landmark detection approach has been used similar to the work done in [8]. With an ImageNet [33] pre-trained ResNet50 [45] backbone, we regress the point locations using a FCN to localize landmarks.
- **HigherHRNet** [46]: Is a bottom-up method for 2D human pose estimation that aims to localise human anatomical keypoints. It learns scale-aware representations using

a high-resolution feature pyramid. This approach can solve the scale variation challenge in bottom-up multi-person pose estimation and localise key points more precisely. In this study, we have used the same approach to localise the caliper endpoints in US images by generating high-resolution spatial heatmaps (see Fig. 8).

- **LCFCN** [43]: LCFCN is based on a semantic segmentation architecture that is similar to FCN [14]. LCFCN method optimizes a loss function that ensures that only a single small blob is predicted around the centre of each object (in our case a caliper point) to stop the model from predicting large blobs that merge several object instances. We take the centroid of each blob to get coordinate locations and measure the distances between coordinate pairs to measure the TrA thickness.
- **LCFCN+CoordConv:** We improved LCFCN [43] prediction by adding a CoordConv [16] layer to the beginning of the FCN8 [14] architecture. The image is first passed through a CoordConv layer to add pixel-wise spatial location information to it. It then passes through the FCN8 layers to allow the CNN to locate caliper endpoints in the image, more efficiently.

### A. Evaluation Procedure

We evaluate our models using Mean Absolute Error (MAE) and standard deviation ( $\sigma$ ) across 10 folds, and Muscle Width Agreement (MWA), as performance metrics.

**MAE** measures the deviation of the predicted and ground-truth endpoints, and is described using (6).

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{E}_i - E_i|, \quad (6)$$

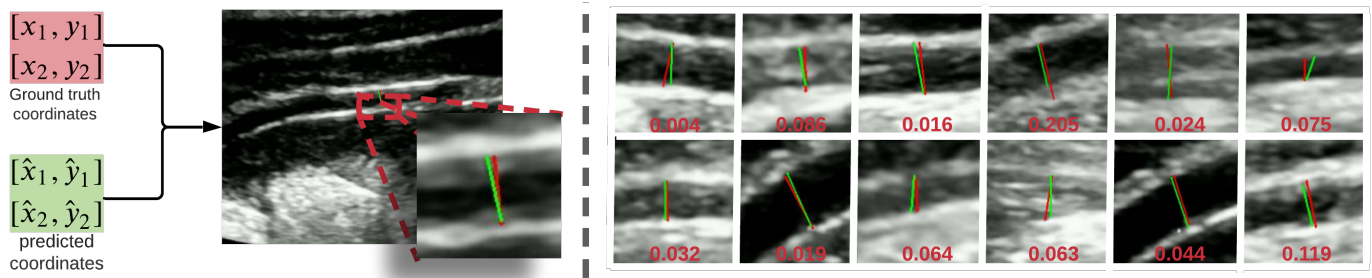
where  $\hat{E}_i$  and  $E_i$  are the predicted and ground truth endpoint coordinates for image  $i$ , respectively.  $N$  is the total number of images considered.

**The standard deviation** ( $\sigma$ ) shows the standard deviation between predicted and ground-truth endpoints across 10 folds.

**MWA** [8] is a metric used to compare human measurements against CNN-predicted measurements. To calculate MWA, here we define an error measurement formula that compares the human-annotated and machine-predicted lines connecting two caliper endpoints using (7)

$$\text{MWA} = \frac{1}{N} \sum_{i=1}^N \left| \frac{\|\bar{t} - \bar{b}\| - \|\hat{t} - \hat{b}\|}{\|\bar{t} - \bar{b}\|} \right|, \quad (7)$$

where  $\bar{t}$  and  $\bar{b}$  are top and bottom ground truth caliper endpoints, respectively, and  $\hat{t}$  and  $\hat{b}$  are top and bottom predicted endpoints, respectively. The error is computed as the sum of the absolute length difference between human measurements and CNN predicted measurements. The 50th, 75th, and 95th error percentiles are tabulated in Table I to show the distribution of errors and evaluate complex predictions more directly. The thickness range of RTrA is between 1.05 and 5.24 mm, while the average is 3.03 mm. The thickness range of CTrA is between 2.07 and 8.02 mm, while the average is 5.25 mm. Figure 7 shows the qualitative results to compare



**Fig. 7:** **Left:** Qualitative results. Here, the prediction results obtained from our novel model with the LCFCN loss and CoordConv layer is shown. **Right:** Sample prediction results from our model compared with ground-truth labels annotated by human experts. The predictions (green line) are in reasonably good agreement with the ground-truth labels (red line). Muscle Width Agreement (MWA) score for each image is shown in red.


Method	MWA			MAE	$\sigma$
	50th	75th	95th	Average	Average
Landmark	6.483	11.563	13.193	0.765	2.454
HigherHRNet	4.025	4.854	7.391	0.415	1.538
LCFCN	3.216	4.976	6.987	0.354	1.604
LCFCN + CoordConv	<b>2.710</b>	<b>4.803</b>	<b>6.641</b>	<b>0.312</b>	<b>1.499</b>

**TABLE I:** MWA, MAE, and standard deviation ( $\sigma$ ) metrics for the test sets across all folds. MWA was computed by comparing human and CNN measurement (7) for different percentiles (50th, 75th, 95th). MAE and  $\sigma$  are the average across 10 folds.

human measurements against CNN-predicted measurements. As the figure demonstrates, the predictions obtained from our proposed network (green line) are in reasonably good agreement with the ground-truth labels (red line) annotated by human experts.

Overall, the results summarized in Table I shows that the proposed LCFCN + CoordConv network outperforms previous localization networks, achieving lower MAE and MWA on the average of all cross-validation test subsets. It is inspiring that close to expert level performance was achieved with a small training dataset. We believe that the accuracy of our FCNs would increase if they are trained on a larger dataset. Overall, LCFCN with CoordConv layer improves results for all test subsets. The final prediction accuracy at the 50th percentile is 2.710% of the measurement length which is within the average inter-observer error of 5% [47], [48].

Statistical significance paired  $t$ -test was carried out to better compare the human measurements and our proposed method's predicted values of muscle thickness. Here, our null hypothesis ( $H_0$ ) was that the true mean difference between the human and our method's measurements is zero. The alternative hypothesis ( $H_A$ ), on the other hand, was that the true mean difference between the two measurements is not equal to zero. The values used in the two-tailed  $P$ -value calculations are  $t$ -score = 0.1371, degrees of freedom = 399, and mean difference of 0.055. The critical two-tailed  $t$ -value was 1.9659 and the calculated  $P$ -value was 0.891021. These show that there is not enough evidence to reject the null hypothesis. This decision is made at significance level  $\alpha = 0.05$  and 95% confidence interval.



**Fig. 8:** A sample US image (left) and its corresponding predicted scale-aware high-resolution heatmaps [46] (right).

## VI. DISCUSSION

The main goal of this study was to investigate various DL approaches and find the best method that can facilitate automatic but accurate measurement of the abdominal muscle thickness in 2D US images. An initial naive method that comes to mind is to develop a DL network that is trained using image-thickness pairs to predict a thickness value for a given ultrasound image. However, this approach cannot be trusted and explained because there maybe no meaningful relationship between the actual muscle thickness and the network-predicted value and the learned task will not certainly be generalizable to similar abdominal ultrasound images. To make sure we have

better explainability of the network outcome, our investigations targeted a way that performs the measurement similar to a human operator, i.e. finding key-point locations of the abdominal muscle to predict its thickness. Hence, we opted to treat the task at hand as a landmark detection problem.

Intrigued by the recent success of the FCN-based deep learning architecture [14] in segmentation problems, we modified this architecture for localization by incorporating LCFCN loss that extend the segmentation loss to perform localization with point supervision [43]. Unlike U-net [26] that requires the full per-pixel instance segmentation labels, FCN uses point-level annotations only. This enabled us to provide localization training points, which in our case are caliper points on an abdominal muscle. We also added a CoordConv [16] layer, which introduces two extra channels to the network and enables the CNN to learn the order of convolutional filters through coordinates, in order to properly learn translation equivairance [16].

After extensively tuning our customized network for the measurement task, our localization FCN achieved competitive localization performance in the experiments even with a limited training dataset. We evaluated the localization performance of our proposed FCN method on 10 folds test set that contains 2D US abdominal muscle images of different anatomical structures, demonstrating the generic applicability of our proposed method. Our results show that when training on ( $n = 320$ ) training set images, our method achieved close to expert ultrasound technicians level performance [47], [48]. We only used a limited number of training set images. A greater number of training set images could lead to a higher generalization of our results.

The 10-fold cross-validation technique used in this paper is one way to overcome the limited number of dataset images and improve prediction credibility over the holdout method. Firstly, the dataset is divided into 10 subsets, and the holdout method is repeated 10 times (see Fig. 6). Secondly, each time, one of the 10 subsets is used as the test set and the other 9 subsets are used to form training and validation sets. Finally, the average error across all 10 predictions is computed. The prediction for each single test subsets was consistent with the average of all cross-validation test subsets. This confirms that our model appropriately generalizes to new data.

A limitation of this study are that all the still images were extracted only from eligible participants, in a single site and with a single acquisition protocol [5]. To avoid overfitting to a single acquisition protocol, exams should be obtained on several sites. Thus, future studies are required to ensure the following:

- The dataset to be collected from several sites;
- All measurements should be performed by a single technician experienced in 2D US measurements;
- The US images to be labelled multiple times by the same expert;
- Consistent US image quality is achieved in all measurements.

As mentioned in Section II, so far no one appears to have applied FCN variants such as an LCFCN or LCFCN+CoordConv for US image muscle thickness measurements. The importance

of our results on using such networks thus lies both in their generality and their relative ease of application to automatic measurements of ultrasound images. We hope that our contribution can be used and extended for other measurement tasks across other medical imaging domains where a localization and/or measurement task is at hand.

## VII. CONCLUSION

In this work, we used the DL localization approach to automate measuring abdominal muscle dimensions in 2D ultrasound imaging. We proposed to detect measurement key-point locations by computing their localized estimates with a FCN. To improve localization performance, we added a CoordConv layer to FCN and incorporated the LCFCN loss function. This novel design resulted in errors below 3% of MWA between the measurements by human and the CNN predictions, which is within normal inter-observer variability. As this measurement is one of the outcome measures used to evaluate abdominal muscle dimensions to aid rehabilitation, our developed tool can assist by automatic measurements of ultrasound images, while reducing the inter-rater variability. This tool can provide more reliable diagnostic and outcome measurements, benefiting both patients and healthcare professional.

## ACKNOWLEDGEMENT

This research is supported by an Australian Research Training Program (RTP) Scholarship. MRA would like to thank Dr Carla Ewels for fruitful discussion regarding the statistical significance of the proposed method compared to human gold standard.

## REFERENCES

- [1] H. So, J. Chen, B. Yiu, and A. Yu, "Medical ultrasound imaging: To gpu or not to gpu?" *IEEE micro*, vol. 31, no. 5, pp. 54–65, 2011.
- [2] A. van Randen, W. Laméris, H. W. van Es, H. P. van Heeswijk, B. van Ramshorst, W. ten Hove, W. H. Bouma, M. S. van Leeuwen, E. M. van Keulen, P. M. Bossuyt *et al.*, "A comparison of the accuracy of ultrasound and computed tomography in common diagnoses causing acute abdominal pain," *European radiology*, vol. 21, no. 7, pp. 1535–1545, 2011.
- [3] G.-V. Vera, F. Aleksandra, K. Dragan, and H. Andrija, "Assessment of genome damage in occupational exposure to ionising radiation and ultrasound," *Mutation Research/Genetic Toxicology and Environmental Mutagenesis*, vol. 395, no. 2-3, pp. 101–105, 1997.
- [4] S. L. Koppenhaver, J. J. Hebert, E. C. Parent, and J. M. Fritz, "Rehabilitative ultrasound imaging is a valid measure of trunk muscle size and activation during most isometric sub-maximal contractions: a systematic review," *Australian Journal of Physiotherapy*, vol. 55, no. 3, pp. 153–169, 2009.
- [5] V. L. Kennedy, C. A. Flavell, and K. Doma, "Intra-rater reliability of transversus abdominis measurement by a novice examiner: comparison of "freehand" to "probe force device" method of real-time ultrasound imaging," *Ultrasound*, vol. 27, no. 3, pp. 156–166, 2019.
- [6] C. A. Flavell, L. G. Marshman, and S. J. Gordon, "Measurement of transversus abdominis activation in chronic low back pain patients using a novel standardized real-time ultrasound imaging method," *Ultrasound*, vol. 27, no. 1, pp. 31–37, 2019.
- [7] C. W. Hoppes, A. D. Sperier, C. F. Hopkins, B. D. Griffiths, M. F. Principe, B. L. Schnall, J. C. Bell, and S. L. Koppenhaver, "Ultrasound imaging measurement of the transversus abdominis in supine, standing, and under loading: a reliability study of novice examiners," *International journal of sports physical therapy*, vol. 10, no. 6, p. 910, 2015.

- [8] M. Sofka, F. Milletari, J. Jia, and A. Rothberg, "Fully convolutional regression network for accurate detection of measurement points," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2017, pp. 258–266.
- [9] A. Gilbert, M. Holden, L. Eikvil, S. A. Aase, E. Samset, and K. McLeod, "Automated left ventricle dimension measurement in 2d cardiac ultrasound via an anatomically meaningful cnn approach," in *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis*. Springer, 2019, pp. 29–37.
- [10] K. O'Shea and R. Nash, "An introduction to convolutional neural networks," *arXiv preprint arXiv:1511.08458*, 2015.
- [11] S. Albawi, T. A. Mohammed, and S. Al-Zawi, "Understanding of a convolutional neural network," in *2017 International Conference on Engineering and Technology (ICET)*. IEEE, 2017, pp. 1–6.
- [12] R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi, "Convolutional neural networks: an overview and application in radiology," *Insights into imaging*, vol. 9, no. 4, pp. 611–629, 2018.
- [13] M. Lin, Q. Chen, and S. Yan, "Network in network," *arXiv preprint arXiv:1312.4400*, 2013.
- [14] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," *CVPR*, 2015.
- [15] M. Drozdal, E. Vorontsov, G. Chartrand, S. Kadoury, and C. Pal, "The importance of skip connections in biomedical image segmentation," in *Deep Learning and Data Labeling for Medical Applications*. Springer, 2016, pp. 179–187.
- [16] R. Liu, J. Lehman, P. Molino, F. P. Such, E. Frank, A. Sergeev, and J. Yosinski, "An intriguing failing of convolutional neural networks and the coordconv solution," in *Advances in Neural Information Processing Systems*, 2018, pp. 9605–9616.
- [17] M. Xian, Y. Zhang, H.-D. Cheng, F. Xu, B. Zhang, and J. Ding, "Automatic breast ultrasound image segmentation: A survey," *Pattern Recognition*, vol. 79, pp. 340–355, 2018.
- [18] S. Y. Shin, S. Lee, I. D. Yun, S. M. Kim, and K. M. Lee, "Joint weakly and semi-supervised deep learning for localization and classification of masses in breast ultrasound images," *IEEE transactions on medical imaging*, vol. 38, no. 3, pp. 762–774, 2018.
- [19] T.-C. Chiang, Y.-S. Huang, R.-T. Chen, C.-S. Huang, and R.-F. Chang, "Tumor detection in automated breast ultrasound using 3-d cnn and prioritized candidate aggregation," *IEEE transactions on medical imaging*, vol. 38, no. 1, pp. 240–249, 2018.
- [20] I. E. Timor-Tritsch, M. N. Yeh, D. B. Peisner, K. B. Lesser, and T. A. Slavik, "The use of transvaginal ultrasonography in the diagnosis of ectopic pregnancy," *American journal of obstetrics and gynecology*, vol. 161, no. 1, pp. 157–161, 1989.
- [21] M. Cameron and P. Moran, "Prenatal screening and diagnosis of neural tube defects," *Prenatal Diagnosis: Published in Affiliation With the International Society for Prenatal Diagnosis*, vol. 29, no. 4, pp. 402–411, 2009.
- [22] C. F. Baumgartner, K. Kamnitsas, J. Matthew, T. P. Fletcher, S. Smith, L. M. Koch, B. Kainz, and D. Rueckert, "Sononet: real-time detection and localisation of fetal standard scan planes in freehand ultrasound," *IEEE transactions on medical imaging*, vol. 36, no. 11, pp. 2204–2215, 2017.
- [23] N. J. Cronin, T. Finni, and O. Seynnes, "Fully automated analysis of muscle architecture from b-mode ultrasound images with deep learning," *arXiv preprint arXiv:2009.04790*, 2020.
- [24] R. J. Cunningham, P. J. Harding, and I. D. Loram, "Real-time ultrasound segmentation, analysis and visualisation of deep cervical muscle structure," *IEEE transactions on medical imaging*, vol. 36, no. 2, pp. 653–665, 2016.
- [25] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in *2016 fourth international conference on 3D vision (3DV)*. IEEE, 2016, pp. 565–571.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [27] C. Payer, D. Štern, H. Bischof, and M. Urschler, "Regressing heatmaps for multiple landmark localization using cnns," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2016, pp. 230–238.
- [28] R. JG van Sloun and L. Demi, "Deep learning for automated detection of b-lines in lung ultrasonography," *The Journal of the Acoustical Society of America*, vol. 144, no. 3, pp. 1668–1668, 2018.
- [29] X. Liu, T. Zhou, M. Lu, Y. Yang, Q. He, and J. Luo, "Deep learning for ultrasound localization microscopy," *IEEE Transactions on Medical Imaging*, 2020.
- [30] M. Fayyaz, M. H. Saffar, M. Sabokrou, M. Fathy, R. Klette, and F. Huang, "Stfcn: spatio-temporal fcn for semantic video segmentation," *arXiv preprint arXiv:1608.05971*, 2016.
- [31] A. B. S. Nayem, A. Sarker, O. Paul, A. Ali, M. Amin, A. Rahman *et al.*, "Lulc segmentation of rgb satellite image using fcn-8," *arXiv preprint arXiv:2008.10736*, 2020.
- [32] W. Zhao, H. Zhang, Y. Yan, Y. Fu, and H. Wang, "A semantic segmentation algorithm using fcn with combination of bslic," *Applied Sciences*, vol. 8, no. 4, p. 500, 2018.
- [33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, 2015.
- [34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [35] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler, "Efficient object localization using convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 648–656.
- [36] S. Lee and C. Lee, "Revisiting spatial dropout for regularizing convolutional neural networks," *Multimedia Tools and Applications*, pp. 1–13, 2020.
- [37] E. Celledoni, M. J. Ehrhardt, C. Etmann, B. Owren, C.-B. Schönlieb, and F. Shery, "Equivariant neural networks for inverse problems," *arXiv preprint arXiv:2102.11504*, 2021.
- [38] I. H. Laradji, N. Rostamzadeh, P. O. Pinheiro, D. Vázquez, and M. Schmidt, "Instance segmentation with point supervision," *arXiv preprint arXiv:1906.06392*, 2019.
- [39] I. Laradji, P. Rodriguez, O. Manas, K. Lensink, M. Law, L. Kurzman, W. Parker, D. Vazquez, and D. Nowrouzezahrai, "A weakly supervised consistency-based learning method for covid-19 segmentation in ct images," in *wacv*, 2021, pp. 1–10.
- [40] I. H. Laradji, R. Pardin, P. Rodriguez, and D. Vazquez, "Looc: Localize overlapping objects with count supervision," in *icip*, 2020, pp. 1–10.
- [41] I. H. Laradji, D. Vazquez, and M. Schmidt, "Where are the masks: Instance segmentation with image-level supervision," in *bmvc*, 2019, pp. 1–10.
- [42] A. Saleh, I. H. Laradji, D. A. Konovalov, M. Bradley, D. Vazquez, and M. Sheaves, "A realistic fish-habitat dataset to evaluate algorithms for underwater visual analysis," *Scientific Reports*, vol. 10, no. 1, pp. 1–10, 2020.
- [43] I. H. Laradji, N. Rostamzadeh, P. O. Pinheiro, D. Vazquez, and M. Schmidt, "Where are the blobs: Counting by localization with point supervision," *ECCV*, 2018.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [45] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [46] B. Cheng, B. Xiao, J. Wang, H. Shi, T. S. Huang, and L. Zhang, "Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5386–5395.
- [47] C. Pirri, S. Todros, C. Fede, S. Pianigiani, C. Fan, C. Foti, C. Stecco, and P. Pavan, "Inter-rater reliability and variability of ultrasound measurements of abdominal muscles and fasciae thickness," *Clinical Anatomy*, vol. 32, no. 7, pp. 948–960, 2019.
- [48] A. Wilson, J. A. Hides, L. Blizzard, M. Callisaya, A. Cooper, V. K. Srikanth, and T. Winzenberg, "Measuring ultrasound images of abdominal and lumbar multifidus muscles in older adults: A reliability study," *Manual therapy*, vol. 23, pp. 114–119, 2016.