

# Automatic Assessment of Full Left Ventricular Coverage in Cardiac Cine Magnetic Resonance Imaging With Fisher-Discriminative 3-D CNN

Le Zhang<sup>1</sup>, Ali Gooya<sup>2</sup>, Marco Pereañez, Bo Dong, Stefan K. Piechnik, Stefan Neubauer<sup>3</sup>, Steffen E. Petersen, and Alejandro F. Frangi<sup>4</sup>, *Fellow, IEEE*

**Abstract**—Cardiac magnetic resonance (CMR) images play a growing role in the diagnostic imaging of cardiovascular diseases. Full coverage of the left ventricle (LV), from base to apex, is a basic criterion for CMR image quality and is necessary for accurate measurement of cardiac volume and functional assessment. Incomplete coverage of the LV is identified through visual inspection, which is time consuming and usually done retrospectively in the assessment of large imaging cohorts. This paper proposes a novel automatic method for determining LV coverage from CMR images by using Fisher-discriminative three-dimensional (FD3D) convolutional neural networks (CNNs). In contrast to our previous method employing 2-D CNNs, this approach utilizes spatial contextual information in CMR volumes, extracts more representative high-level features, and enhances the discriminative capacity of the baseline 2-D CNN learning framework, thus, achieving superior detection accuracy. A two-stage framework is proposed to identify missing basal and apical slices in measurements of CMR volume. First, the FD3D CNN extracts high-level features from the

CMR stacks. These image representations are then used to detect the missing basal and apical slices. Compared to the traditional 3-D CNN strategy, the proposed FD3D CNN minimizes within-class scatter and maximizes between-class scatter. We performed extensive experiments to validate the proposed method on more than 5000 independent volumetric CMR scans from the UK Biobank study, achieving low error rates for missing basal/apical slice detection (4.9%/4.6%). The proposed method can also be adopted for assessing LV coverage for other types of CMR image data.

**Index Terms**—3D convolutional neural network, LV coverage, image-quality assessment, population image analysis, Fisher discriminant criterion.

## I. INTRODUCTION

LEFT ventricular (LV) cardiac anatomy and function are widely used in the field of cardiac medicine for diagnosis and monitoring disease progression and for assessing the patient's response to cardiac surgery and interventional procedures. Cardiac ultrasound (US) and cardiac magnetic resonance (CMR) imaging are arguably the most widespread techniques for diagnostic imaging of the heart. For population imaging studies, however, CMR remains the modality of choice. CMR is a single technique that provides access to cardiac anatomy and non-invasive measurements of cardiac function [1]. In large population imaging studies or assessment of patient cohorts from large clinical trials, the quantification of LV anatomy and function requires automatic image quality assessment and tools for image analysis. One basic criterion for cardiac image quality is LV coverage and detection of missing apical and basal CMR slices [2]. CMR may display incomplete LV coverage because of insufficient radiographer experience in planning a scan, natural cardiac muscle contraction, breathing motion, and imperfect triggering, all of which pose challenges in efforts at quantitative LV characterisation and accurate diagnosis [3]. For example, missing basal slices affect calculations of LV volume and derived LV functional measures such as ejection fraction and cardiac output. Even if scout images are acquired, in order to centre the LV in view and minimize this issue, incomplete coverage may result at any point throughout the cardiac cycle because of changes in patient breathing and cardiac motion. Image quality assessment is traditionally performed by radiographers who ensure that patients do not leave the scanner without providing diagnostically interpretable data. However, there are

Manuscript received February 9, 2018; revised July 3, 2018 and October 1, 2018; accepted November 6, 2018. Date of publication November 21, 2018; date of current version June 21, 2019. This work was supported in part by the Amazon, in part by the NVIDIA, and in part by the British Heart Foundation under Grant PG/14/89/31194. The work of L. Zhang was supported by the China Scholarship Council (Ph.D. studies). The work of M. Pereañez was supported by the VPH-DARE@IT FP7 EC Integrated Project (FP7-ICT-2011-9-601055). The work of S. K. Piechnik and S. Neubauer was supported in part by the National Institute for Health Research Oxford Biomedical Research Center, Oxford University Hospitals Trust, University of Oxford, and in part by the British Heart Foundation Center of Research Excellence. The work of S. E. Petersen was supported in part by the NIHR Biomedical Research Center at Barts, and in part by the SmartHeart EPSRC Programme under Grant EP/P001009/1. The work of A. F. Frangi was supported in part by the VPH-DARE@IT FP7 EC Integrated Project (FP7-ICT-2011-9-601055EPSR) and in part by the C-funded MedIAN Partnership under Grant EP/N026993/1. (Corresponding author: Le Zhang.)

L. Zhang is with the Centre for Computational Imaging and Simulation Technologies in Biomedicine, University of Sheffield, Sheffield, S1 3JD U.K. (e-mail: le.zhang@sheffield.ac.uk).

A. Gooya and B. Dong are with the Centre for Computational Imaging and Simulation Technologies in Biomedicine, University of Sheffield.

M. Pereañez and A. F. Frangi are with the Centre for Computational Imaging and Simulation Technologies in Biomedicine, School of Computing and School of Medicine, University of Leeds.

S. K. Piechnik and S. Neubauer are with the Oxford Centre for Clinical Magnetic Resonance Research, Division of Cardiovascular Medicine, John Radcliffe Hospital, University of Oxford.

S. E. Petersen is with the Cardiovascular Medicine at the William Harvey Research Institute, Queen Mary University of London, and also with the Barts Heart Centre, Barts Health NHS Trust.

Digital Object Identifier 10.1109/TBME.2018.2881952

limits to human attention. With CMR examinations becoming less expensive and increasingly commissioned, scanning loads at some centres may be insufficient to maintain consistent standards. Quality assessment is of particular importance in large-scale population imaging studies, where data are acquired across different imaging sites before core lab analysis. For example, large volumes of data may be stored without being checked by experienced staff prior to analysis [4], [5]. Automatic methods for these repetitive quality assurance tasks provide the required consistency and reliability.

To ensure consistent quantification of CMR data, automatic assessment of complete LV coverage is the first step. LV coverage is assessed by visual inspection of CMR image sequences, which is a subjective, repetitive, error-prone, and time-consuming process [6]. Automatic coverage assessment is required to promptly intervene and correct data acquisition, and/or discard images with incomplete LV coverage whose analysis would otherwise impair any statistics aggregated over the cohort. The most common causes of incomplete LV coverage are lack of a basal slice (no atrial chamber visible in end-systole, hence no certainty that the base of the heart is completely covered) and lack of an apical slice (LV cavity remains visible at end-systole). According to the criteria used in [2] for CMR quality assessment, a missing basal slice carries a higher penalty than a missing apical slice, given its impact on LV volume computation. Although technological developments in magnetic resonance imaging (MRI) hardware and pulse sequences have led to faster CMR acquisitions, challenges remain with regard to ensuring full heart coverage and motion compensation. In the UK Biobank's CMR protocol, for instance, incomplete heart coverage is the reason for flagging 4% of all CMR examinations as providing unreliable or non-analysable image data [7]. While 4% may seem to be a small proportion, the challenge is to automatically sift through the entire database to identify and exclude those cases from further quantitative analysis. Methods for the objective detection of basal and apical imaging planes are relevant in this context, as their absence affects diagnostic accuracy as well as anatomical and functional LV quantification.

In the field of video processing, Automatic Image Quality Assessment (AIQA) is a well-developed corpus of techniques concerned with detecting image distortions characteristic of multimedia communications [8], [9]. These distortions generally differ from those affecting medical images. No-reference-based image quality assessment (NR-IQA) [10], [11] is relevant for medical imaging data. While there is relatively easy access to abundant data sets of mixed quality, it is not possible to collect data without some level of image degradation or artefacts. Practical CMR image-processing applications do not provide perfect versions of incomplete LV coverage images, but rather, only the image to be assessed. While assessments attempt to highlight differences in our assessed data set regarding a hypothetical high-quality image [12], the final image quality is estimated solely based on the characteristics of the assessed image.

The current standard operating procedure in the UK Biobank, for instance, involves the detection of missing basal/apical slices based on visual assessment by experts. Few methods have been developed for automating this process, and prior work mostly

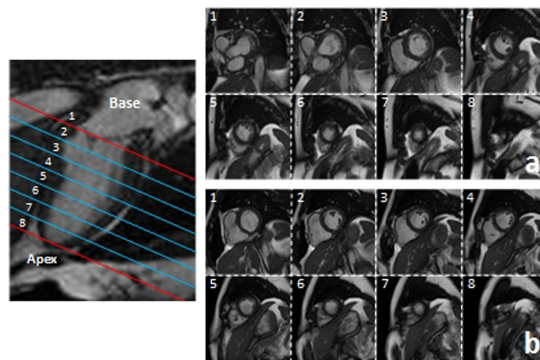


Fig. 1. Left: A typical two-chamber view cardiac MRI with eight slices covering from base to apex. Right: (a) a volume with whole coverage (slice 1 is the basal slice), and (b) a volume with missing basal slice (slice 1 is not the basal slice). In each rectangle, from top to bottom, rows correspond to adjacent axial slices.

adopted approaches that require segmenting short-axis slices of LV [13], [14] or landmark localization [15]–[17]. However, fast full LV coverage detection as the first step of an image quantification pipeline is largely unexplored. Hoffmann *et al.* pioneered this field [15] by initially localizing the heart in raw data prior to applying computer-aided diagnosis algorithms. Lu *et al.* [18] proposed an approach to locate LV and prescribe long/short-axis views before MR image acquisition, which could be used to evaluate cardiac coverage in short-axis views. These methods detect missing basal/apical slices and largely rely on the quality of LV segmentation and localization. de Vos *et al.* [19] proposed a method that automatically identifies a slice of interest (SOI) in 3D images. A ConvNet regressor was trained to determine the distance between each 2D slice and the SOI. However, this solution does not consider 3D contextual information contained across slices.

The characteristics of the LV are useful in identifying the position that the slice belongs to, since the LV in each slice shows a different shape and size. For example, the LV shape is approximately circular in mid-slices, while it is more elliptical in basal slices (Fig. 1). Recent work [20], [21] has focused on learning data-driven features to more accurately detect shape differences. Among them, 3D convolutional neural networks (CNNs) are one of the most regularly used deep-learning schemes to meet the challenges of discriminative shape detection [22], [23]. Roth *et al.* [24] and Prasoos *et al.* [25] adapted 2D CNNs for processing 3D volumetric data. However, these studies reported having difficulties when attempting to employ 3D CNN on their data, since they often lack sufficient training samples and computational resources to learn accurate 3D models. Although some authors [26], [27] have utilized 3D CNNs to process medical images, their architectural settings, convolution kernels, and prediction score volumes have not been disclosed in the detail required to reproduce their results [28]. Some exceptions, however, include the work of Kamnitsas *et al.* [29], who devised an effective dense training scheme based on 3D CNNs for brain lesion segmentation and dealing with the computational burden of processing 3D medical scans. Moreover, the 3D U-Net architecture of Cicek *et al.* [30] takes 3D volumes

as input and produces volumetric image segmentation. The architecture and data augmentation of the U-net allow learning models with very good generalization performance from only a few annotated samples. Owing to the success of 3D deep neural networks in medical image segmentation, we are motivated to devise an end-to-end network optimization without requiring manual annotations of the visual image quality. Meanwhile, we seek features maximally affected by partial image artefacts, which are also not very sensitive to variability related to the intrinsic anatomy or image modality at hand.

In this paper, we focus on the analysis of short-axis (SA) cine MRI, although the technique can also be generalized to long-axis images. We aim to identify missing apical slices (MAS) and/or basal slices (MBS) in 3D cardiac MRI volumes. In our previous work, we used a 2D CNN constructed on single-slice images and processed them sequentially [31]. However, this solution ignores contextual information contained across slices providing inferior performance compared to a 3D analysis. We assume that 3D CNNs can easily and effectively deal with within-class variability and between-class similarity, which are important sources of the detection error [32]. We seek to learn a feature representation that achieves reliable classification results even with a small amount of training data or a small number of iterations. In this paper, we address incomplete LV coverage detection using a Fisher-discriminative 3D (FD3D) CNN, which utilizes 3D convolution kernels and exploits the spatial contextual information in volumetric data. The proposed FD3D CNN uses the Fisher discriminant criterion [33] on the fully connected layer to render features more discriminative and insensitive to geometric structural variations.

To the best of our knowledge, this is the first study tackling the problem of automatic detection of missing basal and apical slices on a CMR dataset as extensive and challenging as the UK Biobank. Besides introducing a novel FD3D CNN architecture, we propose an effective cascaded detection strategy for incomplete coverage identification. In the first stage, we train two separate FD3D CNN classifiers to detect the absence of basal and apical slices. In the second stage, we combine the classification results from stage 1 to determine the type of incomplete coverage found on the image.

The rest of this paper is organized as follows. Section II introduces the proposed FD3D CNN architecture and explains the learning strategy for its parameters. Section III presents experimental materials and metrics. Section IV describes the experimental design and classification results. Further analysis and discussion of the proposed method are provided in Section V. Conclusions are presented in Section VI.

## II. FULL LV COVERAGE DETECTION METHOD

### A. Problem Formulation

During image acquisition, a sufficient margin ought to be left above and below the LV cavity according to the established guidelines [34]. However, some image volumes may lack sufficient information at the apical and basal levels, which can hamper or bias the subsequent statistical analysis of cardiac

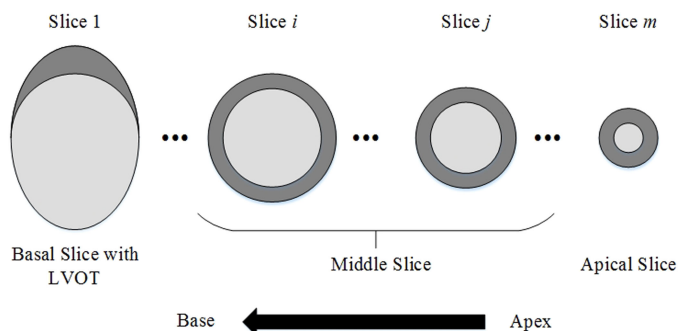


Fig. 2. Schematic LV shapes showing blood pool (light grey) and myocardium (dark grey) for different slices from apex to base. Slice 1 (left) shows LVOT, which identifies the basal slice.

structural and functional parameters in population imaging [35], [36]. In many LV quantification approaches, the LV cross section is approximated using simple quasi-circular models [37], [38]. These methods can produce a good approximation on LV mid-slices, but not on slices containing the left ventricular outflow tract (LVOT), which is at or near the basal slice. Therefore, in our approach, we treat the blood pool cross-section as a distinct model. Fig. 2 depicts the LV shape of several slices in one cardiac volume from the apex to base. In volumes with missing basal slice, LVOT is usually not present.

We use a vector  $\mathbf{s}$  to represent pixel values in each slice. A 3D cardiac MRI volume  $\mathbf{V}$  with full coverage with  $n$  slices can be described as follows:

$$\mathbf{V} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n]. \quad (1)$$

Each cardiac volume,  $\mathbf{V} = [\mathbf{s}_p, \dots, \mathbf{s}_q], p \leq q \in [1, n]$ , can have a different or same number of slices but cover a different portion of the LV.

To guarantee accurate cardiac volumetry and functional measurements [2], full LV coverage is a basic requirement [36]. To address this problem, we propose a two stage detection system that first computes image intensity representations by a FD3D CNN model and then detects missing slices based on these representations. In the first stage, we encode spatial contextual information and hierarchically extract high-level features, which indicate intensity representations. Our FD3D CNN model is equipped with a fully connected Fisher discriminative layer (F2) that takes the output of the fully connected layer (F1) as input. In the second stage, independent detection of any missing basal and apical slices is performed and the results are combined to provide the final coverage assessment.

### B. Three-Dimensional Intensity Representations

Lu *et al.* [39] proposed a pattern recognition technique built on intra-segment correlation, using a normalization scheme, which maps each LV slice to polar coordinates with fixed size, shape level, and position. Intensity information and slice position are relevant even with incomplete LV coverage detection. In our paper, we define intensity representation for the missing slice in a high-level feature space where slices of cardiac MRI are used to construct a representation of intensity. Each slice of the



3D volume is accounted for and the similarity of neighboring slices determines the difference of the 3D intensity distribution. Different characteristics in each slice and contextual information about spatial relation between slices are used to compute intensity representations.

*Which 3D intensity representations?* Our intensity representations are computed as a feature distribution matrix, which integrates information about LV shape and size. We detect incomplete LV coverage by image classification using the distribution matrix. We define two classes: missing apical slice (MAS) and missing basal slice (MBS).

Given a particular describable visual representation, we can formalize our notion of 3D intensity representations based on Eq. (1). For example, if we are looking at the volume from base to apex, MAS and MBS can be formalized as follows:

$$\begin{cases} \mathbf{V}_{MBS} = [\mathbf{s}_q, \dots, \mathbf{s}_n], \\ \mathbf{V}_{MAS} = [\mathbf{s}_1, \dots, \mathbf{s}_p], \end{cases} \quad (2)$$

where,  $p, q \in (1, n)$ ,  $\mathbf{s}_1$  is the basal slice and  $\mathbf{s}_n$  is the apical slice. Our intensity representations classifiers can be thought of functions  $f(\cdot)$  for mapping 3D stacks  $\mathbf{V}$  to real value  $p_i$ . A positive value of  $p_i$  indicates the presence or strength of the  $i$ th representation, while negative values indicate its absence. Considering our intensity representations, if we define  $\mathbf{V}_1$  and  $\mathbf{V}_2$  as MBS and non-MBS samples, respectively, the representation function  $f_{MBS}(\cdot)$  may map  $\mathbf{V}_1$  to a positive value and  $\mathbf{V}_2$  to a negative value. This is a binary classification function. Our 3D intensity representation classifiers are trained on the UK Biobank dataset as they provide reliable ground-truth labels based on visual inspection and manual annotation.

### C. Fisher Discriminative 3D CNN Model

In this subsection, we propose a FD3D CNN (shown in Fig. 3b) to extract high-level features, which represent 3D intensity representations. Our FD3D CNN model is designed by adding a new Fisher-discriminative fully connected layer, F2, which uses the output of the previous layer, F1, as input. The new layer is then stacked onto a conventional 3D CNN. To maximize inter-class distances between learned features while minimizing intra-class distances of learned features, we train the newly added Fisher discriminative layer F2 on CNN features based on a Fisher discriminant criterion [33].

**1) 3D CNN:** Learning feature representations in three dimensions is important for later feature detection and image interpretation tasks in volumetric medical imagery. We employ 3D convolution kernels to encode richer spatial information in volumetric data. Here, feature maps are 3D blocks instead of 2D patches. Conventional 3D convolution is achieved by convolving a 3D kernel, with the cube formed by stacking multiple contiguous slices. With this construction, feature maps in the convolution layer are connected to multiple contiguous frames of the previous layer [40], [41]. Given an input  $\mathbf{v}_k^l$ , the 3D

convolution layer output equates to a filtering operation with a filter  $\mathbf{W}_{ik}^{l+1}$ . Computation of the 3D feature volume  $\mathbf{h}_i^{l+1}$  is given by:

$$\mathbf{h}_i^{l+1} = f \left( \sum_k \sum_{r=0}^{R-1} \sum_{s=0}^{S-1} \sum_{t=0}^{T-1} \mathbf{W}_{ik}^{l+1}(r, s, t) \mathbf{v}_k^l + b_k^{l+1} \right) \quad (3)$$

where  $\mathbf{W}_{ik}^{l+1}(r, s, t)$  is the element-wise weight in the 3D convolution kernel,  $\mathbf{W}_{ik}^{l+1}$  and  $b_k^{l+1}$  are the filter and bias terms connecting the feature maps of adjacent layers, and  $f(\cdot)$  is the element-wise, non-linear activation function.

**2) Fisher Discriminative 3D CNN:** To boost the discriminative power of 3D CNN learned features, we impose a Fisher discrimination criterion [33] on them. Given the 3D input data  $\mathbf{V}_i^t$ , where  $i$  is the representation class, with  $i = \{1, 2\}$ , corresponding to MAS and MBS; the superscript  $t$  in  $\mathbf{V}_i^t$  indicates whether the representation is positive or negative, i.e.,  $t = \{0, 1\}$ ;  $\mathbf{V}_i^t = [\mathbf{v}_{i,1}^t, \mathbf{v}_{i,2}^t, \dots, \mathbf{v}_{i,C}^t]$ ,  $\mathbf{v}_{i,j}^t$  is the input data of  $j$ th sample from class  $i$ , for  $j = 1, 2, \dots, C$ . We denote  $\mathbf{F}_{i,j}^t$  to be features in the fully-connected layer of the 3D CNN for class  $i$  and  $j$ th sample. Using the Fisher criterion, discrimination is achieved by minimizing within-class scatter of  $\mathbf{F}^t$ , denoted by  $S_w(\mathbf{F}^t)$ , and maximizing between-class scatter of  $\mathbf{F}^t$ , denoted by  $S_b(\mathbf{F}^t)$ .  $S_w(\mathbf{F}^t)$  and  $S_b(\mathbf{F}^t)$  are defined as follows:

$$S_w(\mathbf{F}^t) = \sum_{i=1}^I \sum_{\mathbf{F}_{i,j}^t \in \mathbf{F}^t} (\mathbf{F}_{i,j}^t - \mathbf{m}_i^t)(\mathbf{F}_{i,j}^t - \mathbf{m}_i^t)^T, \quad (4)$$

$$S_b(\mathbf{F}^t) = \sum_{i=1}^I n_i (\mathbf{m}_i^t - \mathbf{m}^t)(\mathbf{m}_i^t - \mathbf{m}^t)^T, \quad (5)$$

where  $\mathbf{m}_i^t$  and  $\mathbf{m}^t$  are mean vectors of  $\mathbf{F}_i^t$  and  $\mathbf{F}^t$ , respectively, and  $n_i$  is the number of samples from class  $i$ . The Fisher discriminant regularization term  $\Phi(\mathbf{F}^t)$  is defined as  $\text{tr}(S_w(\mathbf{F}^t)) - \text{tr}(S_b(\mathbf{F}^t))$ . To obtain a discriminative classification result with deep learned features, we propose modifying the objective function of the FD3D CNN model by inserting a Fisher discriminant regularization term:

$$\begin{aligned} \mathbf{J}^*(\mathbf{W}, \mathbf{b}) = \arg \min_{\mathbf{W}, \mathbf{b}} & \frac{1}{I} \sum_{i=1}^I y^t \log a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b}) \\ & + (1 - y^t) \log (1 - a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b})) \\ & + \frac{1}{2} \lambda \|\mathbf{W}\|_2^2 + \frac{1}{2} \eta (\text{tr}(S_w(\mathbf{F}^t)) \\ & - \text{tr}(S_b(\mathbf{F}^t))), \end{aligned} \quad (6)$$

where  $\mathbf{J}^*$  is our new cost function that can minimize within-class scatter and maximize between-class scatter, and  $y$  is the output label. Output activation  $a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b}) = 1/(1 + e^{-\mathbf{W}\mathbf{V}_{i,j}^t - \mathbf{b}})$  is typically restricted to the open interval  $(0, 1)$  by using a logistic sigmoid, which is parametrized by  $\mathbf{W}$  and  $\mathbf{b}$  on the  $j$ th training sample.  $\|\mathbf{W}\|_2^2$  is a penalty term to the loss function that prevents weights from getting too large and helps to prevent over-fitting. Weights in each layer can be adjusted toward target

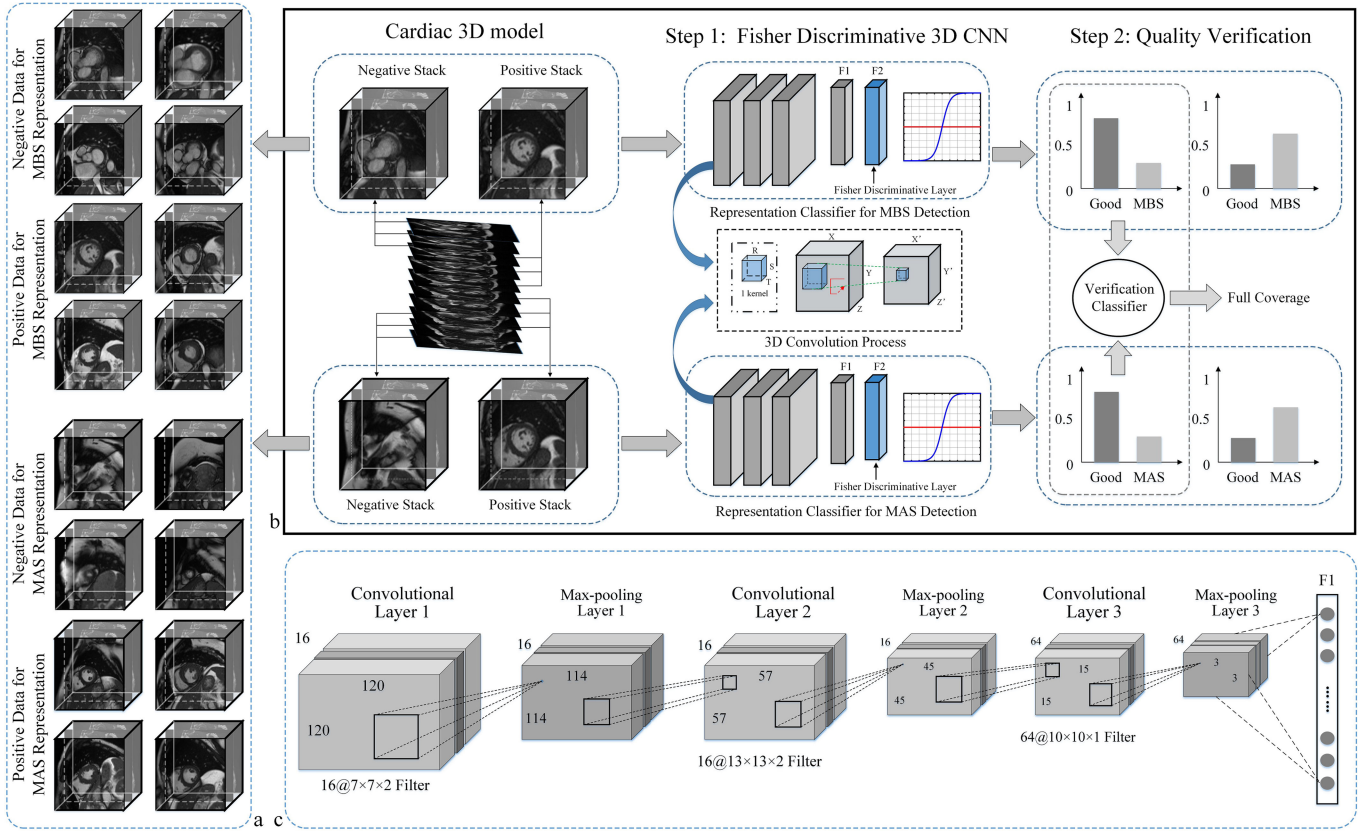


Fig. 3. Whole assessment framework. (a) Positive and negative training data for each representation classifier (MBS and MAS). (b) Framework for our LV coverage assessment process. (c) Structure and parameters of the 3D CNN used in panel (b): Step 1.

classes and utilize input data close to the corresponding classes in case of no large dataset or a small number of iteration. Here,  $\lambda, \eta \in [0, 1]$  are two trade-off parameters that control the relative importance of each term and are usually chosen by experiments, which can differ depending on different databases and network structures.

For intensity representation  $\mathbf{V}_{i,j}^t$ , we define:

$$\mathbf{J}(\mathbf{W}, \mathbf{b}) = y^t \log a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b}) + (1 - y^t) \log(1 - a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b})), \quad (7)$$

$$\Phi(\mathbf{F}^t) = \frac{1}{2} \text{tr}((\mathbf{F}_{i,j}^t - \mathbf{m}_i^t)(\mathbf{F}_{i,j}^t - \mathbf{m}_i^t)^T) - \frac{1}{2} \text{tr}((\mathbf{m}_i^t - \mathbf{m}^t)(\mathbf{m}_i^t - \mathbf{m}^t)^T). \quad (8)$$

Once the new cost function is obtained, we can employ the gradient descent method [42] to solve this optimization problem. Our key problem is to calculate the error of output units, which consists of output errors from two sub-functions  $\mathbf{J}(\mathbf{W}, \mathbf{b})$  and  $\Phi(\mathbf{F}^t)$ . To update parameters  $\mathbf{W}^t$  and  $\mathbf{b}^t$ , we first calculate the error  $\delta_i^{L,t}$  ( $L$  is the output layer) of the output layer with forward propagation, and then adopt the back-propagation method [43] to calculate the error  $\delta_i^{l,t}$  ( $l < L$ ) for other layers. Partial derivatives of the overall cost function  $\mathbf{J}^*(\mathbf{W}, \mathbf{b})$  regarding  $\mathbf{W}^t$  and

TABLE I  
ARCHITECTURE OF THE 3D DISCRIMINATIVE CNN MODEL

Layer	Kernel Size	Stride	Output size	Feature volumes
Input	–	–	$120 \times 120 \times 3$	1
C1	$7 \times 7 \times 2$	1	$114 \times 114 \times 2$	16
M1	$2 \times 2 \times 1$	2	$57 \times 57 \times 2$	16
C2	$13 \times 13 \times 2$	1	$45 \times 45 \times 1$	16
M2	$3 \times 3 \times 1$	1	$15 \times 15 \times 1$	16
C3	$10 \times 10 \times 1$	1	$6 \times 6 \times 1$	64
M3	$2 \times 2 \times 1$	1	$3 \times 3 \times 1$	64
F1	–	1	$1 \times 1 \times 1$	256
F2	–	1	$1 \times 1 \times 1$	4

Note: F2 is the Fisher Discriminant Layer.

$\mathbf{b}^t$  are:

$$\frac{\partial \mathbf{J}^*(\mathbf{W}, \mathbf{b})}{\partial \mathbf{W}^{l,t}} = \sum_{t=0}^C \sum_{F^t \in t} \frac{\partial \mathbf{J}(\mathbf{W}^t, \mathbf{b}^t)}{\partial \mathbf{W}^{l,t}} + \eta \sum_{t=0}^C \sum_{F^t \in t} \frac{\partial \Phi(\mathbf{F}^t)}{\partial \mathbf{W}}, \quad (9)$$

$$\frac{\partial \mathbf{J}^*(\mathbf{W}, \mathbf{b})}{\partial \mathbf{b}^{l,t}} = \sum_{t=0}^C \sum_{F^t \in t} \frac{\partial \mathbf{J}(\mathbf{W}^t, \mathbf{b}^t)}{\partial \mathbf{b}^{l,t}} + \eta \sum_{t=0}^C \sum_{F^t \in t} \frac{\partial \Phi(\mathbf{F}^t)}{\partial \mathbf{b}}. \quad (10)$$

In this stage, we use the 3D CNN model with architecture in Table I. Algorithm 1 provides the pseudo-code to train this new network. In our 3D CNN implementation, a rectifier linear

**Algorithm 1:** FD3D CNN Training.

**Input:** input-target pairs  $(\mathbf{v}_{i,j}^t, \mathbf{y}^t)$ , corresponding  $j^{th}$  pairs from class  $i$ ,  $t$  indicates positive or negative sample;  $\eta$ .

**Output:** FD3D CNN weight and biases, respectively,  
 $\mathbf{W} = [\mathbf{W}^{1,t}, \mathbf{W}^{2,t}, \dots, \mathbf{W}^{l,t}]$  and  
 $\mathbf{b} = [\mathbf{b}^{1,t}, \mathbf{b}^{2,t}, \dots, \mathbf{b}^{l,t}]$ .

**Begin**

Initialize  $\mathbf{W}_{i,j}^t$  and  $\mathbf{b}_{i,j}^t$

**while** *stopping criterion has not been met* **do**

1) Classification error:

$$\operatorname{argmin}_{\mathbf{W}, \mathbf{b}} \sum_{i=1}^I y^t \log a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b}) + (1 - y^t) \log(1 - a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b})).$$

2) Fisher discriminant:

$$\Phi(\mathbf{F}^t) = \operatorname{tr}(S_w(\mathbf{F}^t)) - \operatorname{tr}(S_b(\mathbf{F}^t)).$$

3) Discriminative objective function:

$$\operatorname{argmin}_{\mathbf{W}, \mathbf{b}} \sum_{i=1}^I y^t \log a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b}) + (1 - y^t) \log(1 - a(\mathbf{V}_{i,j}^t; \mathbf{W}, \mathbf{b})) + \frac{1}{2} \lambda \|\mathbf{W}^t\|_2^2 + \frac{1}{2} \eta \Phi.$$

4) Update  $\mathbf{W}_{i,j}^t$  and  $\mathbf{b}_{i,j}^t$  with Eqs. (9) and (10).

**end**

**return**  $\mathbf{W}_{i,j}^t$  and  $\mathbf{b}_{i,j}^t$  until values of  $\mathbf{J}^*(\mathbf{W}, \mathbf{b})$  in successive iterations are close enough or the maximum number of iterations is reached.

**End begin**

unit (ReLU) [44] is utilized as a non-linear activation function in layers C and F1.

### III. MATERIALS AND METRICS

#### A. CMR Acquisition Protocol and Annotation

*UK Biobank CMR Protocol:* UK Biobank's CMR acquisitions are performed on a clinical wide bore 1.5T scanner (MAGNETOM Aera, Syngo Platform VD13A; Siemens Healthcare, Erlangen, Germany) and include piloting, sagittal, transverse, and coronal partial coverage of the chest and abdomen. For measuring the cardiac function, three long-axis cines are acquired (viz. horizontal long-axis (HLA), vertical long-axis (VLA), and LVOT in both sagittal and coronal views). In addition, a complete SA stack is acquired. All acquisitions use balanced steady-state free precession (bSSFP) MRI sequences, attempting full coverage of the LV and right ventricle [45]. In this study, we will focus on SA bSSFP cine CMR data. To date, more than 18,800 volunteers have been scanned. Voxel and matrix size of these CMR images are, respectively,  $1.8 \times 1.8 \times 8.0 \text{ mm}^3$ , and  $208 \times 187$  with, approximately, 10 slices per volume. Each volumetric sequence contains about 50 cardiac phases.

*Gold-Standard Image Quality Annotations:* Quality-scored cardiac MRI data are available for approximately 5,000 volunteers of the UK Biobank (UKBB) imaging resource. Following visual inspection, manual annotation was carried out with a simple three-grade quality score [7]: (1) optimal quality for diagnosis, (2) suboptimal quality yet analysable and (3) bad quality and diagnostically unusable. In 5,065 SA cine CMR from the same number of volunteers, 4,361 sequences correspond to a quality score of 1, an additional 527 sequences have a quality score of

2, and the remaining 177 sequences have a quality score of 3. All datasets with optimal quality (score 1) had full coverage of the heart from base (LVOT existing) to apex (LV cavity still visible at end-systole). These data were used to construct the ground-truth classes for our experiments. Note that having full coverage should not be confused with having top/bottom slices corresponding exactly to the base/apex.

#### B. Training and Testing Set Definitions

*Training Set:* To create a training dataset for learning intensity representations, we extract the three topmost slices as negative samples for MBS detection (i.e. containing the cardiac base), and the three bottom most slices as negative samples for MAS detection. To create positive samples (i.e. not containing the cardiac base/apex), we choose three-slice blocks, each starting from the middle slice towards the base/apex for MBS/MAS detection training. We create the training set from images with optimal quality and with exclusively full coverage.

We train using three-slice stacks (or triplets) to model the 3D context. the average number of slices per image volume is approximately 10. During training, we extract four triplets (two samples including base/apex and two samples excluding the base/apex). To maximize inter-class separation, it is wise to avoid intersection between the training samples; for example, if we use four-slice stacks (for a ten-slice volume), there will be a two-slice overlap between basal positive/negative examples and the apical region. By choosing the proposed slice triplets, we ensure that there is no overlap and increase the discriminative power of the FD3D CNN. Another important observation that supports the choice of slice triplets is that the CMR scan volume is not acquired immediately. Instead, each slice is collected over several cardiac cycles leading to some degree of slice-to-slice misalignment. This effect is minimized when considering only slice triplets in contrast to using the full 3D volume.

*Testing Set:* During testing, we extract every set of three adjacent slices from top to bottom for each volume and apply these triplets to intensity representation classifiers. Data with known MBS/MAS are created by manually removing the three topmost/bottom most slices from images with optimal quality, as in the training set.

During training and testing, three-slice stacks are input to the proposed FD3D CNN. Scores of the output layer can be interpreted as the probability that triplets correspond to negative or positive MBS/MAS. The final output is the combination of two CNN outputs (MBS and MAS). The three slice stacks input into our network are cropped centered images of dimensions  $120 \times 120 \times 3$  to extract the region of interest. Parameter setting of block-size determination is explained in Section IV-A.

#### C. Training Set Augmentation

To prevent over-fitting due to insufficient training data and to improve the detection rate of our algorithm, we employ data augmentation techniques to artificially enlarge our dataset [46], [47]. In our application, we augment the data by applying a discrete set of in-plane rotations and isotropic scalings to the training images. Unlike data augmentation choices made for



natural image datasets where variability in location and pose of objects are relatively high, our data are comparatively constrained due to standard imaging protocols and gross patient positioning on the MRI scanner. We therefore chose a set of realistic rotations and scaling factors for MRI. Based on analysis of the in-plane orientation angle distribution for 5,000 subjects for which manual segmentations were available (and therefore LVRV angle can be computed), we found that LVRV orientation ranges between  $-45^\circ$  and  $45^\circ$ . The set of rotations chosen was accordingly  $-45^\circ$  and  $45^\circ$ , with two scaling factors of 0.75 and 1.25. This increases the number of training samples by a factor of four, while not adding significantly to the convergence time.

After data augmentation, we constructed 845,000 3D stacks comprised of 2D CMR slices from 3,380 sequences each with 50 cardiac phases, with a quality score of 1. These data are used for experiments in Section IV-A, B, and C. We set aside 981 sequences and data with quality scores of 2 and 3 for later use, as described Section IV-D. In our experiments, 10-fold cross-validation [48] was used to evaluate the performance of our system. To the best of our knowledge, this is the largest annotated dataset available to date for automatic CMR quality assessment.

#### D. Learning Performance Metrics

To evaluate the learning process, we use the following established classification metrics:

$$\text{Precision} = TP / (TP + FP), \quad (11)$$

$$\text{Sensitivity} = TP / (TP + FN), \quad (12)$$

$$\text{Error Rate} = (FP + FN) / N, \quad (13)$$

where  $TP$ ,  $FP$  and  $FN$  are numbers of true-positive, false-positive and false-negative samples, respectively, and  $N$  represents the number of subjects in the test set.

### IV. EXPERIMENTS AND RESULTS

#### A. Performance Analysis

We experiment to characterize the performance of our FD3D CNN learning framework. The error (cost) functions used in learning (Eqs. [6] and [7]) remain within this range  $[0, 1]$ . In all experiments, the learning process was terminated when standard deviation of the error function over the last five iterations is smaller than  $\sigma = 0.01$ .

**1) Hyper-Parameter Selection:** LeCun *et al.* [49] and Salah *et al.* [50] used CNN to recognize handwritten digital numbers with different numbers of training samples on the MNIST dataset. Their results illustrated that, when reducing training samples, the recognition rate of the algorithm drops sharply. To demonstrate the behaviour of our FD3D CNN, we experiment with different percentages of training samples. We use *improvement* defined as  $(1 - ER_D / ER_T) \times 100$  to benchmark our method against a traditional 3D CNN, where  $ER_D$  and  $ER_T$  are error rates of our FD3D CNN and the traditional 3D CNN, respectively. Error rates of MBS/MAS representation learning are shown in Fig. 4, where our proposed method appears

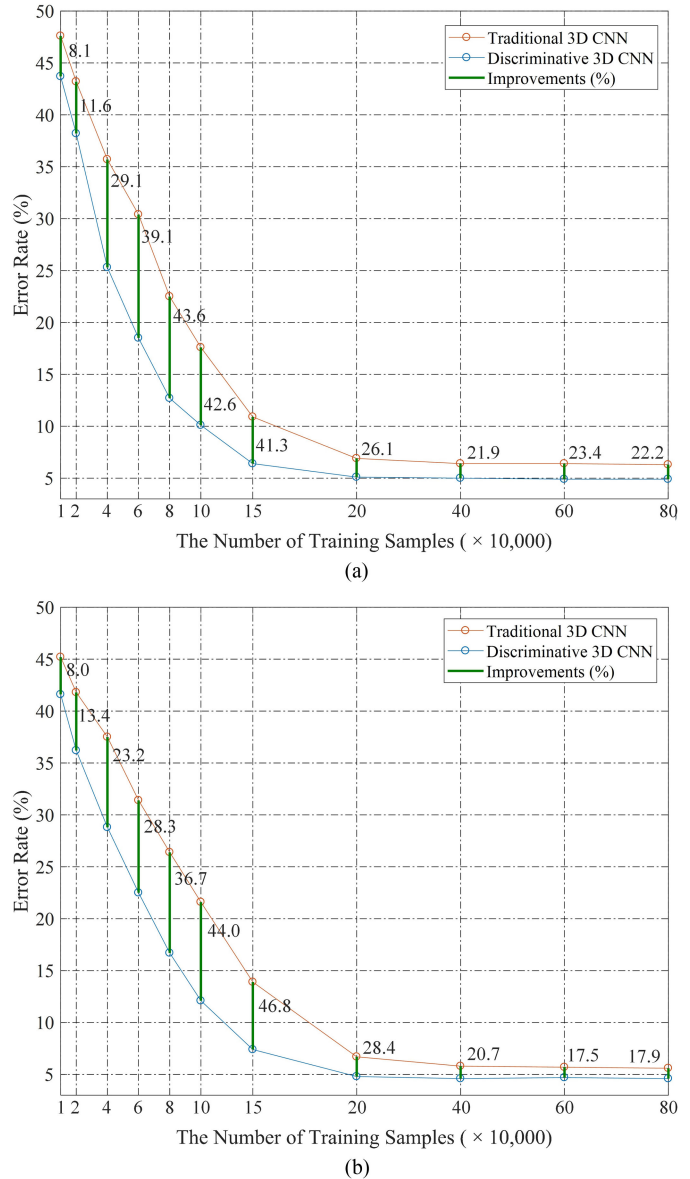


Fig. 4. Error rates and improvements for increasingly larger training sets. (a) MBS detection. (b) MAS detection.

to achieve comparable results with less training data compared to the conventional 3D CNN. We choose 80% of the 845,000 as the training samples and perform testing on the remaining 20%. The results are shown in Table II. Even when trained with fewer iterations, our method achieves better results than the traditional 3D CNN.

With sufficient training samples and iterations, most machine learning methods can improve their accuracy at a higher computational cost. However, we usually want to obtain a trained network as quickly as possible. This is especially important in population imaging as new datasets can become available and retraining might be required. Rapid training is also a desirable feature during algorithmic development since finding an optimal architecture may require multiple training procedures for different parameter settings. We illustrate that our FD3D CNN has better error-reducing performance as a function of the number

**TABLE II**  
ERROR RATES VERSUS LEARNING EPOCHS

Epochs	Error Rate (%)		Improvement (%) (MBS/MAS)
	Traditional 3D CNN (MBS/MAS)	Discriminative 3D CNN (MBS/MAS)	
1	32.4/30.7	28.8/27.4	11.1/10.8
10	25.4/24.2	19.2/17.6	24.4/27.3
20	19.2/18.7	11.3/10.8	41.1/42.2
30	12.7/13.1	8.3/8.6	34.6/34.4
40	6.3/5.6	4.9/4.6	22.2/17.9

**TABLE III**  
PERFORMANCE VERSUS BLOCK SIZE

Block Size	Precision		Sensitivity	
	MAS	MBS	MAS	MBS
$80 \times 80 \times 3$	89.01%	88.36%	88.24%	87.94%
$120 \times 120 \times 3$	<b>91.81%</b>	<b>90.73%</b>	<b>90.92%</b>	<b>90.25%</b>
$180 \times 180 \times 3$	90.12%	89.78%	89.63%	88.92%

of training samples and iterations than other competing techniques.

A 3D CNN requires a suitable receptive field (i.e. input size) to achieve the best discrimination. Based on a random sample of 200 image volumes, we determine the smallest crop size that ensures the coverage of the LV structure compared to three block-size configurations, namely,  $120 \times 120 \times 3$  (which removes redundant background information based on the central point of original images),  $180 \times 180 \times 3$  (which is the original size as extracted and resized from the UK Biobank), and  $80 \times 80 \times 3$ , which mostly contains the LV at the centre. We test sizes smaller than the original block size of the classification model because we want to determine whether a larger input block with more contextual information can enhance the model's discriminative capacity. The results obtained with these settings are shown in Table III. With a block size of  $80 \times 80 \times 3$ , MBS/MAS detection precision rate reaches 89.01% and 88.36%, respectively. The detection performance improves to a precision rate of 91.81% and 90.73% under block size  $120 \times 120 \times 3$ , demonstrating that increasing contextual information can enhance the discriminative capacity of 3D CNN. Without cropping, the detection precision rate decreases to 90.12% and 89.78% for MBS and MAS detection, respectively. This may have been because too much redundant contextual information clutters the actual LV signature, and hence degrades detection performance. Based on these experiments, we set block size to  $120 \times 120 \times 3$ , to achieve optimal detection performance.

Typical classification results using the proposed FD3D CNN architecture are shown in Fig. 5. A few basal stacks (top row) and apical stacks (bottom row) in the test datasets with their AQ or corresponding posterior probability values are shown. High score values on the stack correspond to the likelihood of being a correct basal or apical triplet. Basal slices with existing LVOT indicate higher probability values of being correctly classified.

This shows that the training captures the LVOT as a prominent feature in correctly positioned basal slices.

**2) Comparison to Other Machine Learning Methods:** We compare our framework with a traditional 3D CNN and with our previous 2D CNN study [31]. Table IV lists the results for these architectures. The architecture of a traditional 3D CNN is similar to that of our FD3D CNN, replacing the fisher layer (F2) with a traditional fully connected layer including 256 ReLU activation neurons. We use the same training and testing approaches for the 3D CNN and list the results obtained using the hand crafted features used in [51]. In [51], the basal slice was identified following these steps: 1. Choose the mid-slice image as the start image and process each image sequentially in the basal direction. 2. Apply the optimal threshold method to convert the ROI to a binary image. 3. Identify the binary object with blood pool, which shows an elliptical shape. 4. Calculate the length of the major axis  $L$  of the ellipse that has the same normalized second central moments as the binary object. 5. If the ratio of the current to preceding  $L$  exceeds a predefined threshold (e.g.  $> 1.2$  in this work), then a basal slice is identified; otherwise, the basal slice is missing. We use a similar method to identify the apical slice. We process each image sequentially from base to apex. If the ratio of the current to preceding  $L$  is smaller than a predefined threshold (e.g.  $< 0.2$  in this study), an apical slice is detected; otherwise, the apical slice is missing. We employ this feature extraction procedure for prediction. The proposed FD3D CNN shows the best precision and sensitivity figures in each representation classifier, and full LV coverage detection performance.

### B. Inter-Observer Reliability

To contextualize the results of automatic full LV coverage assessment, we compare it to the inter-observer full LV coverage detection rate obtained by expert readers. The inter-observer agreement [52] of human experts is evaluated by reassessing a subset of 200 random CMR datasets. The quality distribution levels in this randomly selected subset are compared to original data using Pearson's  $\chi^2$  goodness-of-fit test to confirm that it represents the original data distribution ( $p > 0.05$ ). The reassessed samples demonstrate strong agreement with original qualities (Cohen's  $\kappa = 0.76$ ,  $p < 0.05$ ).

To show how our results can be compared to the expected human detection error rates, we present the error rates between an expert cardiologist (VQ1) and another cardiac image expert (VQ2) for 200 re assessed samples. The confusion matrix of VQ1 versus VQ2 is presented in Table V. Use of the confusion matrix reveals 7 among the 200 re assessed samples with inconsistent quality assessment between VQ1 and VQ2. These findings show that the expert cardiologist's visual results conflict with the cardiac image expert's visual assessment only 3% of the time. As shown in Table II  $epoch = 40$ , our automatic algorithm's error rate is just below 5%, which shows excellent agreement with human expert assessments (two percentage points). Some examples of MBS/MAS test images are shown in Fig. 5 (panels a and b correspondingly). We have intentionally



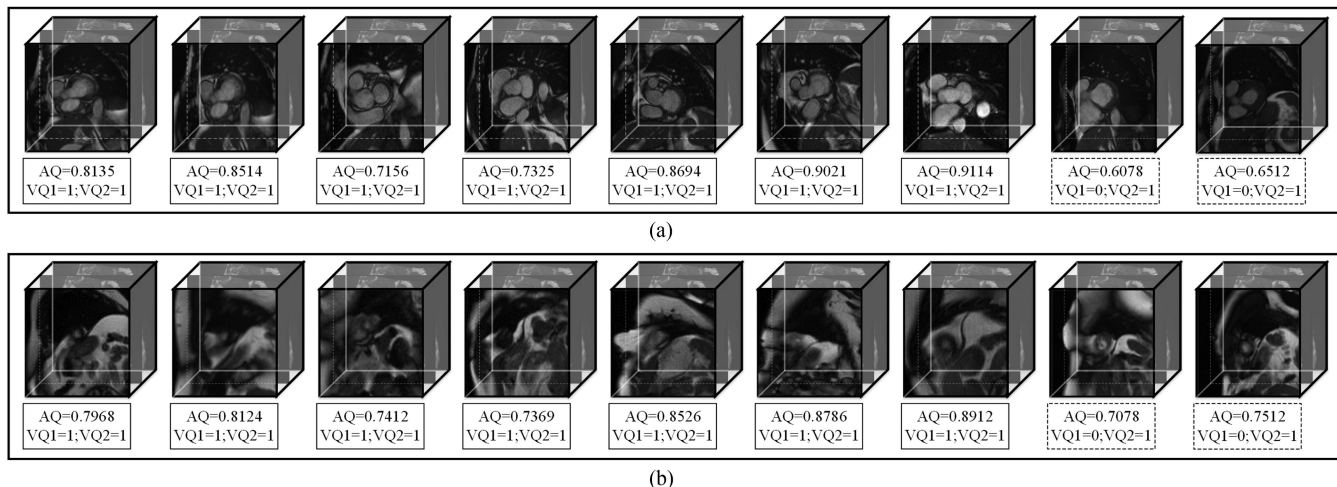


Fig. 5. Sample test volumes and their AQ, expert cardiologist (VQ1) and cardiac image expert’s visual (VQ2) qualities for MBS detection (top row) or MAS detection (bottom row) are shown. The left seven samples in each row show consistency between AQ and VQ1, which means our algorithm yields an accurate prediction. The right two samples in each row show the wrong quality prediction and show inconsistency between VQ1 and VQ2. (a) Sample volumes for MBS testing with automatic quality (AQ), expert cardiologist (VQ1) and cardiac image expert’s visual (VQ2) qualities. (b) Sample volumes for MAS testing with automatic quality (AQ), expert cardiologist (VQ1) and cardiac image expert’s visual (VQ2) qualities.

TABLE IV  
PERFORMANCE COMPARISON OF DIFFERENT LEARNING MODELS WITH LEARNED AND HAND-CRAFTED VISUAL REPRESENTATIONS

Method	Features	Precision (%)			Sensitivity (%)		
		MAS	MBS	MBS ∨ MAS	MAS	MBS	MBS ∨ MAS
FD3D CNN	Learned	<b>91.81 ± 0.21</b>	<b>90.73 ± 0.28</b>	<b>91.12 ± 0.24</b>	<b>90.92 ± 0.26</b>	<b>90.25 ± 0.28</b>	<b>90.15 ± 0.22</b>
3D CNN		89.12 ± 0.36	89.32 ± 0.34	89.20 ± 0.31	89.42 ± 0.31	89.47 ± 0.30	89.25 ± 0.29
2D CNN		81.61 ± 0.56	74.10 ± 0.58	79.42 ± 0.62	88.73 ± 0.49	88.75 ± 0.51	88.01 ± 0.56
Lu <i>et al.</i> [51]	Hand-crafted	37.60 ± 1.22	45.68 ± 1.36	56.92 ± 1.71	67.43 ± 0.92	74.56 ± 1.32	63.25 ± 1.79

TABLE V  
CONFUSION MATRIX OF THE EXPERT CARDIOLOGIST (VQ1) AND CARDIAC IMAGE EXPERT’S VISUAL (VQ2) RESULTS. GREY NUMBERS INDICATE NUMBER AND RATIO OF CORRECT ESTIMATES

		VQ2			Correct
		MBS	MAS	MBS ∨ MAS	
VQ1	MBS	67	0	3	0.96
	MAS	0	65	2	0.97
	MBS ∨ MAS	1	1	61	0.97

TABLE VI  
CROSS-DATASET PERFORMANCE: KAGGLE DATASET

	Error (%)	Precision (%)	Sensitivity (%)
MAS	6.43	86.51	88.74
MBS	7.02	84.03	85.69
MBS ∨ MAS	6.64	85.74	87.01

chosen to show seven inter-observer agreement examples, plus two disagreement examples on each panel.

C. Cross-Database Performance: Sunnybrook Cardiac Dataset

We evaluate the generalization of the performance of our full LV coverage detection system on an independent database. We assess the sensitivity of our system to moderate changes in imaging conditions, scanner vendors, image resolution, etc. To this effect, we use Data Science Bowl Cardiac Challenge Data (Kaggle or Sunnybrook Cardiac dataset) [53]. This dataset comprises 1,120 cardiac MRI volumes. Cine steady state free precession (SSFP) MR short-axis (SAX) images are obtained

with a 1.5T GE Signa CV/i MRI System (General Electric, Milwaukee, WI). All images are obtained during 10–15 second breath-holds with a temporal resolution of 20 cardiac phases over the heart cycle (scanned from the ED phase). Six to twelve SAX images are obtained from the atrioventricular ring to the apex (resolution  $1.25 \times 1.25 \times 8 \text{ mm}^3$ , thickness = 8 mm). Gold-standard full LV coverage is obtained by an experienced reader and checked visually by inspecting slices from base to apex. Original volumes are used for full LV coverage detection and triplets of top and bottom slices are used, respectively, as negative examples for MBS and MAS. Positive examples of MBS/MAS are obtained from triples of mid-slices. This dataset is used as a test set for the FD3D CNN that was pre-trained with 800,000 volumes from the UK Biobank. Values for error, precision and sensitivity under various conditions are shown in Table VI.

TABLE VII  
MISSING SLICE RATE PER VISUAL QUALITY SCORE

Quality Score	MAS (%)	MBS (%)	MBS $\vee$ MAS (%)
1	1.7	0.6	97.7
2	74.7	24.0	1.3
3	18.0	80.4	1.6

TABLE VIII  
EFFECT OF INCOMPLETE CARDIAC COVERAGE (MBS/MAS) ON THE END-DIASTOLIC, END-SYSTOLIC, STROKE VOLUMES AND EJECTION FRACTION. VALUES ARE SHOWN AS MEAN  $\pm$  STANDARD DEVIATIONS

	Full	MBS	Effect(%)	MAS	Effect(%)
LVEDV(ml)	155.8 $\pm$ 35.6	136.1 $\pm$ 33.4	-12.6%	151.5 $\pm$ 35.1	-2.7%
LVESV(ml)	66.8 $\pm$ 21.2	53.0 $\pm$ 19.0	-20.0%	64.3 $\pm$ 20.9	-3.7%
LVSVM(ml)	89.1 $\pm$ 19.8	83.1 $\pm$ 19.7	-6.7%	87.1 $\pm$ 19.6	-2.2%
LVEF(%)	57.1 $\pm$ 0.06	61.0 $\pm$ 0.06	+3.9%	57.5 $\pm$ 0.06	+0.4%

#### D. Missing Slice Rate per Visual Quality Score

To gain insight into the relation between missing slice rates and visual quality scores achieved by experts [7], a third experiment is conducted. The system is trained on 3,380 random volumes from a total of 5,065. The testing set, as earlier indicated, has 1,685 CMR volumes distributed among the quality scores (from 1 to 3: 981, 527 and 177). Table VII gives the percentages of the full LV coverage class for each quality score. CMR data with a quality score of 3 highly correlates with MBS, as missing basal slices highly affect accurate quantitative analysis in CMR.

#### E. Clinical Impact

To assess the impact of incomplete LV coverage in real applications, such as measurement of cardiac function based on blood volumes, we design an experiment where incomplete coverage is simulated and volume differences between full and incomplete coverage are measured. We also compute two commonly used indexes of the cardiac function derived from such volumes viz. stroke volume (SV) and ejection fraction (EF), and similarly report the differences between the full and incomplete coverages. For this experiment, we take 4,737 subjects for which manual annotations are available (both cardiac phase labels and full coverage labels), and systematically remove the basal and apical slices to generate incomplete MBS and MAS volumes. Then, we compute blood pool volumes at the ED and ES phases, and from these, we obtain SV and EF. Finally, the average volumes and indexes are computed across the sample, comparing full coverage and MBS/MAS. Table VIII shows that the largest effect of incomplete coverage is caused by MBS, where the missing slice reduces ED and ES volumes by an average of 12% and 20%, respectively. In turn, these differences cause a decrease in the computed SV by 6.7% and an increase in the EF by 3.9%. The absence of the apical slice has a smaller yet non-negligible impact on the volumes and derived indexes.

#### F. Implementation Considerations

The experiments reported here are conducted using the ConVNet library [54] on an Intel Xeon E5-1620 v3 @3.50 GHz machine running Windows 10 with 32 GB RAM and Nvidia Quadro K620 GPU. The networks are optimized using the gradient descent method [42] with the following hyper-parameters: learning rate = 0.01, momentum = 0.9, drop-out rate = 0.1. Trainable weights are randomly initialized from a Gaussian distribution ( $\mu = 0$ ,  $\sigma = 0.01$ ) and updated with standard back-propagation. Models converge in about 6 hours when training is performed with 800,000 volumes with size  $120 \times 120 \times 3$ . Testing is rapid and can process each volume in 3 seconds.

#### V. DISCUSSION

Automatic identification of CMR volumes with incomplete LV coverage is important in high-throughput image analysis of population imaging. The acquisition of thousands of suboptimal CMR images for later image analysis can be avoided if such quality assessment is performed online and a system provides immediate feedback to technical staff when new images are acquired. Incomplete LV coverage influences the accuracy of anatomical and functional LV parameters of clinical interest. Manual annotation of LV coverage is laborious, time-consuming and error prone in current clinical routines. To automate this labour-intensive task, we propose an efficient and robust two-stage framework for the automatic detection of missing slices at the LV base and apex. In the first stage, we train a FD3D CNN that computes the corresponding intensity representation with high accuracy. It can qualify CMR volumes based on two representations, and can assist radiologists by automatically labelling the potentially incomplete volumes to mark them for closer inspection. The second stage robustly discriminates two quality categories (MBS and MAS), based only on the intensity representation classifiers, which are then used to recognize new cardiac volumes with no further training. Specifically, to use the spatial information in volumetric data, we use 3D CNN with shared 3D convolution kernels. Meanwhile, a Fisher discriminant layer leads to small within-class scatter and large between-class scatter of feature vectors in that layer. Extensive experimental results illustrate the effectiveness and efficiency of our method: its performance is superior to that of other methods with obvious advantages.

In any AIQA system for population imaging, accuracy and robustness are key design criteria. These methods must work without many false positives or false negatives, and must cope with considerable variation in image quality. Most machine learning methods can improve their recognition accuracy by increasing the number of iterations. However, an increasing number of iterations comes at a high computational cost. This can be prohibitive with large databases or when retraining is required as new data become available. In this study, we used a very large dataset comprising more than 5,000 individually annotated cardiac MRI scans of the same number of subjects, which is 50-fold the 100 cases used in our previous study [31]. However, when compared to natural image datasets [42], our cardiac

MRI dataset is still relatively small. We had to design an efficient network taking full advantage of the available data. Considering there were only a few labelled images, there was no point in constructing a network with too many sub-sampling layers; there would have been a higher computational cost with more layers of feature abstraction. Three-dimensional CNN have been among the most promising solutions for object detection tasks. Thus far, most studies have focused on image segmentation and registration, and little effort has been devoted to AIQA. We propose a FD3D CNN with an extra layer using a Fisher discriminant criterion, which tackles the problem of detecting full LV coverage as an important quality criterion. Our method can eliminate redundant convolutional computations during forward propagation and achieve a comparable result with a smaller number of training samples and iterations. Specifically, our FD3D CNN can achieve a high precision rate of nearly 92%/91% for MBS/MAS detection with only 20 epochs, which is better than traditional 3D CNN. Meanwhile, even with a small number of training samples ( $4 \times 10,000$ ), our FD3D CNN can decrease the error rate by approximately 29.1% compared to traditional 3D CNN approaches for MBS detection.

Our proposed automatic assessment framework for full LV coverage has great potential to improve the robustness of subsequent population image parsing. One can imagine an approach whereby image analysis is adaptive to image quality and where different models are used depending on whether the volume under analysis is missing basal or apical slices. In our architecture, we focus on learning intensity representations and develop a FD3D CNN to describe those that best discriminate the missing apical or basal slices. We then use the computed representation classifiers to identify the final image quality. The advantages of a representation-based method for vision tasks are manifold: they can be composed to create descriptions at various levels of specificity; they are generalizable, as they can be learned once and then applied to recognize new objects or categories with no further training and are efficient, possibly requiring exponentially fewer representations than explicitly naming each category. In the future, we plan to investigate the possibility of detecting full LV coverage for all slices, rather than just for basal/apical slices, so we can directly predict visual quality scores. The difficulty of detecting missing middle slices lies in the similar shape of contiguous LV slices, which makes training the representation classifier a non-trivial task. Another future work is to extend deep-learning methods for multi-plane estimation, that is, regressing one 3D volume to estimating missing slices acquired from different positions. This is a limitation of our two-stage framework, which can only estimate the basal and apical planes. One way to achieve 3D CNN for multi-plane estimation would be to apply regression on each plane separately and then combine all regression results into a single estimation.

## VI. CONCLUSION

In this study, we tackled the problem of detecting incomplete LV coverage in large population image databases. We illustrated the concept by proposing a Fisher discriminative 3D CNN tested on CMR data from the UK Biobank. Our FD3D

CNN was proposed by adding a new Fisher-discriminative fully connected layer into the network, which achieved a significant improvement in intensity representation. The learned representation classifiers were computed for candidates of corresponding quality categories. We also validated our model by training with the UK Biobank dataset and cross-evaluating with data from the Data Science Bowl Cardiac Challenge dataset. The proposed model shows high consistency with human perception and is superior to state-of-the-art methods, showing its high potential. Our proposed FD3D CNN can also be easily applied and boosts results for other detection and segmentation tasks in medical image analysis.

## ACKNOWLEDGMENT AND COMPETING INTERESTS

This research has been conducted using the UK Biobank Resource under Application 2964. The authors wish to thank all UK Biobank participants and staff. S. E. Petersen provides consultancy to Circle Cardiovascular Imaging Inc., Calgary, AB, Canada, and the access application 11350 of UK Biobank.

## REFERENCES

- [1] S. E. Petersen *et al.*, "Imaging in population science: Cardiovascular magnetic resonance in 100,000 participants of UK Biobank—Rationale, challenges and approaches," *J. Cardiovascular Magn. Reson.*, vol. 15, pp. 46–56, 2013.
- [2] V. Klinke *et al.*, "Quality assessment of cardiovascular magnetic resonance in the setting of the European CMR registry: Description and validation of standardized criteria," *J. Cardiovascular Magn. Reson.*, vol. 15, pp. 55–67, 2013.
- [3] E. Pusey *et al.*, "Magnetic resonance imaging artifacts: Mechanism and clinical significance," *Radiographics*, vol. 6, pp. 891–911, 1986.
- [4] P. Ferreira *et al.*, "Cardiovascular magnetic resonance artefacts," *J. Cardiovascular Magn. Reson.*, vol. 15, pp. 41–79, 2013.
- [5] Z. Wang *et al.*, "Quality-aware images," *IEEE Trans. Image Process.*, vol. 15, no. 6, pp. 1680–1689, Jun. 2006.
- [6] A. Attili *et al.*, "Quantification in cardiac MRI: Advances in image acquisition and processing," *Int. J. Cardiovascular Imag.*, vol. 26, pp. 27–40, 2010.
- [7] V. Carapella *et al.*, "Towards the semantic enrichment of free-text annotation of image quality assessment for UK Biobank cardiac cine MRI scans," in *Proc. Deep Learn. Data Labeling Med. Appl.*, 2016, pp. 238–248.
- [8] M. Saad *et al.*, "Blind image quality assessment: A natural scene statistics approach in the DCT domain," *IEEE Trans. Image Process.*, vol. 21, no. 8, pp. 3339–3352, Aug. 2012.
- [9] W. Xue *et al.*, "Blind image quality assessment using joint statistics of gradient magnitude and Laplacian features," *IEEE Trans. Image Process.*, vol. 23, no. 11, pp. 4850–4862, Nov. 2014.
- [10] L. He *et al.*, "Sparse representation for blind image quality assessment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012, pp. 1146–1153.
- [11] A. Moorthy and A. Bovik, "Blind image quality assessment: From natural scene statistics to perceptual quality," *IEEE Trans. Image Process.*, vol. 20, no. 12, pp. 3350–3364, Dec. 2011.
- [12] L. Kang *et al.*, "Convolutional neural networks for no-reference image quality assessment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1733–1740.
- [13] O. Bernard *et al.*, "Standardized evaluation system for left ventricular segmentation algorithms in 3D echocardiography," *IEEE Trans. Med. Imag.*, vol. 35, no. 4, pp. 967–977, Apr. 2016.
- [14] X. Zhen *et al.*, "Multi-scale deep networks and regression forests for direct bi-ventricular volume estimation," *Med. Image Anal.*, vol. 30, pp. 120–129, 2016.
- [15] R. Hoffmann *et al.*, "Automated heart localization in cardiac cine MR data," in *Proc. Bildverarbeitung für die Medizin*, 2016, pp. 116–121.
- [16] X. Lu *et al.*, "Robust 3D organ localization with dual learning architectures and fusion," in *Proc. Int. Workshop Large-Scale Annotation Biomed. Data Expert Label Synthesis.*, 2016, pp. 12–20.



- [17] B. D. de Vos *et al.*, "Convnet-based localization of anatomical structures in 3D medical images," *IEEE Trans. Med. Imag.*, vol. 36, no. 7, pp. 1470–1481, Jul. 2017.
- [18] X. Lu *et al.*, "Automatic view planning for cardiac MRI acquisition," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention.*, 2011, pp. 479–486.
- [19] B. D. de Vos *et al.*, "Automatic slice identification in 3D medical images with a ConvNet regressor," in *Proc. Int. Workshop Large-Scale Annotation Biomed. Data Expert Label Synthesis.*, 2016, pp. 161–169.
- [20] D. Fouhey *et al.*, "3D Shape Attributes," in *Proc. IEEE Conf. Comp. Vis. Pattern Recognit.*, 2016, pp. 1516–1524.
- [21] J. Xie *et al.*, "Deepshape: Deep-learned shape descriptor for 3D shape retrieval," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 7, pp. 1335–1345, Jul. 2017.
- [22] D. Tran *et al.*, "Learning spatiotemporal features with 3D convolutional networks," in *Proc. IEEE Conf. Comp. Vis.*, 2015, pp. 4489–4497.
- [23] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, 2017.
- [24] H. Roth *et al.*, "Improving computer-aided detection using convolutional neural networks and random view aggregation," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1170–1181, May 2016.
- [25] A. Prasoon *et al.*, "Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2013, pp. 246–253.
- [26] K. Kamnitsas *et al.*, "Multi-scale 3D convolutional neural networks for lesion segmentation in brain MRI," in *Proc. Ischemic Stroke Lesion Segmentation*, 2015, vol. 13, pp. 46–49.
- [27] G. Urban *et al.*, "Multi-modal brain tumor segmentation using deep convolutional neural networks," in *Proc. MICCAI Brain Tumor Segmentation Challenge*, 2014, pp. 31–35.
- [28] Q. Dou *et al.*, "Automatic detection of cerebral micro bleeds from MR images via 3D convolutional neural networks," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1182–1195, May 2016.
- [29] K. Kamnitsas *et al.*, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Med. Image Anal.*, vol. 36, pp. 61–78, 2017.
- [30] Ö. Çiçek *et al.*, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2016, pp. 424–432.
- [31] L. Zhang *et al.*, "Automated quality assessment of cardiac MR images using convolutional neural networks," in *Proc. Int. Workshop Simul. Synthesis Med. Imag.*, 2016, pp. 138–145.
- [32] G. Cheng *et al.*, "RIFD-CNN: Rotation-invariant and Fisher discriminative convolutional neural networks for object detection," in *Proc. IEEE Conf. Comp. Vis. Pattern Recognit.*, 2016, pp. 2884–2893.
- [33] M. Yang *et al.*, "Sparse representation based Fisher discrimination dictionary learning for image classification," *Int. J. Comput. Vis.*, vol. 109, pp. 209–232, 2014.
- [34] J. Schulz-Menger *et al.*, "Standardized image interpretation and post processing in cardiovascular magnetic resonance: Society for Cardiovascular Magnetic Resonance (SCMR) board of trustees task force on standardized post processing," *J. Cardiovascular Magn. Reson.*, vol. 15, pp. 35–53, 2013.
- [35] S. Petersen *et al.*, "Reference ranges for cardiac structure and function using cardiovascular magnetic resonance (CMR) in Caucasians from the UK Biobank population cohort," *J. Cardiovascular Magn. Reson.*, vol. 19, pp. 18–36, 2017.
- [36] J. Marcus *et al.*, "The influence of through-plane motion on left ventricular volumes measured by magnetic resonance imaging: Implications for image acquisition and analysis," *J. Cardiovascular Magn. Reson.*, vol. 1, pp. 1–6, 1999.
- [37] H. Childs *et al.*, "Comparison of long and short axis quantification of left ventricular volume parameters by cardiovascular magnetic resonance, with ex-vivo validation," *J. Cardiovascular Magn. Reson.*, vol. 13, pp. 40–48, 2011.
- [38] C. Miller *et al.*, "Quantification of left ventricular indices from SSFP cine imaging: Impact of real-world variability in analysis methodology and utility of geometric modeling," *J. Magn. Reson. Imag.*, vol. 37, pp. 1213–1222, 2013.
- [39] Y. Lu *et al.*, "Pattern recognition of abnormal left ventricle wall motion in cardiac MR," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2009, pp. 750–758.
- [40] S. Ji *et al.*, "3D convolutional neural networks for human action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 221–231, Jan. 2013.
- [41] F. Huang *et al.*, "Unsupervised learning of invariant feature hierarchies with applications to object recognition," in *Proc. IEEE Conf. Comp. Vis. Pattern Recognit.*, 2007, pp. 1–8.
- [42] A. Krizhevsky *et al.*, "ImageNet classification with deep convolutional neural networks," *Adv. Neural Inf. Proc. Syst.*, vol. 25, pp. 1097–1105, 2012.
- [43] C. Ionescu *et al.*, "Matrix backpropagation for deep networks with structured layers," in *Proc. IEEE Conf. Comput. Vis.*, 2015, pp. 2965–2973.
- [44] Y. LeCun *et al.*, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015.
- [45] S. E. Petersen *et al.*, "UK Biobank's cardiovascular magnetic resonance protocol," *J. Cardiovascular Magn. Reson.*, vol. 18, no. 1, pp. 8–14, 2015.
- [46] A. Hoogi *et al.*, "Adaptive estimation of active contour parameters using convolutional neural networks and texture analysis," *IEEE Trans. Med. Imag.*, vol. 36, no. 3, pp. 781–791, Mar. 2017.
- [47] A. Mortazi *et al.*, "Multi-planar deep segmentation networks for cardiac substructures from MRI and CT," in *Int. Workshop Statist. Atlases and Comput. Models Heart*, Springer, pp. 199–206, 2017.
- [48] N. Kumar *et al.*, "Describable visual attributes for face verification and image search," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 10, pp. 1962–1977, Oct. 2011.
- [49] P. Sermanet *et al.*, "Convolutional neural networks applied to house numbers digit classification," in *Proc. IEEE 21st Int. Conf. Pattern Recognit.*, 2012, pp. 3288–3291.
- [50] A. Salah *et al.*, "A selective attention-based method for visual pattern recognition with application to handwritten digit recognition and face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 3, pp. 420–425, Mar. 2002.
- [51] Y. Lu *et al.*, "Watershed segmentation of basal left ventricle for quantitation of cine cardiac MRI function," *J. Cardiovascular Magn. Reson.*, vol. 13, pp. 1–2, 2011.
- [52] K. L. Gwet, "Intrarater reliability," in *Wiley Encyclopedia Clinical Trials*. Hoboken, NJ, USA: Wiley, 2008, pp. 1–14.
- [53] Kaggle, "Data science bowl cardiac challenge data," 2016. [Online]. Available: [www.kaggle.com/c/second-annual-data-science-bowl/data](http://www.kaggle.com/c/second-annual-data-science-bowl/data). Accessed on: Oct. 15, 2017.
- [54] S. Demyanov, "ConvNet library for Matlab." 2014. [Online]. Available: <https://github.com/sdemyanov/ConvNet>. Accessed on: Oct. 15, 2017.