OPEN

Automatic segmentation with detection of local segmentation failures in cardiac MRI

Jörg Sander^{1,2,3✉}, Bob D. de Vos^{1,2} & Ivana Išgum^{1,2,3,4}

Segmentation of cardiac anatomical structures in cardiac magnetic resonance images (CMRI) is a prerequisite for automatic diagnosis and prognosis of cardiovascular diseases. To increase robustness and performance of segmentation methods this study combines automatic segmentation and assessment of segmentation uncertainty in CMRI to detect image regions containing local segmentation failures. Three existing state-of-the-art convolutional neural networks (CNN) were trained to automatically segment cardiac anatomical structures and obtain two measures of predictive uncertainty: entropy and a measure derived by MC-dropout. Thereafter, using the uncertainties another CNN was trained to detect local segmentation failures that potentially need correction by an expert. Finally, manual correction of the detected regions was simulated in the complete set of scans of 100 patients and manually performed in a random subset of scans of 50 patients. Using publicly available CMR scans from the MICCAI 2017 ACDC challenge, the impact of CNN architecture and loss function for segmentation, and the uncertainty measure was investigated. Performance was evaluated using the Dice coefficient, 3D Hausdorff distance and clinical metrics between manual and (corrected) automatic segmentation. The experiments reveal that combining automatic segmentation with manual correction of detected segmentation failures results in improved segmentation and to 10-fold reduction of expert time compared to manual expert segmentation.

To perform diagnosis and prognosis of cardiovascular disease (CVD) medical experts depend on the reliable quantification of cardiac function¹. Cardiac magnetic resonance imaging (CMRI) is currently considered the reference standard for quantification of ventricular volumes, mass and function². Short-axis CMR imaging, covering the entire left and right ventricle (LV resp. RV) is routinely used to determine quantitative parameters of both ventricle's function. This requires manual or semi-automatic segmentation of corresponding cardiac tissue structures for end-diastole (ED) and end-systole (ES).

Existing semi-automated or automated segmentation methods for CMRIs regularly require (substantial) manual intervention caused by lack of robustness. Manual or semi-automatic segmentation across a complete cardiac cycle, comprising 20 to 40 phases per patient, enables computation of parameters quantifying cardiac motion with potential diagnostic implications but due to the required workload, this is practically infeasible. Consequently, segmentation is often performed at end-diastole and end-systole precluding comprehensive analysis over complete cardiac cycle.

Recently^{3,4}, deep learning segmentation methods have shown to outperform traditional approaches such as those exploiting level set, graph-cuts, deformable models, cardiac atlases and statistical models^{5,6}. However, recent comparison of a number of automatic methods showed that even the best performing methods generated anatomically implausible segmentations in more than 80% of the CMRIs⁷. Such errors do not occur when experts perform segmentation. To achieve acceptance in clinical practice these shortcomings of the automatic approaches need to be alleviated by further development. This can be achieved by generating more accurate segmentation result or by development of approaches that automatically detect segmentation failures.

¹Department of Biomedical Engineering and Physics, Amsterdam University Medical Centers - location AMC, University of Amsterdam, Amsterdam, The Netherlands. ²Amsterdam Cardiovascular Sciences, Amsterdam University Medical Centers - location AMC, University of Amsterdam, Amsterdam, The Netherlands. ³Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands. ⁴Department of Radiology and Nuclear Medicine, Amsterdam University Medical Centers - location AMC, University of Amsterdam, Amsterdam, The Netherlands. ✉email: j.sander1@amsterdamumc.nl

In manual and automatic segmentation of short-axis CMRI, largest segmentation inaccuracies are typically located in the most basal and apical slices due to low tissue contrast ratios⁸. To increase segmentation performance, several methods have been proposed^{9–12}. Tan et al.⁹ used a convolutional neural network (CNN) to regress anatomical landmarks from long-axis views (orthogonal to short-axis). They exploited the landmarks to determine most basal and apical slices in short-axis views and thereby constraining the automatic segmentation of CMRIs. This resulted in increased robustness and performance. Other approaches leverage spatial¹⁰ or temporal^{11,12} information to increase segmentation consistency and performance in particular in the difficult basal and apical slices.

An alternative approach to preventing implausible segmentation results is by incorporating knowledge about the highly constrained shape of the heart. Oktay et al.¹³ developed an anatomically constrained neural network (NN) that infers shape constraints using an auto-encoder during segmentation training. Duan et al.¹⁴ developed a deep learning segmentation approach for CMRIs that used atlas propagation to explicitly impose a shape refinement. This was especially beneficial in the presence of image acquisition artifacts. Recently, Painchaud et al.¹⁵ developed a post-processing approach to detect and transform anatomically implausible cardiac segmentations into valid ones by defining cardiac anatomical metrics. Applying their approach to various state-of-the-art segmentation methods the authors showed that the proposed method provides strong anatomical guarantees without hampering segmentation accuracy.

A different research trend focuses on detecting segmentation failures, i.e. on automated quality control for image segmentation. These methods can be divided in those that predict segmentation quality using image at hand or corresponding automatic segmentation result, and those that assess and exploit predictive uncertainties to detect segmentation failure.

Recently, two methods were proposed to detect segmentation failures in large-scale cardiac MR imaging studies to remove these from subsequent analysis^{16,17}. Robinson et al.¹⁷ using the approach of Reverse Classification Accuracy (RCA)¹⁸ predicted CMRI segmentation metrics to detect failed segmentations. They achieved good agreement between predicted metrics and visual quality control scores. Alba et al.¹⁶ used statistical, pattern and fractal descriptors in a random forest classifier to directly detect segmentation contour failures without intermediate regression of segmentation accuracy metrics.

Methods for automatic quality control were also developed for other applications in medical image analysis. Frounchi et al.¹⁹ extracted features from the segmentation results of the left ventricle in CT scans. Using the obtained features the authors trained a classifier that is able to discriminate between consistent and inconsistent segmentations. To distinguish between acceptable and non-acceptable segmentations Kohlberger et al.²⁰ proposed to directly predict multi-organ segmentation accuracy in CT scans using a set of features extracted from the image and corresponding segmentation.

A number of methods aggregate voxel-wise uncertainties into an overall score to identify insufficiently accurate segmentations. For example, Nair et al.²¹ computed an overall score for target segmentation structure from voxel-wise predictive uncertainties. The method was tested for detection of Multiple Sclerosis in brain MRI. The authors showed that rejecting segmentations with high uncertainty scores led to increased detection accuracy indicating that correct segmentations contain lower uncertainties than incorrect ones. Similarly, to assess segmentation quality of brain MRIs Jungo et al.²² aggregated voxel-wise uncertainties into a score per target structure and showed that the computed uncertainty score enabled identification of erroneous segmentations.

Unlike approaches evaluating segmentation directly, several methods use predictive uncertainties to predict segmentation metrics and thereby evaluate segmentation performance^{23,24}. For example, Roy et al.²³ aggregated voxel-wise uncertainties into four scores per segmented structure in brain MRI. The authors showed that computed scores can be used to predict the Intersection over Union and hence, to determine segmentation accuracy. Similar idea was presented by DeVries et al.²⁴ that predicted segmentation accuracy per patient using an auxiliary neural network that leverages the dermoscopic image, automatic segmentation result and obtained uncertainties. The researchers showed that a predicted segmentation accuracy is useful for quality control.

We build on our preliminary work where automatic segmentation of CMR images using a dilated CNN was combined with assessment of two measures of segmentation uncertainties²⁵. For the first measure the multi-class entropy per voxel (entropy maps) was computed using the output distribution. For the second measure Bayesian uncertainty maps were acquired using Monte Carlo dropout (MC-dropout)²⁶. In²⁵ we showed that the obtained uncertainties almost entirely cover the regions of incorrect segmentation i.e. that uncertainties are calibrated. In the current work we extend our preliminary research in two ways. First, we assess impact of CNN architecture on the segmentation performance and calibration of uncertainty maps by evaluating three existing state-of-the-art CNNs. Second, we employ an auxiliary CNN (detection network) that processes a cardiac MRI and corresponding spatial uncertainty map (Entropy or Bayesian) to automatically detect segmentation failures. We differentiate errors that may be within the range of inter-observer variability and hence do not necessarily require correction (tolerated errors) from the errors that an expert would not make and hence require correction (segmentation failures). Given that overlap measures do not capture fine details of the segmentation results and preclude us to differentiate two types of segmentation errors, in this work, we define segmentation failure using a metric of boundary distance. In²⁵ we found that degree of calibration of uncertainty maps is dependent on the loss function used to train the CNN. Nevertheless, in the current work we show that uncalibrated uncertainty maps are useful to detect local segmentation failures. In contrast to previous methods that detect segmentation failure per-patient or per-structure^{23,24}, we propose to detect segmentation failures per image region. We expect that inspection and correction of segmentation failures using image regions rather than individual voxels or images would simplify correction process. To show the potential of our approach and demonstrate that combining automatic segmentation with manual correction of the detected segmentation failures per region results in higher segmentation performance we performed two additional experiments. In the first experiment, correction of detected segmentation failures was simulated in the complete data set. In the second experiment, correction

Figure 1. Overview of proposed two step approach. Step 1 (left): Automatic CNN segmentation of CMR images combined with assessment of segmentation uncertainties. Step 2 (right): Differentiate tolerated errors from segmentation failures (to be detected) using distance transform maps based on reference segmentations. Detection of image regions containing segmentation failures using CNN which takes CMR images and segmentation uncertainties as input. Manual corrected segmentation failures (green) based on detected image regions.

was performed by an expert in a subset of images. Using publicly available set of CMR scans from MICCAI 2017 ACDC challenge⁷, the performance was evaluated before and after simulating the correction of detected segmentation failures as well as after manual expert correction.

Data

In this study data from the MICCAI 2017 Automated Cardiac Diagnosis Challenge (ACDC)⁷ was used. The dataset consists of cardiac cine MR images (CMRIs) from 100 patients uniformly distributed over normal cardiac function and four disease groups: dilated cardiomyopathy, hypertrophic cardiomyopathy, heart failure with infarction, and right ventricular abnormality. Detailed acquisition protocol is described by Bernard et al.⁷. Briefly, short-axis CMRIs were acquired with two MRI scanners of different magnetic strengths (1.5 and 3.0 T). Images were made during breath hold using a conventional steady-state free precession (SSFP) sequence. CMRIs have an in-plane resolution ranging from 1.37 to 1.68 mm (average reconstruction matrix 243×217 voxels) with slice spacing varying from 5 to 10 mm. Per patient 28 to 40 volumes are provided covering partially or completely one cardiac cycle. Each volume consists of on average ten slices covering the heart. Expert manual reference segmentations are provided for the LV cavity, RV endocardium and LV myocardium (LVM) for all CMRI slices at ED and ES time frames. To correct for intensity differences among scans, voxel intensities of each volume were scaled to the $[0.0, 1.0]$ range using the minimum and maximum of the volume. Furthermore, to correct for differences in-plane voxel sizes, image slices were resampled to $1.4 \times 1.4 \text{ mm}^2$.

Methods

To investigate uncertainty of the segmentation, anatomical structures in CMR images are segmented using a CNN. To investigate whether the approach generalizes to different segmentation networks, three state-of-the-art CNNs were evaluated. For each segmentation model two measures of predictive uncertainty were obtained per voxel. Thereafter, to detect and correct local segmentation failures an auxiliary CNN (detection network) that analyzes a cardiac MRI was used. Finally, this leads to the uncertainty map allowing detection of image regions that contain segmentation failures. Fig. 1 visualizes this approach.

Automatic segmentation of cardiac MRI. To perform segmentation of LV, RV, and LVM in cardiac MR images i.e. 2D CMR scans, three state-of-the-art CNNs are trained. Each of the three networks takes a CMR image as input and has four output channels providing probabilities for the three cardiac structures (LV, RV, LVM) and background. Softmax probabilities are calculated over the four tissue classes. Patient volumes at ED and ES are processed separately. During inference the 2D automatic segmentation masks are stacked into a 3D volume per patient and cardiac phase. After segmentation, the largest 3D connected component for each class is retained and volumes are resampled to their original voxel resolution. Segmentation networks differ substantially regarding architecture, number of parameters and receptive field size. To assess predictive uncertainties from the segmentation models *Monte Carlo dropout* (MC-dropout) introduced by Gal and Ghahramani²⁶ is implemented in every network. The following three segmentation networks were evaluated: Bayesian Dilated CNN, Bayesian Dilated Residual Network, Bayesian U-net.

Bayesian dilated CNN (DN). The Bayesian DN architecture comprises a sequence of ten convolutional layers. Layers 1 to 8 serve as feature extraction layers with small convolution kernels of size 3×3 voxels. No padding is applied after convolutions. The number of kernels increases from 32 in the first eight layers, to 128 in the final two fully connected classification layers, implemented as 1×1 convolutions. The dilation level is successively increased between layers 2 and 7 from 2 to 32 which results in a receptive field for each voxel of 131×131 voxels, or 18.3×18.3 cm². All trainable layers except the final layer use rectified linear activation functions (ReLU). To enhance generalization performance, the model uses batch normalization in layers 2 to 9. In order to convert the original DN²⁷ into a Bayesian DN dropout is added as the last operation in all but the final layer and 10 percent of a layer's hidden units are randomly switched off.

Bayesian dilated residual network (DRN). The Bayesian DRN is based on the original DRN from Yu et al.²⁸ for image segmentation. More specifically, the DRN-D-22²⁸ is used which consists of a feature extraction module with output stride eight followed by a classifier implemented as fully convolutional layer with 1×1 convolutions. Output of the classifier is upsampled to full resolution using bilinear interpolation. The convolutional feature extraction module comprises eight levels where the number of kernels increases from 16 in the first level, to 512 in the two final levels. The first convolutional layer in level 1 uses 16 kernels of size 7×7 voxels and zero-padding of size 3. The remaining trainable layers use small 3×3 voxel kernels and zero-padding of size 1. Level 2 to 4 use a strided convolution of size 2. To further increase the receptive field convolutional layers in level 5, 6 and 7 use a dilation factor of 2, 4 and 2, respectively. Furthermore, levels 3 to 6 consist of two residual blocks. All convolutional layers of the feature extraction module are followed by batch normalization, ReLU function and dropout. Adding dropout and switching off 10 percent of a layer's hidden units converts the original DRN²⁸ into a Bayesian DRN.

Bayesian U-net (U-net). The standard architecture of the U-net²⁹ is used. The network is fully convolutional and consists of a contracting, bottleneck and expanding path. The contracting and expanding path each consist of four blocks i.e. resolution levels which are connected by skip connections. The first block of the contracting path contains two convolutional layers using a kernel size of 3×3 voxels and zero-padding of size 1. Downsampling of the input is accomplished by employing a max pooling operation in block 2 to 4 of the contracting path and the bottleneck using a convolutional kernel of size 2×2 voxels and stride 2. Upsampling is performed by a transposed convolutional layer in block 1 to 4 of the expanding path using the same kernel size and stride as the max pooling layers. Each downsampling and upsampling layer is followed by two convolutional layers using 3×3 voxel kernels with zero-padding size 1. The final convolutional layer of the network acts as a classifier and uses 1×1 convolutions to reduce the number of output channels to the number of segmentation classes. The number of kernels increases from 64 in the first block of the contracting path to 1024 in the bottleneck. In contrast, the number of kernels in the expanding path successively decreases from 1024 to 64. In deviation to the standard U-net instance normalization is added to all convolutional layers in the contracting path and ReLU non-linearities are replaced by LeakyReLU functions because this was found to slightly improve segmentation performance. In addition, to convert the deterministic model into a Bayesian neural network dropout is added as the last operation in each block of the contracting and expanding path and 10 percent of a layer's hidden units are randomly switched off.

Assessment of predictive uncertainties. To detect failures in segmentation masks generated by CNNs in testing, spatial uncertainty maps of the obtained segmentations are generated. For each voxel in the image two measures of uncertainty are calculated. First, a computationally cheap and straightforward measure of uncertainty is the entropy of softmax probabilities over the four tissue classes which are generated by the segmentation networks. Using these, normalized entropy maps $\mathbf{E} \in [0, 1]^{H \times W}$ (e-map) are computed where H and W denote the height and width of the original CMRI, respectively.

Second, by applying MC-dropout in testing, softmax probabilities with a number of samples T per voxel are obtained. As an overall measure of uncertainty the mean standard deviation of softmax probabilities per voxel over all tissue classes C is computed

$$\mathbf{B}(I)^{(x,y)} = \frac{1}{C} \sum_{c=1}^C \sqrt{\frac{1}{T-1} \sum_{t=1}^T (p_t(I)^{(x,y,c)} - \hat{\mu}^{(x,y,c)})^2}, \quad (1)$$

where $\mathbf{B}(I)^{(x,y)} \in [0, 1]$ denotes the normalized value of the Bayesian uncertainty map (b-map) at position (x, y) in 2D slice I , C is equal to the number of classes, T is the number of samples and $p_t(I)^{(x,y,c)}$ denotes the softmax probability at position (x, y) in image I for class c . The predictive mean per class $\hat{\mu}^{(x,y,c)}$ of the samples is computed as follows:

$$\hat{\mu}^{(x,y,c)} = \frac{1}{T} \sum_{t=1}^T p_t(I)^{(x,y,c)}. \quad (2)$$

In addition, the predictive mean per class is used to determine the tissue class per voxel.

Calibration of uncertainty maps. Ideally, incorrectly segmented voxels as defined by the reference labels should be covered by higher uncertainties than correctly segmented voxels. In such a case the spatial uncertainty maps are perfectly calibrated. *Risk-coverage curves* introduced by Geifman et al.³⁰ visualize whether incorrectly

segmented voxels are covered by higher uncertainties than those that are correctly segmented. Risk-coverage curves convey the effect of avoiding segmentation of voxels above a specific uncertainty value on the reduction of segmentation errors (i.e. risk reduction) while at the same time quantifying the voxels that were omitted from the classification task (i.e. coverage).

To generate risk-coverage curves first, each patient volume is cropped based on a minimal enclosing parallelepiped bounding box that is placed around the reference segmentations to reduce the number of background voxels. Note that this is only performed to simplify the analysis of the risk-coverage curves. Second, voxels of the cropped patient volume are ranked based on their uncertainty value in descending order. Third, to obtain uncertainty threshold values per patient volume the ranked voxels are partitioned into 100 percentiles based on their uncertainty value. Finally, per patient volume each uncertainty threshold is evaluated by computing a coverage and a risk measure. Coverage is the percentage of voxels in a patient volume at ED or ES that is automatically segmented. Voxels in a patient volume above the threshold are discarded from automatic segmentation and would be referred to an expert. The number of incorrectly segmented voxels per patient volume is used as a measure of risk. Using bilinear interpolation risk measures are computed per patient volume between [0, 100] percent.

Detection of segmentation failures. To detect segmentation failures uncertainty maps are used but direct application of uncertainties is infeasible because many correctly segmented voxels, such as those close to anatomical structure boundaries, have high uncertainty. Hence, an additional patch-based CNN (detection network) is used that takes a cardiac MR image together with the corresponding spatial uncertainty map as input. For each patch of 8×8 voxels the network generates a probability indicating whether it contains segmentation failure. In the following, the terms patch and region are used interchangeably.

The detection network is a shallow Residual Network (S-ResNet)³¹ consisting of a feature extraction module with output stride eight followed by a classifier indicating the presence of segmentation failure. The first level of the feature extraction module consists of two convolutional layers. The first layer uses 16 kernels of 7×7 voxels and zero-padding of size 3 and second layer 32 kernels of 3×3 voxels and zero-padding of 1 voxel. Level 2 to 4 each consist of one residual block that contains two convolutional layers with 3×3 voxels kernels with zero-padding of size 1. The first convolutional layer of each residual block uses a strided convolution of 2 voxels to downsample the input. All convolutional layers of the feature extraction module are followed by batch normalization and ReLU function. The number of kernels in the feature extraction module increases from 16 in level 1 to 128 in level 4. The network is a 2D patch-level classifier and requires that the size of the two input slices is a multiple of the patch-size. The final classifier consists of three fully convolutional layers, implemented as 1×1 convolutions, with 128 feature maps in the first two layers. The final layer has two channels followed by a softmax function which indicates whether the patch contains segmentation failure. Furthermore, to regularize the model dropout layers ($p = 0.5$) were added between the residual blocks and the fully convolutional layers of the classifier.

Evaluation

Automatic segmentation performance, as well as performance after simulating the correction of detected segmentation failures and after manual expert correction was evaluated. For this, the 3D Dice-coefficient (DC) and 3D Hausdorff distance (HD) between manual and (corrected) automatic segmentation were computed. Furthermore, the following clinical metrics were computed for manual and (corrected) automatic segmentation: left ventricle end-diastolic volume (EDV); left ventricle ejection fraction (EF); right ventricle EDV; right ventricle ejection fraction; and left ventricle myocardial mass. Following Bernard et al.⁷ for each of the clinical metrics three performance indices were computed using the measurements based on manual and (corrected) automatic segmentation: Pearson correlation coefficient; mean difference (bias and standard deviation); and mean absolute error (MAE).

To evaluate detection performance of the automatic method precision-recall curves of identification of slices that require correction were computed. A slice is considered positive in case it consists of at least one image region with a segmentation failure. To achieve accurate segmentation in clinic, identification of slices that contain segmentation failures might ease manual correction of automatic segmentations in daily practice. To further evaluate detection performance detection rate of segmentation failures was assessed on a voxel level. More specific, sensitivity against the number of false positive regions was evaluated because manual correction is presumed to be performed at this level.

Finally, after simulation and manual correction of the automatically detected segmentation failures, segmentation was re-evaluated and significance of the difference between the DCs, HDs and clinical metrics was tested with a Mann-Whitney U test.

Experiments

To use stratified four-fold cross-validation the dataset was split into training (75%) and test (25%) set. The splitting was done on a patient level, so there was no overlap in patient data between training and test sets. Furthermore, patients were randomly chosen from each of the five patient groups w.r.t. disease. Each patient has one volume for ED and ES time points, respectively.

Training segmentation networks. DRN and U-net were trained with a patch size of 128×128 voxels which is a multiple of their output stride of the contracting path. In the training of the dilated CNN (DN) images with 151×151 voxel samples were used. Zero-padding to 281×281 was performed to accommodate the 131×131 voxel receptive field that is induced by the dilation factors. Training samples were randomly chosen from training set and augmented by 90 degree rotations of the images. All models were initially trained with three loss

functions: soft-Dice³³ (SD); cross-entropy (CE); and Brier loss³⁴. However, for the evaluation of the combined segmentation and detection approach for each model architecture the two best performing loss functions were chosen: soft-Dice for all models; cross-entropy for DRN and U-net and Brier loss for DN. For completeness, we provide the equations for all three used loss functions.

$$\text{soft-Dice}_c = \frac{\sum_{i=1}^N R_c(i) A_c(i)}{\sum_{i=1}^N R_c(i) + \sum_{i=1}^N A_c(i)}, \quad (3)$$

where N denotes the number of voxels in an image, R_c is the binary reference image for class c and A_c is the probability map for class c .

$$\text{Cross-Entropy}_c = - \sum_{i=1}^N t_{ic} \log p(y_i = c|x_i), \text{ where } t_{ic} = 1 \text{ if } y_i = c, \text{ and } 0 \text{ otherwise.} \quad (4)$$

$$\text{Brier}_c = \sum_{i=1}^N (t_{ic} - p(y_i = c|x_i))^2, \text{ where } t_{ic} = 1 \text{ if } y_i = c, \text{ and } 0 \text{ otherwise.} \quad (5)$$

where N denotes the number of voxels in an image and p denotes the probability for a specific voxel x_i with corresponding reference label y_i for class c .

Choosing Brier loss to train the DN model instead of CE was motivated by our preliminary work which showed that segmentation performance of DN model was best when trained with Brier loss²⁵.

All models were trained for 100,000 iterations. DRN and U-net were trained with a learning rate of 0.001 which decayed with a factor of 0.1 after every 25,000 steps. Training DN used the snapshot ensemble technique³⁵, where after every 10,000 iterations the learning rate was reset to its original value of 0.02.

All three segmentation networks were trained using mini-batch stochastic gradient descent using a batch size of 16. Network parameters were optimized using the Adam optimizer³⁶. Furthermore, models were regularized with weight decay to increase generalization performance.

Training detection network. To train the detection model a subset of the errors performed by the segmentation model is used. Segmentation errors that presumably are within the range of inter-observer variability and therefore do not inevitably require correction (tolerated errors) are excluded from the set of errors that need to be detected and corrected (segmentation failures). To distinguish between tolerated errors and the set of segmentation failures S_I the Euclidean distance of an incorrectly segmented voxel to the boundary of the reference target structure is used. For each anatomical structure a 2D distance transform map is computed that provides for each voxel the distance to the anatomical structure boundary. To differentiate between tolerated errors and the set of segmentation failures S_I an acceptable tolerance threshold is applied. A more rigorous threshold is used for errors located inside compared to outside of the anatomical structure because automatic segmentation methods have a tendency to undersegment cardiac structures in CMRI. Hence, in all experiments the acceptable tolerance threshold was set to three voxels (equivalent to on average 4.65 mm) and two voxels (equivalent to on average 3.1 mm) for segmentation errors located outside and inside the target structure. Furthermore, a segmentation error only belongs to S_I if it is part of a 2D 4-connected cluster of minimum size 10 voxels. This value was found in preliminary experiments by evaluating values {1, 5, 10, 15, 20}. However, for apical slices all segmentation errors are included in S_I regardless of fulfilling the minimum size requirement because in these slices anatomical structures are relatively small and manual segmentation is prone to large inter-observer variability⁷. Finally, segmentation errors located in slices above the base or below the apex are always included in the set of segmentation failures.

Using the set S_I a binary label t_j is assigned to each patch $P_j^{(I)}$ indicating whether $P_j^{(I)}$ contains at least one voxel belonging to set S_I where $j \in \{1 \dots M\}$ and M denotes the number of patches in a slice I .

The detection network is trained by minimizing a weighted binary cross-entropy loss:

$$\mathcal{L}_{DT} = - \sum_{j \in P^{(I)}} w_{pos} t_j \log p_j + (1 - t_j) \log(1 - p_j), \quad (6)$$

where w_{pos} represents a scalar weight, t_j denotes the binary reference label and p_j is the softmax probability indicating whether a particular image region $P_j^{(I)}$ contains at least one segmentation failure. The average percentage of regions in a patient volume containing segmentation failures ranges from 1.5 to 3 percent depending on the segmentation architecture and loss function used to train the segmentation model. To train a detection network w_{pos} was set to the ratio between the average percentage of negative samples divided by the average percentage of positive samples.

Each fold was trained using spatial uncertainty maps and automatic segmentation masks generated while training the segmentation networks. Hence, there was no overlap in patient data between training and test set across segmentation and detection tasks. In total 12 detection models were trained and evaluated resulting from the different combination of 3 model architectures (DRN, DN and U-net), 2 loss functions (DRN and U-net with CE and soft-Dice, DN with Brier and soft-Dice) and 2 uncertainty maps (e-maps, b-maps).

The patches used to train the network were selected randomly ($\frac{2}{3}$), or were forced ($\frac{1}{3}$) to contain at least one segmentation failure by randomly selecting a scan containing segmentation failure, followed by random sampling of a patch containing at least one segmentation failure. During training the patch size was fixed to $80 \times$

80 voxels. To reduce the number of background voxels during testing, inputs were cropped based on a minimal enclosing, rectangular bounding box that was placed around the automatic segmentation mask. Inputs always had a minimum size of 80×80 voxels or were forced to a multiple of the output grid spacing of eight voxels in both direction required by the patch-based detection network. The patches of size 8×8 voxels did not overlap. In cases where the automatic segmentation mask only contains background voxels (scans above the base or below apex of the heart) input scans were center-cropped to a size of 80×80 voxels.

Models were trained for 20,000 iterations using mini-batch stochastic gradient descent with batch-size 32 and Adam as optimizer³⁶. Learning rate was set to 0.0001 and decayed with a factor of 0.1 after 10,000 steps. Furthermore, dropout percentage was set to 0.5 and weight decay was applied to increase generalization performance.

Segmentation using correction of the detected segmentation failures. To investigate whether correction of detected segmentation failures increases segmentation performance two scenarios were performed. In the first scenario manual correction of the detected failures by an expert was simulated for all images at ED and ES time points of the ACDC dataset. For this purpose, in image regions that were detected to contain segmentation failure predicted labels were replaced with reference labels. In the second scenario manual correction of the detected failures was performed by an expert in a random subset of 50 patients of the ACDC dataset. The expert was shown CMRI slices for ED and ES time points together with corresponding automatic segmentation masks for the RV, LV and LV myocardium. Image regions detected to contain segmentation failures were indicated in slices and the expert was only allowed to change the automatic segmentations in these indicated regions. Annotation was performed following the protocol described in⁷. Furthermore, expert was able to navigate through all CMRI slices of the corresponding ED and ES volumes.

Results

In this section we first present results for the segmentation-only task followed by description of the combined segmentation and detection results.

Segmentation-only approach. Table 1 lists quantitative results for segmentation-only and combined segmentation and detection approach in terms of Dice coefficient and Hausdorff distance. These results show that DRN and U-net achieve similar Dice coefficients and outperformed the DN network for all anatomical structures at end-systole. Differences in the achieved Hausdorff distances among the methods are present for all anatomical structures and for both time points. The DRN model achieved the highest and the DN network the lowest Hausdorff distance.

Table 2 lists results of the evaluation in terms of clinical metrics. These results reveal noticeable differences between models for ejection fraction (EF) of left and right ventricle, respectively. We can observe that U-net trained with the soft-Dice and the Dilated Network (DN) trained with Brier or soft-Dice loss achieved considerable lower accuracy for LV and RV ejection fraction compared to DRN. Overall, the DRN model achieved highest performance for all clinical metrics.

Effect of model architecture on segmentation. Although quantitative differences between models are small, qualitative evaluation discloses that automatic segmentations differ substantially between the models. Figure 2 shows that especially in regions where the models perform poorly (apical and basal slices) the DN model more often produced anatomically implausible segmentations compared to the DRN and U-net. This seems to be correlated with the performance differences in Hausdorff distance.

Effect of loss function on segmentation. The results indicate that the choice of loss function only slightly affects the segmentation performance. DRN and U-net perform marginally better when trained with soft-Dice compared to cross-entropy whereas DN performs better when trained with Brier loss than with soft-Dice. For DN this is most pronounced for the RV at ES.

A considerable effect of the loss function on the accuracy of the LV and RV ejection fraction can be observed for the U-net model. On both metrics U-net achieved the lowest accuracy of all models when trained with the soft-Dice loss.

Effect of MC dropout on segmentation. The results show that enabling MC-dropout during testing seems to result in slightly improved HD while it does not affect DC.

Detection of segmentation failures. *Detection of segmentation failures on voxel level.* To evaluate detection performance of segmentation failures on voxel level Fig. 3a shows average voxel detection rate as a function of false positively detected regions. This was done for each combination of model architecture and loss function exploiting e- (Fig. 3a, left) or b-maps (Fig. 3a, right). These results show that detection performance of segmentation failures depends on segmentation model architecture, loss function and uncertainty map.

The influence of (segmentation) model architecture and loss function on detection performance is slightly stronger when e-maps were used as input for the detection task compared to b-maps. Detection rates are consistently lower when segmentation failures originate from segmentation models trained with soft-Dice loss compared to models trained with CE or Brier loss. Overall, detection rates are higher when b-maps were exploited for the detection task compared to e-maps.

Figure 2. Examples of automatic segmentations generated by different segmentation models for two cardiac MRI scans (rows) at ES at the base of the heart.

Detection of slices with segmentation failures. To evaluate detection performance w.r.t. slices containing segmentation failures precision-recall curves for each combination of model architecture and loss function using e-maps (Fig. 3b, left) or b-maps (Fig. 3b, right) are shown. The results show that detection performance of slices containing segmentation failures is slightly better for all models when using e-maps. Furthermore, the detection network achieves highest performance using uncertainty maps obtained from the DN model and the lowest when exploiting e- or b-maps obtained from the DRN model. Table 3 shows the average precision of detected slices with segmentation failures per patient, as well as the average percentage of slices that do contain segmentation failures (reference for detection task). The results illustrate that these measures are positively correlated i.e. that precision of detected slices in a patient volume is higher if the volume contains more slices that need correction. On average the DN model generates cardiac segmentations that contain more slices with at least one segmentation failure compared to U-net (ranks second) and DRN (ranks third). A higher number of detected slices containing segmentation failures implies an increased workload for manual correction.

Calibration of uncertainty maps. Figure 4 shows risk-coverage curves for each combination of model architectures, uncertainty maps and loss functions (Fig. 4 left: CE or Brier loss, Fig. 4 right: soft-Dice). The results show an effect of the loss function on slope and convergence of the curves. Segmentation errors of models trained with the soft-Dice loss are less frequently covered by higher uncertainties than models trained with CE or Brier loss (steeper slope and lower minimum are better). This difference is more pronounced for e-maps. Models trained with the CE or Brier loss only slightly differ concerning convergence and their slopes are approximately identical. In contrast, the curves of the models trained with the soft-Dice differ regarding their slope and achieved minimum. Comparing e- and b-map of the DN-SD and U-net-SD models the results reveal that the curve for b-map has a steeper slope and achieves a lower minimum compared to the e-map. For the DRN-SD model these differences are less striking. In general for a specific combination of model and loss function the risk-coverage curves using b-maps achieve a lower minimum compared to e-maps.

Correction of automatically identified segmentation failures. *Simulated correction.* The results listed in Tables 1 and 2 show that the proposed method consisting of segmentation followed by simulated manual correction of detected segmentation failures delivers accurate segmentation for all tissues over ED and ES points. Correction of detected segmentation failures improved the performance in terms of DC, HD and clinical metrics for all combinations of model architectures, loss functions and uncertainty measures. Focusing on the DC after correction of detected segmentation failures the results reveal that performance differences between evaluated models decreased compared to the segmentation-only task. This effect is less pronounced for HD where the DRN network clearly achieved superior results in the segmentation-only and combined approach. The DN performs the least of all models but achieves the highest absolute DC performance improvements in the combined approach for RV at ES. Overall, the results in Table 1 disclose that improvements attained by the combined approach are almost all statistically significant ($p \leq 0.05$) at ES and frequently at ED (96% resp. 83% of the cases). Moreover, improvements are in 99% of the cases statistically significant for HD compared to 81% of the cases for DC.

Model	Uncertainty	End-diastole			End-systole		
	map for detection	LV	RV	LVM	LV	RV	LVM
(a) Dice coefficient							
DN-Brier		0.962 ± 0.02	0.928 ± 0.04	0.875 ± 0.03	0.901 ± 0.11	0.832 ± 0.10	0.884 ± 0.04
	e-map	*0.965 ± 0.01	*0.949 ± 0.02	*0.885 ± 0.03	*0.937 ± 0.06	*0.905 ± 0.05	*0.909 ± 0.03
DN-Brier+MC		0.961 ± 0.02	0.922 ± 0.04	0.875 ± 0.04	0.912 ± 0.08	0.839 ± 0.11	0.882 ± 0.04
	b-map	*0.966 ± 0.01	*0.950 ± 0.01	*0.886 ± 0.03	*0.942 ± 0.03	*0.916 ± 0.04	*0.912 ± 0.03
DN-soft-Dice		0.960 ± 0.02	0.921 ± 0.04	0.870 ± 0.04	0.909 ± 0.08	0.812 ± 0.12	0.879 ± 0.04
	e-map	*0.965 ± 0.01	*0.945 ± 0.02	*0.879 ± 0.04	*0.938 ± 0.03	*0.891 ± 0.06	*0.905 ± 0.03
DN-soft-Dice+MC		0.958 ± 0.02	0.913 ± 0.05	0.868 ± 0.04	0.907 ± 0.07	0.818 ± 0.12	0.875 ± 0.04
	b-map	*0.964 ± 0.01	*0.944 ± 0.02	*0.877 ± 0.04	*0.939 ± 0.03	*0.900 ± 0.05	*0.904 ± 0.03
DRN-CE		0.961 ± 0.02	0.929 ± 0.03	0.878 ± 0.03	0.912 ± 0.06	0.850 ± 0.09	0.891 ± 0.03
	e-map	0.964 ± 0.01	*0.943 ± 0.02	*0.886 ± 0.03	*0.937 ± 0.03	*0.899 ± 0.04	*0.908 ± 0.03
DRN-CE+MC		0.961 ± 0.02	0.926 ± 0.03	0.877 ± 0.03	0.913 ± 0.06	0.847 ± 0.10	0.890 ± 0.03
	b-map	*0.965 ± 0.01	*0.948 ± 0.01	*0.887 ± 0.03	*0.939 ± 0.03	*0.911 ± 0.04	*0.909 ± 0.03
DRN-soft-Dice		0.964 ± 0.01	0.937 ± 0.02	0.888 ± 0.03	0.919 ± 0.06	0.856 ± 0.09	0.900 ± 0.03
	e-map	0.967 ± 0.01	*0.945 ± 0.02	0.893 ± 0.03	0.934 ± 0.04	*0.892 ± 0.06	*0.911 ± 0.03
DRN-soft-Dice+MC		0.963 ± 0.02	0.935 ± 0.03	0.886 ± 0.03	0.921 ± 0.06	0.857 ± 0.09	0.899 ± 0.03
	b-map	0.967 ± 0.01	*0.947 ± 0.02	0.893 ± 0.03	*0.938 ± 0.03	*0.907 ± 0.04	*0.912 ± 0.03
U-net-CE		0.962 ± 0.02	0.923 ± 0.05	0.878 ± 0.03	0.907 ± 0.07	0.840 ± 0.08	0.885 ± 0.03
	e-map	0.966 ± 0.01	*0.946 ± 0.02	*0.890 ± 0.03	*0.935 ± 0.04	*0.901 ± 0.06	*0.909 ± 0.03
U-net-CE+MC		0.962 ± 0.02	0.926 ± 0.04	0.879 ± 0.03	0.909 ± 0.07	0.849 ± 0.07	0.887 ± 0.03
	b-map	0.967 ± 0.01	*0.954 ± 0.02	*0.893 ± 0.03	*0.940 ± 0.04	*0.920 ± 0.04	*0.914 ± 0.03
U-net-soft-Dice		0.965 ± 0.02	0.928 ± 0.04	0.888 ± 0.03	0.914 ± 0.08	0.844 ± 0.09	0.896 ± 0.03
	e-map	0.968 ± 0.01	*0.943 ± 0.03	*0.898 ± 0.03	0.930 ± 0.05	*0.886 ± 0.07	*0.911 ± 0.03
U-net-soft-Dice+MC		0.965 ± 0.02	0.929 ± 0.04	0.889 ± 0.03	0.911 ± 0.10	0.845 ± 0.09	0.897 ± 0.03
	b-map	0.968 ± 0.01	*0.948 ± 0.03	*0.900 ± 0.03	0.928 ± 0.09	*0.895 ± 0.06	*0.914 ± 0.03
Isensee et al.		0.966	0.941	0.899	0.924	0.875	0.908
(b) Hausdorff Distance							
DN-Brier		6.7 ± 3.1	13.5 ± 5.9	10.2 ± 6.9	10.7 ± 7.7	16.7 ± 6.8	12.3 ± 5.8
	e-map	*5.7 ± 2.7	*11.7 ± 5.2	*8.3 ± 5.9	*8.0 ± 6.5	*14.2 ± 5.6	*9.7 ± 5.0
DN-Brier+MC		6.9 ± 3.3	13.1 ± 5.2	9.9 ± 5.9	9.9 ± 5.7	15.0 ± 6.1	12.0 ± 5.2
	b-map	*5.5 ± 2.6	*10.6 ± 5.1	*7.4 ± 4.2	*7.5 ± 6.0	*12.6 ± 5.6	*8.8 ± 4.0
DN-soft-Dice		7.1 ± 3.5	14.8 ± 6.8	11.0 ± 6.6	10.2 ± 5.6	17.7 ± 7.8	12.9 ± 6.2
	e-map	*5.6 ± 2.8	*12.6 ± 5.5	*8.6 ± 4.6	*8.0 ± 5.0	*14.6 ± 5.9	*9.6 ± 4.5
DN-soft-Dice+MC		7.7 ± 3.9	14.4 ± 6.0	10.5 ± 4.9	10.1 ± 5.3	17.2 ± 8.0	12.5 ± 5.3
	b-map	*6.3 ± 3.4	*11.5 ± 4.0	*8.6 ± 4.8	*7.8 ± 4.6	*13.6 ± 4.9	*9.6 ± 4.7
DRN-CE		5.5 ± 2.6	11.7 ± 5.4	8.2 ± 6.2	9.1 ± 6.4	13.7 ± 5.6	8.9 ± 5.3
	e-map	*4.5 ± 1.9	*9.0 ± 4.5	*6.3 ± 4.1	*6.2 ± 4.4	*11.1 ± 5.3	*6.7 ± 4.2
DRN-CE+MC		5.6 ± 2.6	11.9 ± 5.5	8.0 ± 5.9	8.7 ± 5.5	13.5 ± 5.9	8.5 ± 4.5
	b-map	*4.2 ± 1.6	*8.1 ± 3.7	*6.1 ± 4.2	*5.4 ± 3.6	*10.1 ± 5.5	*6.8 ± 3.8
DRN-soft-Dice		5.5 ± 2.8	11.9 ± 6.1	7.7 ± 5.9	8.5 ± 5.0	13.5 ± 5.5	8.9 ± 5.1
	e-map	*4.6 ± 2.2	*9.4 ± 4.5	6.7 ± 4.7	*6.7 ± 4.4	*11.6 ± 5.4	*7.0 ± 3.3
DRN-soft-Dice+MC		5.7 ± 3.2	11.5 ± 5.1	8.0 ± 5.5	8.3 ± 4.5	13.3 ± 5.1	8.9 ± 5.1
	b-map	*4.5 ± 2.2	*9.3 ± 4.5	*6.3 ± 4.0	*6.2 ± 4.1	*10.4 ± 5.0	*7.0 ± 3.4
U-net-CE		6.4 ± 4.3	15.7 ± 8.6	9.0 ± 6.0	9.7 ± 5.3	17.0 ± 7.7	12.7 ± 8.2
	e-map	*4.9 ± 3.9	*12.2 ± 8.1	*7.1 ± 5.6	*6.1 ± 3.2	*12.6 ± 6.5	*8.4 ± 6.3
U-net-CE+MC		6.2 ± 4.2	15.3 ± 8.4	8.8 ± 5.8	9.2 ± 5.0	16.5 ± 7.6	12.0 ± 8.0
	b-map	*4.3 ± 1.6	*9.9 ± 6.6	*6.7 ± 4.8	*5.4 ± 2.8	*10.3 ± 4.7	*7.6 ± 6.2
U-net-soft-Dice		6.1 ± 3.9	14.1 ± 7.6	10.6 ± 8.4	9.2 ± 7.1	16.3 ± 7.5	12.6 ± 9.6
	e-map	*4.6 ± 2.3	*11.3 ± 7.2	*7.5 ± 5.5	*7.3 ± 6.5	*13.7 ± 7.6	*9.8 ± 8.0
U-net-soft-Dice+MC		6.2 ± 3.9	14.1 ± 7.7	10.5 ± 8.7	9.0 ± 7.0	15.8 ± 7.5	12.1 ± 9.2
	b-map	*4.5 ± 2.1	*10.4 ± 7.2	*7.6 ± 7.0	*7.3 ± 6.9	*12.9 ± 6.6	*9.8 ± 8.4
Isensee et al.		7.1	14.3	8.9	9.8	16.3	10.4

Table 1. Segmentation performance of different combination of model architectures, loss functions and evaluation modes (without or with MC dropout enabled during testing) in terms of Dice coefficient (top) and Hausdorff distance (bottom) (mean $\pm$ standard deviation). Each combination comprises a block of two rows. A row in which column *Uncertainty map for detection* indicates e- or b-map shows results for the combined segmentation and detection approach. Numbers accentuated in bold are ranked first in the segmentation only task whereas numbers accentuated in italics are ranked first in the combined segmentation & detection task. The last row states the performance of the winning model in the ACDC challenge (on 100 patient images)³². Number with asterisk indicates statistical significant at 5% level w.r.t. the segmentation-only approach.

Method	Uncertainty map for detection	LV _{EDV}			LV _{EF}			RV _{EDV}			RV _{EF}			LVM _{Mass}		
		ρ	bias $\pm\sigma$	MAE	ρ	bias $\pm\sigma$	MAE	ρ	bias $\pm\sigma$	MAE	ρ	bias $\pm\sigma$	MAE	ρ	bias $\pm\sigma$	MAE
DN-Brier		0.997	0.0 $\pm$ 6.1	4.5	0.892	2.2 $\pm$ 9.2	4.2	0.977	-0.2 $\pm$ 11.8	8.5	0.834	5.3 $\pm$ 10.3	8.5	0.984	-2.7 $\pm$ 9.0	7.0
	e-map	0.997	0.0 $\pm$ 5.5	4.0	0.982	0.1 $\pm$ 3.8	2.2	0.992	0.0 $\pm$ 6.9	5.2	0.955	1.9 $\pm$ 5.5	4.1	0.986	-2.1 $\pm$ 8.4	6.6
DN-Brier+MC		0.997	1.6 $\pm$ 6.0	4.4	0.921	1.1 $\pm$ 7.9	3.9	0.975	6.7 $\pm$ 12.4	9.6	0.854	3.5 $\pm$ 9.9	7.7	0.984	0.7 $\pm$ 9.2	7.1
	b-map	0.998	1.0 $\pm$ 5.3	3.9	0.991	0.0 $\pm$ 2.7	1.9	0.993	3.2 $\pm$ 6.7	5.7	0.975	0.8 $\pm$ 4.0	3.0	0.987	0.1 $\pm$ 8.3	6.5
DN-soft-Dice		0.996	1.2 $\pm$ 6.5	4.9	0.918	1.5 $\pm$ 8.0	3.9	0.972	0.2 $\pm$ 13.0	9.6	0.802	7.2 $\pm$ 11.3	10.2	0.982	-4.5 $\pm$ 9.6	8.5
	e-map	0.997	1.0 $\pm$ 5.5	4.2	0.989	0.2 $\pm$ 3.0	2.2	0.990	0.2 $\pm$ 7.6	5.9	<i>0.940</i>	3.3 $\pm$ 6.2	5.2	0.983	-4.3 $\pm$ 9.3	8.2
DN-soft-Dice+MC		0.996	3.2 $\pm$ 7.1	5.6	0.958	0.4 $\pm$ 5.7	3.6	0.964	8.1 $\pm$ 14.9	12.3	0.827	4.8 $\pm$ 11.0	8.9	0.978	-0.7 $\pm$ 10.7	8.3
	b-map	0.997	2.2 $\pm$ 5.6	4.4	0.988	-0.2 $\pm$ 3.1	2.2	0.990	4.0 $\pm$ 7.7	7.0	0.959	1.8 $\pm$ 5.1	4.1	0.982	-1.4 $\pm$ 9.5	7.6
DRN-CE		0.997	-0.2 $\pm$ 5.5	4.1	0.968	1.2 $\pm$ 5.0	3.5	0.976	1.5 $\pm$ 12.1	8.5	0.870	1.3 $\pm$ 9.2	6.9	0.980	0.6 $\pm$ 10.2	7.8
	e-map	0.998	0.2 $\pm$ 4.5	3.5	0.992	0.2 $\pm$ 2.5	1.9	0.988	1.4 $\pm$ 8.5	6.2	0.952	0.8 $\pm$ 5.6	4.2	0.985	0.4 $\pm$ 8.7	6.8
DRN-CE+MC		0.998	1.0 $\pm$ 4.9	3.9	0.972	0.8 $\pm$ 4.6	3.1	0.973	4.8 $\pm$ 12.8	9.4	0.876	0.4 $\pm$ 9.1	6.6	0.981	1.9 $\pm$ 9.9	7.6
	b-map	0.998	0.7 $\pm$ 4.6	3.6	0.992	-0.1 $\pm$ 2.5	1.8	0.992	2.9 $\pm$ 6.9	5.7	0.967	0.6 $\pm$ 4.6	3.4	0.987	1.2 $\pm$ 8.3	6.6
DRN-soft-Dice		0.998	0.8 $\pm$ 5.1	4.0	0.976	0.2 $\pm$ 4.4	3.0	0.980	0.2 $\pm$ 11.0	7.5	0.882	3.1 $\pm$ 8.7	6.8	0.984	-3.5 $\pm$ 9.1	7.5
	e-map	0.998	0.7 $\pm$ 4.4	3.5	0.987	-0.1 $\pm$ 3.1	2.2	0.987	0.1 $\pm$ 9.1	6.4	0.938	1.9 $\pm$ 6.3	4.9	0.986	-3.5 $\pm$ 8.7	7.1
DRN-soft-Dice+MC		0.998	1.8 $\pm$ 5.1	3.9	0.979	-0.3 $\pm$ 4.1	2.9	0.977	3.5 $\pm$ 11.7	8.1	0.868	1.7 $\pm$ 9.5	6.8	0.983	-1.4 $\pm$ 9.5	7.4
	b-map	0.998	1.7 $\pm$ 4.7	3.7	0.990	-0.2 $\pm$ 2.9	2.1	0.989	2.3 $\pm$ 8.1	5.8	0.959	0.8 $\pm$ 5.2	3.8	0.986	-1.3 $\pm$ 8.5	6.8
U-net-CE		0.995	-4.7 $\pm$ 7.2	6.1	0.954	4.1 $\pm$ 6.0	5.1	0.963	-7.6 $\pm$ 15.2	12.1	0.870	5.6 $\pm$ 9.0	8.1	0.971	-8.5 $\pm$ 12.2	11.5
	e-map	0.998	-3.2 $\pm$ 4.8	4.4	0.992	1.7 $\pm$ 2.6	2.4	0.987	-4.1 $\pm$ 9.1	6.7	0.957	2.6 $\pm$ 5.2	4.1	0.983	-5.7 $\pm$ 9.3	8.2
U-net-CE+MC		0.995	-4.3 $\pm$ 7.2	5.9	0.958	3.8 $\pm$ 5.8	4.9	0.968	-4.8 $\pm$ 14.1	10.7	0.867	5.0 $\pm$ 9.1	7.9	0.972	-8.1 $\pm$ 12.0	11.1
	b-map	0.997	-3.5 $\pm$ 5.5	4.9	0.990	1.6 $\pm$ 2.9	2.6	0.992	-1.8 $\pm$ 7.0	4.9	0.974	1.6 $\pm$ 4.1	3.3	0.981	-6.8 $\pm$ 10.0	9.4
U-net-soft-Dice		0.997	-2.0 $\pm$ 6.0	4.5	0.853	3.6 $\pm$ 10.9	5.0	0.968	-1.0 $\pm$ 14.1	10.0	0.782	4.8 $\pm$ 11.6	9.0	0.985	-7.7 $\pm$ 8.8	9.2
	e-map	0.997	-1.7 $\pm$ 5.3	4.1	0.969	1.9 $\pm$ 4.9	3.3	0.981	-0.1 $\pm$ 10.9	7.5	0.919	3.3 $\pm$ 7.0	5.9	0.984	-6.6 $\pm$ 9.0	8.7
U-net-soft-Dice+MC		0.997	-1.8 $\pm$ 5.9	4.4	0.941	3.0 $\pm$ 6.7	4.4	0.969	0.6 $\pm$ 13.9	9.8	0.792	4.4 $\pm$ 11.3	8.7	0.985	-7.2 $\pm$ 8.9	8.9
	b-map	0.997	-1.5 $\pm$ 5.3	4.1	0.979	1.1 $\pm$ 4.1	2.9	0.985	1.2 $\pm$ 9.4	6.5	0.939	2.9 $\pm$ 6.2	4.9	0.984	-5.9 $\pm$ 9.0	8.5

Table 2. Segmentation performance of different combination of model architectures, loss functions and evaluation modes (without or with MC dropout (MC) enabled during testing) in terms of clinical metrics: left ventricle (LV) end-diastolic volume (EDV); LV ejection fraction (EF); right ventricle (RV) EDV; RV ejection fraction; and LV myocardial mass. Quantitative results compare clinical metrics based on reference segmentations with (1) automatic segmentations and (2) simulated manual correction of automatic segmentations using spatial uncertainty maps. ρ denotes the Pearson correlation coefficient, *bias* denotes the mean difference between the two measurements (mean $\pm$ standard deviation) and *MAE* denotes the mean absolute error between the two measurements. Each combination comprises a block of two rows. A row in which column *Uncertainty map for detection* indicates e- or b-map shows results for the combined segmentation and detection approach. Numbers accentuated in bold are ranked first in the segmentation only task. Numbers in italics indicate statistical significant at 5% level w.r.t. the segmentation-only approach for the specific clinical metric.

Results in terms of clinical metrics shown in Table 2 are inline with these findings. We observe that segmentation followed by simulated manual correction of detected segmentation failures resulted in considerably higher accuracy for LV and RV ejection fraction. Achieved improvements for clinical metrics are only statistically significant ($p \leq 0.05$) in one case for RV ejection fraction.

Figure 3. Detection performance of segmentation failures generated by different combination of segmentation architectures and loss functions. **(a)** Sensitivity for detection of segmentation failures on voxel level (y-axis) as a function of number of false positive image regions (x-axis). **(b)** Precision-recall curve for detection of slices containing segmentation failures (where AP denotes average precision). Results are split between entropy and Bayesian uncertainty maps. Each figure contains a curve for the six possible combination of models (three) and loss functions (two). SD denotes soft-Dice and CE cross-entropy, respectively.

Model	Average precision		% of slices with segmentation failures	
	e-map	b-map	e-map	b-map
DN-Brier	84.0	83.0	53.7	52.4
DN-SD	87.0	85.0	58.3	58.1
DRN-CE	75.0	69.0	39.5	39.4
DRN-SD	67.0	67.0	34.9	33.7
U-net-CE	81.0	75.0	54.8	52.5
U-net-SD	76.0	76.0	46.7	45.5

Table 3. Average precision and percentage of slices with segmentation failures generated by Dilated Network (DN), Dilated Residual Network (DRN) and U-net when trained with soft-Dice (SD), CE or Brier loss. Per patient, average precision of detected slices with failure using e- or b-maps (2nd and 3rd columns). Per patient, average percentage of slices containing segmentation failures (reference for detection task) (4th and 5th columns).

In general, the effect of correction of detected segmentation failures is more pronounced in cases where the segmentation-only approach achieved relatively low accuracy (e.g. DN-SD for RV at ES). Furthermore, performance gains are largest for RV and LV at ES and for ejection fraction of both ventricles.

The best overall performance is achieved by the DRN model trained with cross-entropy loss while exploiting entropy maps in the detection task. Moreover, the proposed two step approach attained slightly better results using Bayesian maps compared to entropy maps.

Manual correction. Table 4 lists results for the combined automatic segmentation and detection approach followed by *manual* correction of detected segmentation failures by an expert. The results show that this correction led to improved segmentation performance in terms of DC, HD and clinical metrics. Improvements in terms of HD are in 50 percent of the cases statistically significant ($p \leq 0.05$) and most pronounced for RV and LV at end-systole.

Qualitative examples of the proposed approach are visualized in Figs. 5 and 6 for simulated correction and manual correction of the automatically detected segmentation failures, respectively. For the illustrated cases (simulated) manual correction of detected segmentation failures leads to increased segmentation performance. On average manual correction of automatic segmentations took less than 2 min for ED and ES volumes of one patient compared to 20 min that is typically needed by an expert for the same task.

Ablation study

To demonstrate the effect of different hyper-parameters in the method, a number of experiments were performed. In the following text these are detailed.

Impact of number of Monte Carlo samples on segmentation performance. To investigate the impact of the number of Monte Carlo (MC) samples on the segmentation performance validation experiments were performed for all three segmentation architectures (Dilated Network, Dilated Residual Network and U-net)

Scenario	End-diastole			End-systole		
	LV	RV	LVM	LV	RV	LVM
(a) Dice coefficient: Mean $\pm$ standard deviation for left ventricle (LV), right ventricle (RV) and left ventricle myocardium (LVM)						
Auto-only	0.964 $\pm$ 0.02	0.927 $\pm$ 0.04	0.883 $\pm$ 0.03	0.916 $\pm$ 0.05	0.854 $\pm$ 0.08	0.886 $\pm$ 0.04
Auto+simulation	0.967 $\pm$ 0.01	*0.948 $\pm$ 0.03	*0.894 $\pm$ 0.03	*0.939 $\pm$ 0.03	*0.915 $\pm$ 0.04	*0.910 $\pm$ 0.03
Auto+expert	0.965 $\pm$ 0.02	0.940 $\pm$ 0.03	0.885 $\pm$ 0.03	0.927 $\pm$ 0.04	0.868 $\pm$ 0.07	0.894 $\pm$ 0.03
Scenario	End-diastole			End-systole		
	LV	RV	LVM	LV	RV	LVM
(b) Hausdorff Distance: Mean $\pm$ standard deviation for left ventricle (LV), right ventricle (RV) and left ventricle myocardium (LVM)						
Auto-only	5.6 $\pm$ 3.3	15.7 $\pm$ 9.7	8.5 $\pm$ 6.4	9.2 $\pm$ 5.8	16.5 $\pm$ 8.8	13.4 $\pm$ 10.5
Auto+simulation	4.5 $\pm$ 2.1	*9.0 $\pm$ 4.6	*5.9 $\pm$ 3.4	*5.2 $\pm$ 2.5	*10.3 $\pm$ 3.7	*6.6 $\pm$ 2.9
Auto+expert	4.9 $\pm$ 2.8	*9.8 $\pm$ 4.3	7.3 $\pm$ 4.3	7.2 $\pm$ 3.3	*12.5 $\pm$ 4.7	*8.3 $\pm$ 3.5
Scenario	LV _{EDV}			RV _{EDV}		
	ρ	bias $\pm \sigma$	MAE	ρ	bias $\pm \sigma$	MAE
(c) Clinical metrics: (a) left ventricle (LV) end-diastolic volume (EDV), (b) LV ejection fraction (EF), (c) right ventricle (RV) EDV, (d) RV ejection fraction, (e) LV myocardial mass. Quantitative results compare clinical metrics based on reference segmentations with (1) automatic segmentations; (2) simulated manual correction and (3) manual expert correction of automatic segmentations using spatial uncertainty maps. ρ denotes the Pearson correlation coefficient, $bias$ denotes the mean difference between the two measurements (mean $\pm$ standard deviation) and MAE denotes the mean absolute error between the two measurements						
Auto-only	0.995	-4.4 $\pm$ 7.0	5.7	0.927	5.0 $\pm$ 7.1	5.8
Auto+simulation	0.998	-3.9 $\pm$ 5.2	4.8	0.989	2.3 $\pm$ 2.9	2.9
Auto+expert	0.996	-4.3 $\pm$ 6.5	5.5	0.968	2.7 $\pm$ 4.8	4.3
Scenario	RV _{EF}			LVM _{Mass}		
	ρ	bias $\pm \sigma$	MAE	ρ	bias $\pm \sigma$	MAE
Auto-only	0.979	-6.4 $\pm$ 16.2	11.9	0.878	5.8 $\pm$ 8.7	8.0
Auto+simulation	0.983	-3.7 $\pm$ 10.4	6.8	0.954	2.7 $\pm$ 5.5	4.5
Auto+expert	0.980	-3.2 $\pm$ 12.9	8.3	0.883	5.1 $\pm$ 8.6	7.7

Table 4. Comparing performance of segmentation-only approach (auto-only) with combined segmentation and detection approach for two scenarios: simulated correction of detected segmentation failures (auto+simulation); and manual correction of detected segmentation failures by an expert (auto+expert). Automatic segmentations were obtained from a U-net trained with cross-entropy. Evaluation was performed on a subset of 50 patients from the ACD dataset. Scenarios are compared against segmentation-only approach (auto-only) in terms of (a) Dice Coefficient (b) Hausdorff Distance and (c) Clinical metrics. Results obtained from simulated manual correction represent an upper bound on the maximum achievable performance. Detection network was trained with e-maps. Number with asterisk indicates statistical significant at 5% level w.r.t. the segmentation-only approach.

Figure 4. Comparison of risk-coverage curves for different combination of model architectures, loss functions and uncertainty maps. Results are separated for loss functions (left cross-entropy and Brier, right soft-Dice loss). 100% coverage means that none of the voxels is discarded based on its uncertainty whereas a coverage of 0% denotes the scenario in which all predictions are replaced by their reference labels. Note, all models were trained with two different loss functions (1) soft-Dice (SD) for all models (2) cross-entropy (CE) for DRN and U-net and Brier loss for DN.

Number of MC samples	DRN-CE	U-net-CE	DN-soft-Dice
(a) Dice coefficient			
1	0.894 ± 0.07	0.896 ± 0.07	0.871 ± 0.09
3	0.900 ± 0.07	0.901 ± 0.07	0.883 ± 0.08
5	0.902 ± 0.07	0.901 ± 0.07	0.887 ± 0.08
7	0.903 ± 0.07	0.901 ± 0.07	0.888 ± 0.08
10	0.904 ± 0.06	0.902 ± 0.07	0.890 ± 0.08
20	0.904 ± 0.07	0.902 ± 0.07	0.890 ± 0.08
30	0.904 ± 0.07	0.902 ± 0.07	0.891 ± 0.08
60	0.904 ± 0.07	0.902 ± 0.07	0.891 ± 0.08
(b) Hausdorff Distance			
1	9.88 ± 5.76	11.79 ± 8.23	13.54 ± 7.14
3	9.70 ± 6.13	11.40 ± 7.78	12.71 ± 6.79
5	9.54 ± 6.07	11.37 ± 7.81	12.06 ± 6.29
7	9.38 ± 5.86	11.29 ± 7.86	12.08 ± 6.38
10	9.38 ± 5.91	11.24 ± 7.71	11.85 ± 6.34
20	9.37 ± 5.83	11.27 ± 7.79	11.90 ± 6.52
30	9.39 ± 5.91	11.32 ± 7.93	11.90 ± 6.48
60	9.39 ± 5.93	11.22 ± 7.83	11.89 ± 6.56

Table 5. Effect of number of Monte Carlo (MC) samples on segmentation performance in terms of (a) dice coefficient (DC) and (b) Hausdorff distance (HD) (mean ± standard deviation). Higher DC and lower HD is better. Abbreviations: Cross-Entropy (CE), Dilated Residual Network (DRN) and Dilated Network (DN).

using $T \in \{1, 3, 5, 7, 10, 20, 30, 60\}$ samples. Results of these experiments are listed in Table 5. We observe that segmentation performance started to converge using seven samples only. Performance improvements using an increased number of MC samples were largest for the Dilated Network. Overall, using more than ten samples did not increase segmentation performance. Hence, in the presented work T was set to 10.

Effect of patch-size on detection performance. The combined segmentation and detection approach detects segmentation failures on region level. To investigate the effect of patch-size on detection performance three different patch-sizes were evaluated: 4×4 , 8×8 , and 16×16 voxels. The results are shown in Fig. 7. We

Figure 5. Three patients showing results of combined segmentation and detection approach consisting of segmentation followed by simulated manual correction of detected segmentation failures. First column shows MRI (top) and reference segmentation (bottom). Results for automatic segmentation and simulated manual correction respectively achieved by: Dilated Network (DN-Brier, 2nd and 5th columns); Dilated Residual Network (DRN-soft-Dice, 3rd and 6th columns); and U-net (soft-Dice, 4th and 7th columns).

can observe in Fig. 7a that larger patch-sizes result in a lower number of false positive regions. The result is potentially caused by the decreasing number of regions in an image when using larger patch-sizes compared to smaller patch-sizes. Furthermore, Fig. 7b reveals that slice detection performance is only slightly influenced by patch-size. To ease manual inspection and correction by an expert, it is desirable to keep region-size i.e. patch-size small. Therefore, in the experiments a patch-size of 8×8 voxels was used.

Impact of tolerance threshold on number of segmentation failures. To investigate the impact of the tolerance threshold separating segmentation failures and tolerable segmentation errors, we calculated the

Figure 6. Four patients showing results of combined segmentation and detection approach consisting of segmentation followed by manual expert correction of detected segmentation failures. Expert was only allowed to adjust the automatic segmentations in regions where the detection network predicted segmentation failures (orange contour shown in 2nd column). Automatic segmentations were generated by a U-net trained with the cross-entropy loss. Segmentation failure detection was performed using entropy maps.

Figure 7. Detection performance for three different patch-sizes specified in voxels. **(a)** Sensitivity for detection of segmentation failures on voxel level (y-axis) versus number of false positive image regions (x-axis). **(b)** Precision-recall curve for detection of slices containing segmentation failures (where AP denotes average precision). Results are split between entropy and Bayesian uncertainty maps. In the experiments patch-size was set to 8×8 voxels.

ratio of the number of segmentation failures and all errors i.e. the sum of tolerable errors and segmentation failures. Fig. 8 shows the results. We observe that at least half of the segmentation failures are located within a tolerance threshold i.e. distance of two to three voxels of the target structure boundary as defined by the reference annotation. Furthermore, the mean percentage of failures per volume is considerably lower for the Dilated Residual Network (DRN) and highest for the Dilated Network. This result is inline with our earlier finding (see Table 3) that average percentage of slices that do contain segmentation failures is lowest for the DRN model.

Discussion

We have described a method that combines automatic segmentation and assessment of uncertainty in cardiac MRI with detection of image regions containing segmentation failures. The results show that combining automatic segmentation with manual correction of detected segmentation failures results in higher segmentation performance. In contrast to previous methods that detected segmentation failures per patient or per structure, we showed that it is feasible to detect segmentation failures per image region. In most of the experimental settings, simulated manual correction of detected segmentation failures for LV, RV and LVM at ED and ES led to statistically significant improvements. These results represent the upper bound on the maximum achievable performance for the manual expert correction task. Furthermore, results show that manual expert correction of detected segmentation failures led to consistently improved segmentations. However, these results are not on par with the simulated expert correction scenario. This is not surprising because inter-observer variability is high for the presented task and annotation protocols may differ between clinical environments. Moreover, qualitative results of the manual expert correction reveal that manual correction of the detected segmentation failures can prevent anatomically implausible segmentations (see Fig. 6). Therefore, the presented approach can

Figure 8. Mean percentage of the segmentation failures per volume (y-axis) in the set of all segmentation errors (tolerable errors+segmentation failures) depending on the tolerance threshold (x-axis). Red, dashed vertical line indicates threshold value that was used throughout the experiments. Results are split between segmentation errors located (a) outside and (b) inside the target structure. Each figure contains a curve for U-net, Dilated Network (DN) and Dilated Residual Network (DRN) trained with the soft-Dice (SD) loss. Segmentation errors located in slices above the base or below the apex are always included in the set of segmentation failures and therefore, they are independent of the applied tolerance threshold.

potentially simplify and accelerate correction process and has the capacity to increase the trustworthiness of existing automatic segmentation methods in daily clinical practice.

The proposed combined segmentation and detection approach was evaluated using three state-of-the-art deep learning segmentation architectures. The results suggest that our approach is generic and applicable to different model architectures. Nevertheless, we observe noticeable differences between the different combination of model architectures, loss functions and uncertainty measures. In the segmentation-only task the DRN clearly outperforms the other two models in the evaluation of the boundary of the segmented structure. Moreover, qualitative analysis of the automatic segmentation masks suggests that DRN generates less often anatomically implausible and fragmented segmentations than the other models. We assume that clinical experts would prefer such segmentations although they are not always perfect. Furthermore, even though DRN and U-net achieve similar performance in regard to DC we assume that less fragmented segmentation masks would increase trustworthiness of the methods.

In agreement with our preliminary work we found that uncertainty maps obtained from a segmentation model trained with soft-Dice loss have a lower degree of uncertainty calibration compared to when trained with one of the other two loss functions (cross-entropy and Brier)²⁵. Nevertheless, the results of the combined segmentation and detection approach showed that a lower degree of uncertainty calibration only slightly deteriorated the detection performance of segmentation failures for the larger segmentation models (DRN and U-net) when exploiting uncertainty information from e-maps. Hendrycks and Gimpel³⁷ showed that softmax probabilities generated by deep learning networks have poor direct correspondence to confidence. However, in agreement with Geifman et al.³⁰ we presume that probabilities and hence corresponding entropies obtained from softmax function are ranked consistently i.e. entropy can potentially be used as a relative uncertainty measure in deep learning. In addition, we detect segmentation failures per image region and therefore, our approach does not require perfectly calibrated uncertainty maps. Furthermore, results of the combined segmentation and detection approach revealed that detection performance of segmentation failures using b-maps is almost independent of the loss function used to train the segmentation model. In line with Jungo et al.³⁸ we assume that by enabling MC-dropout in testing and computing the mean softmax probabilities per class leads to better calibrated probabilities and b-maps. This assumption is in agreement with Srivastava et al.³⁹ where a CNN with dropout used at testing is interpreted as an ensemble of models.

Quantitative evaluation in terms of Dice coefficient and Hausdorff distance reveals that proposed combined segmentation and detection approach leads to significant performance increase. However, the results also demonstrate that the correction of the detected failures allowed by the combined approach does not lead to statistically significant improvement in clinical metrics. This is not surprising because state-of-the-art automatic segmentation methods are not expected to lead to large volumetric errors⁷ and standard clinical measures are not sensitive to small segmentation errors. Nevertheless, errors of the current state-of-the-art automatic segmentation methods may lead to anatomically implausible segmentations⁷ that may cause distrust in clinical application. Besides increase in trustworthiness of current state-of-the-art segmentation methods for cardiac MRIs, improved segmentations are a prerequisite for advanced functional analysis of the heart e.g. motion analysis⁴⁰ and very detailed morphology analysis such as myocardial trabeculae in adults⁴¹.

For the ACDC dataset used in this manuscript, Bernard et al.⁷ reported inter-observer variability ranging from 4 to 14.1 mm (equivalent to on average 2.6 to 9 voxels). To define the set of segmentation failures, we employed a strict tolerance threshold on distance metric to distinguish between tolerated segmentation errors

and segmentation failures (see Ablation study). Stricter tolerance threshold was used because the thresholding is performed in 2D, while evaluation of segmentation is done in 3D. Large slice thickness in cardiac MR could lead to a discrepancy between the two. As a consequence of this strict threshold results listed in Table 3 show that almost all patient volumes contain at least one slice with a segmentation failure. This might render the approach less feasible in clinical practice. Increasing the threshold decreases the number of segmentation failures and slices containing segmentation failures (see Fig. 8) but also lowers the upper bound on the maximum achievable performance. Therefore, to show the potential of our proposed approach we chose to apply a strict tolerance threshold. Nevertheless, we realize that although manual correction of detected segmentation failures leads to increased segmentation accuracy the performance of precision-recall is limited (see Fig. 3) and hence, should be a focus of future work.

The presented patch-based detection approach combined with (simulated) manual correction can in principle lead to stitching artefacts in the resulting segmentation masks. A voxel-based detection approach could potentially solve this. However, voxel-based detection methods are more challenging to train due to the very small number of voxels in an image belonging to the set of segmentation failures.

Evaluation of the proposed approach for 12 possible combinations of segmentation models (three), loss functions (two) and uncertainty maps (two) resulted in an extensive number of experiments. Nevertheless, future work could extend evaluation to other segmentation models, loss functions or combination of losses. Furthermore, our approach could be evaluated using additional uncertainty estimation techniques e.g. by means of ensembling of networks⁴² or variational dropout⁴³. In addition, previous work by Kendall and Gal⁴⁴, Tanno et al.⁴⁵ has shown that the quality of uncertainty estimates can be improved if model (epistemic) and data (aleatoric) uncertainty are assessed simultaneously with separate measures. The current study focused on the assessment of model uncertainty by means of MC-dropout and entropy which is a combination of epistemic and aleatoric uncertainty. Hence, future work could investigate whether additional estimation of aleatoric uncertainty improves the detection of segmentation failures.

Furthermore, to develop an end-to-end approach future work could incorporate the detection of segmentation failures into the segmentation network. Besides, adding the automatic segmentations to the input of the detection network could increase the detection performance.

Finally, the proposed approach is not specific to cardiac MRI segmentation. Although data and task specific training would be needed the approach could potentially be applied to other image modalities and segmentation tasks.

Conclusion

A method combining automatic segmentation and assessment of segmentation uncertainty in cardiac MR with detection of image regions containing local segmentation failures has been presented. The combined approach, together with simulated and manual correction of detected segmentation failures, increases performance compared to segmentation-only. The proposed method has the potential to increase trustworthiness of current state-of-the-art segmentation methods for cardiac MRIs.

Code availability

All models were implemented using the PyTorch⁴⁶ framework and trained on one Nvidia GTX Titan X GPU with 12 GB memory. The code to replicate the study is publicly available at <https://github.com/toologicbv/cardiacSegUncertainty>.

Received: 2 June 2020; Accepted: 13 November 2020

Published online: 10 December 2020

References

- White, H. D. *et al.* Left ventricular end-systolic volume as the major determinant of survival after recovery from myocardial infarction. *Circulation* **76**, 44–51 (1987).
- Grothues, F. *et al.* Comparison of interstudy reproducibility of cardiovascular magnetic resonance with two-dimensional echocardiography in normal subjects and in patients with heart failure or left ventricular hypertrophy. *Am. J. Cardiol.* **90**, 29–34 (2002).
- Litjens, G. *et al.* A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017).
- Leiner, T. *et al.* Machine learning in cardiovascular magnetic resonance: basic concepts and applications. *J. Cardiovasc. Magnet. Reson.* **21**, 61 (2019).
- Petitjean, C. & Dacher, J.-N. A review of segmentation methods in short axis cardiac mr images. *Med. Image Anal.* **15**, 169–184 (2011).
- Peng, P. *et al.* A review of heart chamber segmentation for structural and functional analysis using cardiac magnetic resonance imaging. *Magnet. Reson. Mater. Phys. Biol. Med.* **29**, 155–195 (2016).
- Bernard, O. *et al.* Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? Dataset available at <https://www.creatis.insa-lyon.fr/Challenge/acdc/databases.html> (2018).
- Suinesiaputra, A. *et al.* Quantification of lv function and mass by cardiovascular magnetic resonance: multi-center variability and consensus contours. *J. Cardiovasc. Magnet. Reson.* **17**, 63 (2015).
- Tan, L. K., McLaughlin, R. A., Lim, E., Abdul Aziz, Y. F. & Liew, Y. M. Fully automated segmentation of the left ventricle in cine cardiac MRI using neural network regression. *J. Magnet. Reson. Imaging* **48**, 140–152 (2018).
- Zheng, Q., Delingette, H., Duchateau, N. & Ayache, N. 3-d consistent and robust segmentation of cardiac images by deep learning with spatial propagation. *IEEE Trans. Med. Imaging* **37**, 2137–2148 (2018).
- Savioli, N., Vieira, M. S., Lamata, P. & Montana, G. Automated segmentation on the entire cardiac cycle using a deep learning work-flow. in *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, 153–158 (IEEE, 2018).
- Bai, W. *et al.* Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *J. Cardiovasc. Magnet. Reson.* **20**, 65 (2018).

13. Oktay, O. *et al.* Anatomically constrained neural networks (acnns): application to cardiac image enhancement and segmentation. *IEEE Trans. Med. Imaging* **37**, 384–395 (2017).
14. Duan, J. *et al.* Automatic 3D bi-ventricular segmentation of cardiac images by a shape-refined multi-task deep learning approach. *IEEE Trans. Med. Imaging* (2019).
15. Painchaud, N. *et al.* Cardiac MRI segmentation with strong anatomical guarantees. in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 632–640 (Springer, 2019).
16. Albà, X. *et al.* Automatic initialization and quality control of large-scale cardiac MRI segmentations. *Med. Image Anal.* **43**, 129–141 (2018).
17. Robinson, R. *et al.* Automated quality control in image segmentation: application to the uk biobank cardiovascular magnetic resonance imaging study. *J. Cardiovasc. Magnet. Reson.* **21**, 18 (2019).
18. Valindria, V. V. *et al.* Reverse classification accuracy: predicting segmentation performance in the absence of ground truth. *IEEE Trans. Med. Imaging* **36**, 1597–1606 (2017).
19. Frounchi, K., Briand, L. C., Grady, L., Labiche, Y. & Subramanyan, R. Automating image segmentation verification and validation by learning test oracles. *Inf. Softw. Technol.* **53**, 1337–1348 (2011).
20. Kohlberger, T., Singh, V., Alvino, C., Bahlmann, C. & Grady, L. Evaluating segmentation error without ground truth. in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 528–536 (Springer, 2012).
21. Nair, T., Precup, D., Arnold, D. L. & Arbel, T. Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 655–663 (Springer, 2018).
22. Jungo, A., Meier, R., Ermis, E., Herrmann, E. & Reyes, M. Uncertainty-driven sanity check: Application to postoperative brain tumor cavity segmentation. in *Medical Imaging with Deep Learning Conference* (2018).
23. Roy, A. G. *et al.* Bayesian quicknat: Model uncertainty in deep whole-brain segmentation for structure-wise quality control. *NeuroImage* **195**, 11–22 (2019).
24. DeVries, T. & Taylor, G. W. Leveraging uncertainty estimates for predicting segmentation quality. arXiv preprint [arXiv:1807.00502](https://arxiv.org/abs/1807.00502) (2018).
25. Sander, J., de Vos, B. D., Wolterink, J. M. & Išgum, I. Towards increased trustworthiness of deep learning segmentation methods on cardiac MRI. in *Medical Imaging 2019: Image Processing*, Vol. 10949, 1094919 (International Society for Optics and Photonics, 2019).
26. Gal, Y. & Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. in *International Conference on Machine Learning (ICML)* 1050–1059 (2016).
27. Wolterink, J. M., Leiner, T., Viergever, M. A. & Išgum, I. Automatic segmentation and disease classification using cardiac cine MR images. in *International Workshop on Statistical Atlases and Computational Models of the Heart*, 101–110 (Springer, 2017).
28. Yu, F., Koltun, V. & Funkhouser, T. Dilated residual networks. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 472–480 (2017). Code available at <https://github.com/fyu/drn>.
29. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241 (Springer, 2015).
30. Geifman, Y. & El-Yaniv, R. Selective classification for deep neural networks. *Adv. Neural Inf. Process. Syst.* 4878–4887 (2017).
31. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (2016).
32. Isensee, F. *et al.* Automatic cardiac disease assessment on cine-MRI via time-series segmentation and domain specific features. in *International Workshop on Statistical Atlases and Computational Models of the Heart*, 120–129 (Springer, 2017).
33. Milletari, F., Navab, N. & Ahmadi, S.-A. V-net: Fully convolutional neural networks for volumetric medical image segmentation. in *2016 Fourth International Conference on 3D Vision (3DV)*, 565–571 (IEEE, 2016).
34. Brier, G. W. Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* **78**, 1–3 (1950).
35. Huang, G. *et al.* Snapshot ensembles: Train 1, get m for free. arXiv preprint [arXiv:1704.00109](https://arxiv.org/abs/1704.00109) (2017).
36. Kingma, D. & Ba, J. Adam: A method for stochastic optimization. *ICLR* **5**, (2015).
37. Hendrycks, D. & Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint [arXiv:1610.02136](https://arxiv.org/abs/1610.02136) (2016).
38. Jungo, A. & Reyes, M. Assessing reliability and challenges of uncertainty estimations for medical image segmentation. arXiv preprint [arXiv:1907.03338](https://arxiv.org/abs/1907.03338) (2019).
39. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
40. Bello, G. A. *et al.* Deep-learning cardiac motion analysis for human survival prediction. *Nat. Mach. Intell.* **1**, 95–104 (2019).
41. Meyer, H. V. *et al.* Genetic and functional insights into the fractal structure of the heart. *Nature* 1–6 (2020).
42. Lakshminarayanan, B., Pritzel, A. & Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Adv. Neural Inf. Process. Syst.* 6402–6413 (2017).
43. Kingma, D. P., Salimans, T. & Welling, M. Variational dropout and the local reparameterization trick. *Adv. Neural Inf. Process. Syst.* 2575–2583 (2015).
44. Kendall, A. & Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision?. *Adv. Neural Inf. Process. Syst.* 5574–5584 (2017).
45. Tanno, R. *et al.* Uncertainty quantification in deep learning for safer neuroimage enhancement. arXiv preprint [arXiv:1907.13418](https://arxiv.org/abs/1907.13418) (2019).
46. Paszke, A. *et al.* Automatic differentiation in PyTorch. in *NIPS Autodiff Workshop* (2017).

Acknowledgements

This study was performed within the DLMedIA program (P15-26) funded by Dutch Technology Foundation with participation of PIE Medical Imaging.

Author contributions

J.S., B.D.V. and I.I. designed the concept of the study. J.S. conducted the experiments. J.S., B.D.V. and I.I. wrote the manuscript. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to J.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020