


RESEARCH

Open Access



Automated quality control in image segmentation: application to the UK Biobank cardiovascular magnetic resonance imaging study

Robert Robinson^{1*} , Vanya V. Valindria¹, Wenjia Bai¹, Ozan Oktay¹, Bernhard Kainz¹, Hideaki Suzuki², Mihir M. Sanghvi^{4,5}, Nay Aung^{4,5}, José Miguel Paiva⁴, Filip Zemrak^{4,5}, Kenneth Fung^{4,5}, Elena Lukaschuk⁶, Aaron M. Lee^{4,5}, Valentina Carapella⁶, Young Jin Kim^{6,7}, Stefan K. Piechnik⁶, Stefan Neubauer⁶, Steffen E. Petersen^{4,5}, Chris Page³, Paul M. Matthews^{2,8}, Daniel Rueckert¹ and Ben Glocker¹

Abstract

Background: The trend towards large-scale studies including population imaging poses new challenges in terms of quality control (QC). This is a particular issue when automatic processing tools such as image segmentation methods are employed to derive quantitative measures or biomarkers for further analyses. Manual inspection and visual QC of each segmentation result is not feasible at large scale. However, it is important to be able to automatically detect when a segmentation method fails in order to avoid inclusion of wrong measurements into subsequent analyses which could otherwise lead to incorrect conclusions.

Methods: To overcome this challenge, we explore an approach for predicting segmentation quality based on Reverse Classification Accuracy, which enables us to discriminate between successful and failed segmentations on a per-cases basis. We validate this approach on a new, large-scale manually-annotated set of 4800 cardiovascular magnetic resonance (CMR) scans. We then apply our method to a large cohort of 7250 CMR on which we have performed manual QC.

Results: We report results used for predicting segmentation quality metrics including Dice Similarity Coefficient (DSC) and surface-distance measures. As initial validation, we present data for 400 scans demonstrating 99% accuracy for classifying low and high quality segmentations using the predicted DSC scores. As further validation we show high correlation between real and predicted scores and 95% classification accuracy on 4800 scans for which manual segmentations were available. We mimic real-world application of the method on 7250 CMR where we show good agreement between predicted quality metrics and manual visual QC scores.

Conclusions: We show that Reverse classification accuracy has the potential for accurate and fully automatic segmentation QC on a per-case basis in the context of large-scale population imaging as in the UK Biobank Imaging Study.

Keywords: Automatic quality control, Population imaging, Segmentation

*Correspondence: r.robinson16@imperial.ac.uk

¹Biomedical Image Analysis Group, Department of Computing, Imperial College London, Queen's Gate, SW7 2AZ London, UK

Full list of author information is available at the end of the article



Background

Biomedical image data are increasingly processed with automated image analysis pipelines which employ a variety of tools to extract clinically useful information. It is important to understand the limitations of such pipelines and assess the quality of the results being reported. This is a particular issue when we consider large-scale population imaging databases comprising thousands of images such as the UK Biobank (UKBB) Imaging Study [1]. There are often many modules in automated pipelines [2] where each may contribute to inaccuracies in the final output and reduce the overall quality of the analysis, e.g. intensity normalisation, segmentation, registration and feature extraction. On a large scale, it is infeasible to perform a manual, visual inspection of all outputs, and even more difficult to perform quality control (QC) within the pipeline itself. We break down this challenge and focus on the automated QC of image segmentation.

Image segmentation is the process of partitioning an image into several parts where each of these parts is a collection of pixels (or voxels) corresponding to a particular structure. The purpose of segmentation is to derive quantitative measures of these structures, e.g. calculating ventricular volume or vessel thickness. Automated segmentation is desired to reduce workload for this tedious, time-consuming and error prone task. A number of these methods have been developed, ranging from basic region-growing techniques and graph cuts to more advanced algorithms involving machine learning [3] and, more recently, Deep Learning in the form of Convolutional Neural Networks (CNNs) [4].

Segmentation performance is traditionally evaluated on a labelled validation dataset, which is a subset of the dataset that the algorithm does not see during training. This evaluation is done using a series of metrics to compare the predicted segmentation and a reference 'ground truth' (GT). Popular metrics include volumetric overlap [5], surface distances or other statistical measures [6]. Due to the lack of actual GT, manual expert annotations are used as reference, despite inter- and intra-rater variability. Once a segmentation method is deployed in clinical practice no such quantitative evaluation can be carried out routinely.

Evaluating the average performance of an algorithm on validation data is arguably less important than being able to assess the quality on a per-case basis, and it is crucial to identify cases where the segmentation has failed. We show that we can effectively predict the per-case quality of automated segmentations of 3D cardiovascular magnetic resonance (CMR) from the UKBB which enables fully automated QC in large-scale population studies and clinical practice.

In this article we will first present related work that attempts to address the problem of automated QC at

large-scale. Our method and datasets are then described in detail before we present our results and discuss their implications.

Related work

Despite its practical importance, there is relatively little work on automatically predicting performance of image analysis methods. Much of the prior work on automated quality control has focused on the quality of images themselves. This focus on image quality assessment (IQA) is also true in the medical-imaging community [7, 8]. In the context of image segmentation, there exist only a few methods outlined here.

Algorithms often rely on 'labels' to support their training. In our case, each label would indicate the quality of each segmentation, either by categorical label, e.g. 0 for 'poor' and 1 for 'good', or by continuous value such as a Dice Similarity Coefficient (DSC). In cases where such labelled data is scarce, Reverse Validation [9] and Reverse Testing [10] use labels generated by one model, trained on a subset of available data, to train another model which is evaluated on the remaining data. This is effectively cross-validation where the amount of labelled data is limited. In Reverse Testing, 'some rules' are created to assess the performance of and rank the different models. In our context, this would involve creating a segmentation quality model from a subset of CMR scans, and their corresponding segmentations, which can then be tested on the remaining images. Different models would be created and tested in order to choose the best model. The difficulty in these methods is that we require all of the scans to be accurately segmented in order to train, and to evaluate, a good model. That is, we need a large, fully-annotated training dataset which is often not available in our field. Additionally, Reverse Validation and Reverse Testing do not allow us to identify individual cases where a segmentation may have failed; instead they focus upon the segmentation method as a whole.

In a method proposed by Kohlberger et al., the quality of segmentations is assessed on a per-case basis using machine learning. The group used 42 different hand-crafted statistics about the intensity and appearance of multi-organ computed-tomography (CT) scans to inform their model. Whilst this method achieved good performance metrics and an accuracy of around 85%, it requires a lot of training data with both good and bad segmentations which is non-trivial to obtain.

In this work, we adopt the recently-proposed approach of Reverse Classification Accuracy (RCA) [11]. Unlike Reverse Validation and Reverse Testing, RCA can accurately predict the quality of a segmentation on a case-by-case basis only requiring a relatively small set of accurately segmented reference images. In RCA, the predicted segmentation being assessed is used to create a small model

to re-segment the reference images for which segmentations are available. If at least one image in the reference set is re-segmented well, the predicted segmentation, that we wish to assess, must have been of good quality. We employ RCA to perform segmentation quality analysis on a per-case basis while only requiring a small set of reference images and segmentations.

Methods and data

Our purpose is to have a system that is able to predict the per-case quality of a segmentation produced by any algorithm deployed in clinical practice. We want our method to not only give us a prediction of the quality of the segmentation, but to be able to identify if that segmentation has failed. To this end, we employ RCA which will give a prediction about the quality of individual segmentations.

Reverse classification accuracy

In RCA the idea is to build a model, also known as an ‘RCA classifier’, solely using one test image and its predicted segmentation which acts as pseudo ground truth. This classifier is then evaluated on a reference dataset for which segmentations are available. There are two possible outcomes to this procedure:

- **Case 1:** assuming that the predicted segmentation is of *good quality*, the created model should be able to segment at least one of the reference images with high accuracy. This is likely to be a reference image which is similar to the test image.
- **Case 2:** if none of the reference images are segmented successfully, then the predicted segmentation is likely to be of *poor quality*.

These assumptions are valid if the reference dataset is representative of the test data. This is usually the case in the context of machine learning where the reference data could have been used in the first place to train the automated method for which we want to predict test performance. If the test data were very different, the automated method would in any case not perform well, and RCA scores would reflect this. It is a great advantage that the same reference dataset can be used to train an automated segmentation method, and also afterwards serves as the reference database enabling prediction of performance after deployment of the segmentation method.

The performance of the RCA classifier on the reference set is measured with any chosen quality metric, e.g., the DSC. The highest score among all reference images determines the quality *estimate* for the predicted segmentation obtained for a test image.

The original work on RCA [11] explored a variety of possible classifiers that could be trained on a single test-image and its segmentation including Atlas Forests (AF)

[12] and CNNs. In this context, and throughout this paper, an ‘atlas’ refers to an image-segmentation pair whose segmentation has been verified by a manual annotator. In Valindria’s paper, a simple single-atlas registration classifier outperformed both the AF and CNN approaches in predicting segmentation accuracy. For this reason, we chose to use this simple approach for the model in our work. Registration is the process of aligning two or more images based upon similar content within them, e.g. structures or intensities. Rigid registration restricts the images to move only by linear translations and rotations. More complex non-rigid registration methods exist that allow for differences in scale between the images and for more complex distortions. The single-atlas registration classifier in RCA works by performing non-rigid registration of the test-image to a set of individual reference-images. The resulting transformations are then used to warp the test-segmentation. This yields a set of warped segmentations which are quantitatively compared to the reference segmentations. The overlap between the pairs is calculated as the DSC whilst boundary agreement is computed using surface-distance metrics. The best metric values among the reference set are taken to be the prediction for the quality of the test-segmentation.

We chose to modify the single-atlas registration classifier from that used in Valindria et al.’s proposal of the RCA method [11]. Processing and modifying the test-segmentation is not usually desirable as this may introduce discretization artefacts adding false-positives into the binary labelmap. We choose to perform the single-atlas registration in reverse: we register the reference-images to the test-image and use this transformation to warp the reference segmentations. This results in a set of warped segmentations in the test-image space which are then compared to the test-segmentation. Figure 1 gives an overview of RCA as applied in our study. We now set out our framework more formally.

For the RCA reference images, we use a set $\mathbf{R}_i = \{r_i^1, \dots, r_i^N\}$ of N cardiac atlases with reference segmentations $\mathbf{R}_s = \{r_s^1, \dots, r_s^N\}$. We have a test set $\mathbf{T}_i = \{t_i^1, \dots, t_i^M\}$ of M images with automatically generated predicted segmentations $\mathbf{T}_s = \{t_s^1, \dots, t_s^M\}$ whose quality we would like to assess. If the GT segmentations $\mathbf{T}_{gt} = \{t_{gt}^1, \dots, t_{gt}^M\}$ for \mathbf{T}_i exist, one can evaluate the accuracy of these quality assessments. Using RCA we estimate the quality of those predicted segmentations and compare the estimates to the real quality with respect to the GT.

In the case where $m = k$, we take the k^{th} test image t_i^k and its predicted segmentation t_s^k . To apply RCA, all reference images \mathbf{R}_i are first registered to t_i^k by performing a rigid registration in the form of a centre of mass (CoM) alignment. Initial versions of our work [13] used landmark

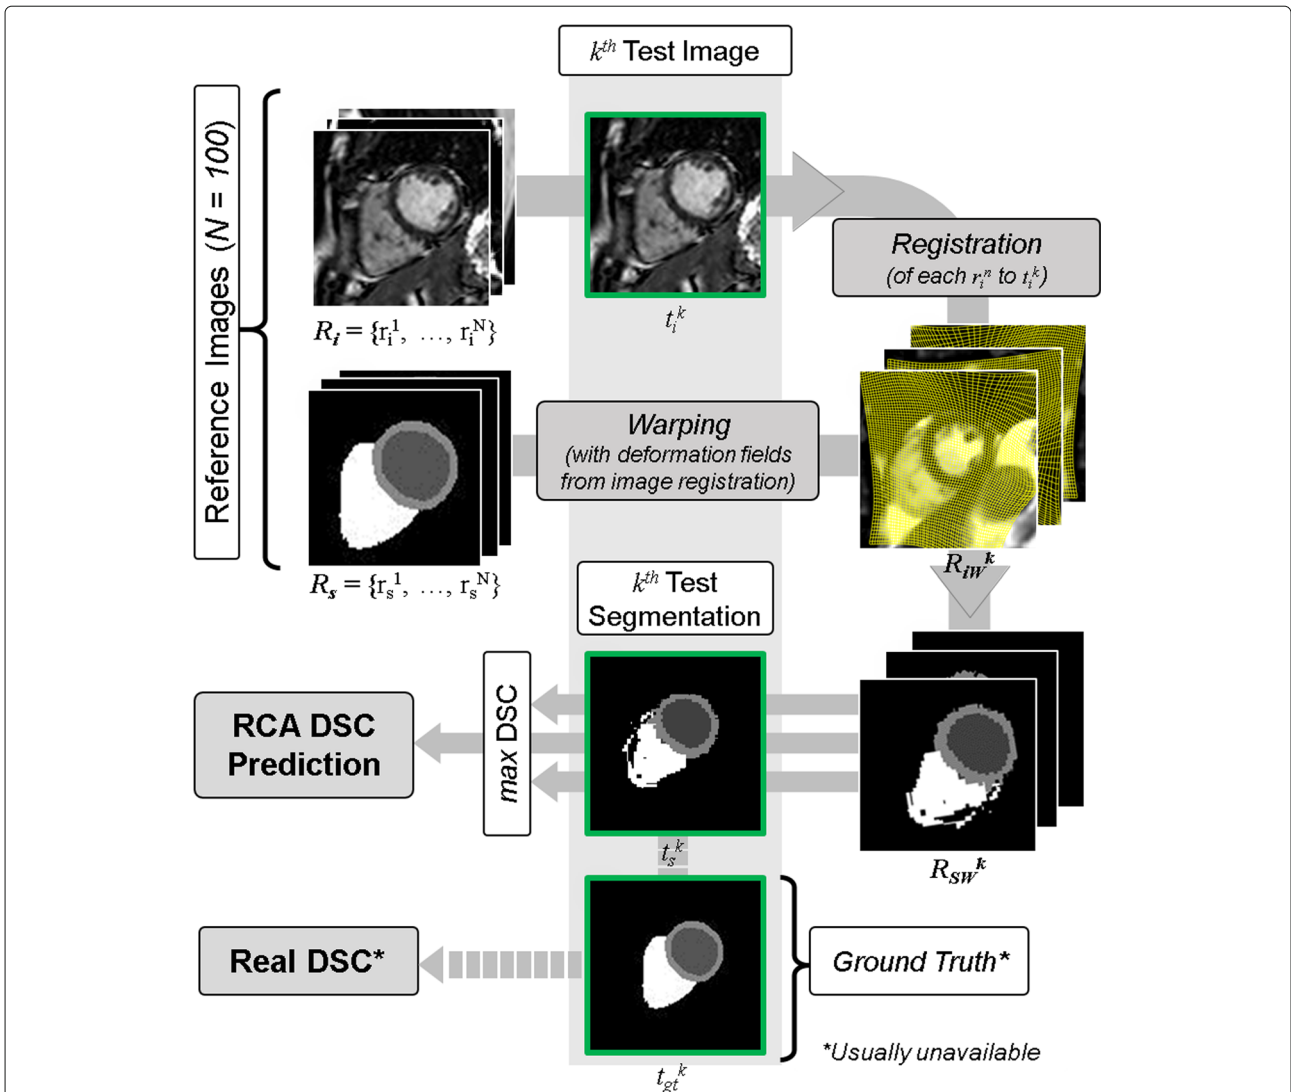


Fig. 1 Reverse Classification Accuracy - Single-atlas Registration Classifier. Reverse Classification Accuracy (RCA), with single-atlas registration classifier, as applied in our study. A set of reference images are first registered to the test-image before the resulting transformations are used to warp the corresponding reference segmentations. Dice Similarity Coefficient (DSC) is calculated between the warped segmentations and the test-segmentation with the maximum DSC taken as a proxy for the accuracy of the test-segmentation. Note that in practice, the ground truth test-segmentation is absent. Images and segmentation annotated as referred to in the text

registration [14] at this stage, but we now opt for CoM alignment to reduce computational cost. We then perform non-linear registration of each aligned reference image in R_i to the test image to get warped reference images R_{iW}^k . The same transformations are used to warp the GT reference segmentations R_s to get the set R_{sW}^k . For each warped segmentation in R_{sW}^k we compare against the predicted segmentation t_s^k by evaluating a set of metrics detailed below. The best value for each metric over all warped reference segmentations is taken to be the *prediction* of segmentation accuracy for t_s^k . In our validation studies, we can compute the *real* metrics by comparing the predicted segmentation t_s^k with its GT t_{gt}^k .

Evaluation of predicted accuracy

The segmentation quality metrics predicted with RCA include the DSC, mean surface distance (MSD), root-mean-square surface distance (RMSD) and Hausdorff distance (HD). For two segmentations, **A** and **B**, DSC is a measure of overlap given by $DSC = 2|A \cap B| / (|A| + |B|)$. The surface distance between a point a on the surface of **A** and the surface of **B** is given by the minimum of the euclidean norm $\min_{b \in B} \|a - b\|_2$ for all points b in the surface of **B**. The total surface distance is the sum of the surface distances for all points in **A**. We don't assume symmetry in these calculations, so the surface distance is also calculated from **B** to **A**. By taking the

mean over all points we get the MSD. RMSD is calculated by taking the square of surface distances, averaging and taking the square root. Finally, the HD is taken to be the maximum surface distance.

For each test image, we report the evaluation metrics for each class label: left ventricular (LV) cavity, LV myocardium (LVM) and right ventricular (RV) cavity (RVC). We incorporate the voxels of the papillary muscles into the LV cavity class. The RV myocardium is difficult to segment because it is thin, therefore it is seldom seen in short-axis (SAX) CMR segmentations and not considered in this paper. For each evaluation metric (DSC and surface distances), we could report two difference average values: either a whole-heart average by combining all class labels into a single ‘whole-heart’ (WH) class or, second, by taking the mean across the individual class scores. The WH-class average is usually higher because a voxel attributed to an incorrect class will reduce the mean calculated across the classes, but will actually be considered correct in the single WH-class case.

Experimental setup

We perform three investigations in this work which are summarised in Table 1: A) an initial small-scale validation study on 400 test contours of 80 images from an internal cardiac atlas dataset; B) a large-scale validation study on another 4805 UKBB images with manual ground truth and C) a real-world application to a large set of 7250 UKBB 3D CMR segmentations.

Reference dataset, $N = 100$

The reference image set is the same in all of our studies. We use 100 2D-stack SAX end-diastolic (ED) CMR scans that were automatically segmented and validated by expert clinicians at Hammersmith Hospital, London. Note that the reference set is distinct from all other datasets used. Compared with data from the UKBB, the reference set are of higher in-plane resolution at 1.25×1.25 mm and have a smaller slice thickness of 2 mm. These images are not used for any purpose other than for this reference set. When choosing a reference set, one should ensure that it is representative of the dataset on which it

Table 1 A summary of the experiments performed in this study

Experiment	Dataset	Size	GT	Seg. Method
A	Hammersmith	100	Yes	RF
B	UKBB-2964	4805	Yes	RF and CNN
C	UKBB-18545	7250	No	Multi-Atlas

Experiment A uses data from an internal dataset which is segmented with a multi-atlas segmentation approach and manually validated by experts at Hammersmith Hospital, London. These manual validations are counted as ‘ground truth’ (GT) and 100 of them are taken for the reference set used in all experiments. UKBB datasets are shown with their application numbers. In experiment C we segment with both random forests (RF) and a convolutional neural network (CNN). In C the CNN from Bai [4] is used

is being used i.e. it should be of the same domain (SAX CMR in this case) and large enough to capture some variability across the dataset. A reference set that is too small may underestimate the RCA prediction. Though we argue that this may be better than overestimating the quality of a segmentation. Conversely, too large a reference set will cause a significant lengthening of RCA execution time. We have explored the effect of the RCA reference set size on the prediction accuracy as part of our evaluation which we present in our [Discussion](#).

Experiment A: Initial validation study, $N = 400$

Data: We validate RCA on predicting cardiac image segmentation quality using 100 manually verified image-segmentation pairs (different from the reference dataset). Each atlas contains a SAX ED 3D (2D-stack) CMR and its manual segmentation. The images have a pixel-resolution of $1.25 \times 1.25 \times 2.0$ mm and span $256 \times 256 \times 56$ voxels. Each manual segmentation identifies voxels belonging to the LV cavity, LV myocardium and RV cavity separating the heart from the background class.

For validation, we generate automatic segmentations of our atlases with varying quality. We employ Random Forests with $T = 500$ trees and a maximum depth of $D = 40$ trained on the same set of 100 cardiac atlases used for testing RCA in this experiment. Random Forests allow us to produce a variety of test segmentations with intentionally degraded segmentation quality by limiting the depth of the trees during test time. We obtain 4 sets of 100 segmentations by using depths of 5, 20, 30 and 40. Thus, a total of 400 segmentations are used in our initial validation study.

Evaluation: We perform RCA on all 400 segmentations to yield predictions of segmentation quality. The manual segmentations allow us to evaluate the real metrics for each automated segmentation. We compare these to the quality predicted by RCA. To identify individual cases where segmentation has failed, we implement a simple classification strategy similar to that in Valindria’s work [11]. We consider a 2-group binary classification where DSC scores in the range [0.0 0.7] are considered ‘poor’ and in the range [0.7 1.0] are ‘good’. These boundaries are somewhat arbitrary and would be adjusted for a particular use-case. Other strategies could be employed on a task-specific basis, e.g. formulation as outlier detection with further statistical measures. The thresholding approach allows us to calculate true (TPR) and false (FPR) positive rates for our method as well as an overall accuracy from the confusion matrix.

Experiment B: Large-scale validation on manually-segmented UKBB Data, $N = 4805$

In this experiment we demonstrate that RCA is robust for employment in large-scale studies, and indeed produces

accurate predictions of segmentation quality on an individual basis. As part of our collaboration under UK Biobank Application 2964 we have access to 4805 CMR images with manually drawn contours.

Data: In the context of UKBB data, 3D means a stack of 2D acquisitions with slice-thickness of 8.0 mm and slice-gap of 2 mm [15]. The CMR scans have in-plane resolution of 1.83×1.83 mm and span around 192×208 pixels per slice. The number of slices per scan varies between 4-14 with the majority (89%) having 9-12 slices.

Petersen and colleagues [16, 17] manually segmented all slices of each 3D CMR scan available under the data access application. Several annotators were employed following a standard operating procedure to generate almost 5000 high-quality segmentations. With these manual segmentations acting as GT we directly compare predicted segmentation scores with real scores at large scale.

We use the same Random Forests trained in Experiment A to perform automated segmentations at various tree depths chosen randomly across the 4805 scans yielding segmentations of varying quality. In addition to our Random Forests segmentations, we evaluate RCA with 900 segmentations generated with a recent deep learning based approach. As part of the UKBB Application 2964, Bai et al. [4] have trained a CNN on 3900 manually segmented images. The remaining 900 were then automatically segmented using the trained network. The results of Bai et al.'s CNN approach reflect the state-of-the-art in automated 3D CMR segmentation with an accuracy matching the performance human experts [4].

Evaluation: We perform RCA on all 4805 Random Forests segmentations to yield predictions of segmentation quality. We also perform RCA separately on the 900 CNN segmentations produced by a state-of-the-art deep learning approach. With the availability of GT manual segmentations, we can evaluate this experiment in the same way as Experiment A.

Experiment C: Automatic quality control in the UKBB imaging study, $N = 7250$

Having evaluated RCA in the previous two experiments, this experiment mimics how our approach would behave in a real-world application where the GT is unavailable. We apply RCA to segmentations of CMR images from the UKBB.

Data: In total, 7250 CMR images were available to us through the UKBB resource. Each image has been automatically segmented using a multi-atlas segmentation approach [18]. As part of a genome-wide association study (GWAS), each automatic segmentation has been checked manually to confirm segmentation quality. As there is no GT segmentation, we rely on manual QC scores for these segmentations assessed by a clinical expert. The manual QC is based only on visual inspection of the basal, mid

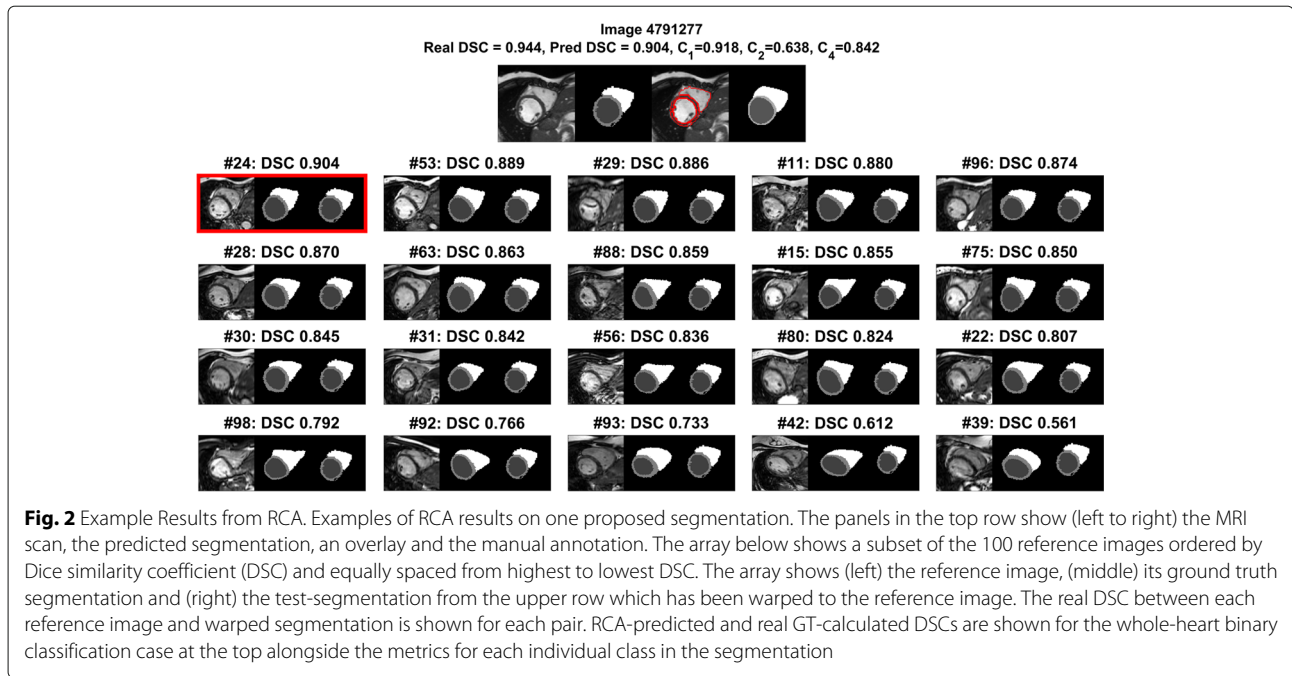
and apical layers. For each layer a score between 0 and 2 is assigned based on the quality of only the LV myocardium segmentation. The total QC score is thus between 0 and 6, where a 6 would be considered as a highly accurate segmentation. Scores for individual layers were not recorded. Where the UKBB images had a poor field-of-view (FOV), the segmentations were immediately discarded for use in the GWAS study: we have given these images a score of -1. For the GWAS study, poor FOV meant any image in which the entire heart was not visible. We expect that despite the poor FOV of these images, the segmentations themselves may still be of good quality as the algorithms can still see most of the heart. Out of the 7250 segmented images, 152 have a bad FOV (QC = -1) and 42 have an obviously poor segmentation (QC = 0). There are 2, 14, 44, 300, 2866 and 3830 images having QC scores 1 to 6 respectively. This investigation explored how well RCA-based quality predictions correlate with those manual QC scores.

Evaluation: We perform RCA on all 7250 segmentations to yield predictions of segmentation quality for the LVM. With the absence of GT segmentations, we are unable to perform the same evaluation as in Experiments A and B. In this case, we determine the correlation between the predicted scores from RCA (for LV myocardium) and the manual QC scores. A visual inspection of individual cases is also performed at quality categories.

Results

Here we present results from our three investigations: (A) the initial small-scale validation study; (B) application to a large set of UKBB CMR with visual QC scores; and (C) a further large-scale validation study on UKBB with manual expert segmentations.

Quantitative results for the experiments are presented in each section. Figure 2 demonstrates additional qualitative inspection that is performed on a per-case basis during RCA. The top row of Fig. 2 shows the mid-ventricular slice of an ED CMR scan and a Random Forest-generated segmentation which is under test. An overlay of the two are also shown alongside the manual reference segmentation which is not available in practice. Below this, an array of further panels is shown. Each of these panels presents one of the 100 reference images used, its corresponding reference segmentation and the result of warping the segmentation-under-test (top-panel, second image) to this reference image. The calculated DSC between the reference image's GT and the warped segmentation is displayed above each panel. The array shows the reference image with the highest (top-left) and the lowest (bottom-right) calculated DSC with the remaining panels showing DSCs that are uniformly spaced amongst the remaining 98 reference images. We can see in this example that there is a large range of predicted DSC values, but



only the maximum prediction, selected in red, is used as the prediction of segmentation quality. For the example in Fig. 2 we show a ‘good’ quality segmentation-under-test for which we are predicting a DSC of 0.904 using RCA. The real DSC between the segmentation-under-test and the GT manual segmentation is 0.944. Note that in this case these values are calculated for the ‘whole-heart’

Table 2 Initial reverse classification accuracy validation on 400 random forest segmentations

Class	Acc.	TPR	FPR	MAE			
				DSC	MSD	RMS	HD
				<i>mm</i>			
LVC	0.973	0.977	0.036				
	0.980	0.975	0.019	0.020	4.104	5.593	14.15
LVM	0.815	0.947	0.215				
	0.990	0.987	0.008	0.044	3.756	4.741	13.08
RVC	0.985	0.923	0.012				
	0.943	0.914	0.047	0.030	4.104	5.022	16.63
Av.	0.924	0.949	0.089				
	0.971	0.959	0.025	0.031	3.988	5.119	14.62
WH	0.988	0.979	0.000				
	0.948	0.886	0.047	0.029	4.445	5.504	15.11

Classes are LV Cavity (LVC), LV Myocardium (LVM), RV Cavity (RVC), An average over the classes (Av.) and a binary segmentation of the whole heart (WH). First row for each class shows the binary classification accuracy for ‘poor’ and ‘good’ segmentations in the Dice Similarity Coefficient (DSC) ranges [0.0 0.7] and [0.7 1.0] respectively. Second row for each class shows the binary classification accuracy for ‘poor’ and ‘good’ segmentations in the Mean Surface Distance (MSD) ranges [$> 2.0\text{mm}$] and [0.0mm 2.0mm] respectively. True-positive and false-positive rates are also shown. We report mean absolute errors (MAE) on the predictions of DSC and additional surface-distance metrics: root-mean-squared surface distance (RMS) and Hausdorff distance (HD)

where individual class labels are merged into one. These values are shown above the top panel along with the DSC calculated on a per-class basis.

For considerations of space, we do not show more visual examples but note that a visualisation as in Fig. 2 could be produced on a per-case basis in a deployed system aiding interpretability and visual means for manual validation by human experts.

(A) Initial validation study

A summary of the results is shown in Table 2. We observe low mean absolute error (MAE) across all evaluation metrics and all class labels. The scatter plots in Fig. 3 on real and predicted scores illustrate the very good performance of RCA in predicting segmentation quality scores. We also find that from the 400 test segmentations, RCA is able to classify ‘good’ ($DSC \in [0.7 1.0]$) and ‘poor’ ($DSC \in [0.1 0.7]$) segmentations with an accuracy of 99%. From 171 poor segmentations at this threshold, 166 could be correctly identified by RCA, i.e. 97.1%. 100% of good-quality segmentations were correctly labelled. Additionally, we find binary classification accuracy of 95% when applying a threshold of 2.0 mm on the MSD. From 365 poor segmentations at this threshold, 348 could be correctly identified by RCA, i.e. 95.3%. Similarly, 31 from 35 (88.6%) good-quality segmentations were correctly labelled. For all evaluation metrics, there is a strong, positive linear relationship between predicted and real values with $r \in [0.95 0.99]$ and $p < 0.0001$. Further analysis of our data shows increasing absolute error in each metric as the real score gets worse, e.g. the error for MSD increases with increasing surface distance. This correlates larger

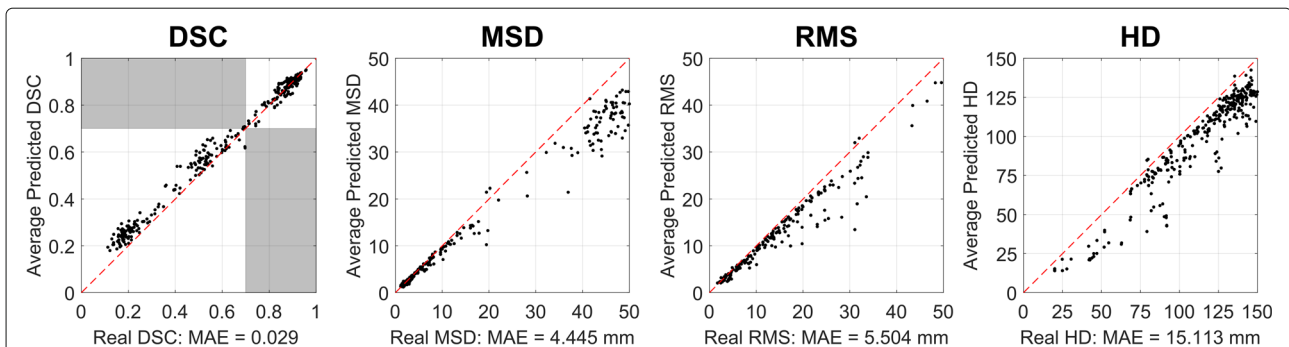


Fig. 3 RCA Validation on 400 cardiac MRI. 400 cardiac MRI segmentations were generated with a Random Forest classifier. 500 trees and depths in the range [5, 40] were used to simulate various degrees of segmentation quality. RCA with single-atlas classifier was used to predict the Dice Similarity Coefficient (DSC), mean surface distance (MSD), root mean-squared surface distance (RMS) and Hausdorff distance (HD). Ground truth for the scans is known so real metrics are also calculated. All calculations on the whole-heart binary classification task. We report low mean absolute error (MAE) for all metrics and 99% binary classification accuracy (TPR = 0.98, FPR = 0.00) with a DSC threshold of 0.70. High accuracy for individual segmentation classes. Absolute error for each image is shown for each metric. We note increasing error with decreasing quality of segmentation based on the real metric score

MAE with lower segmentation quality. In addition, when we consider only those segmentations where the real metric is 30 or less, the MAE drop significantly to 0.65, 1.71 and 6.78 mm for MSD, RMS and HD respectively. We are not concerned with greater errors for poor segmentations as they are still likely to be identified by RCA as having failed.

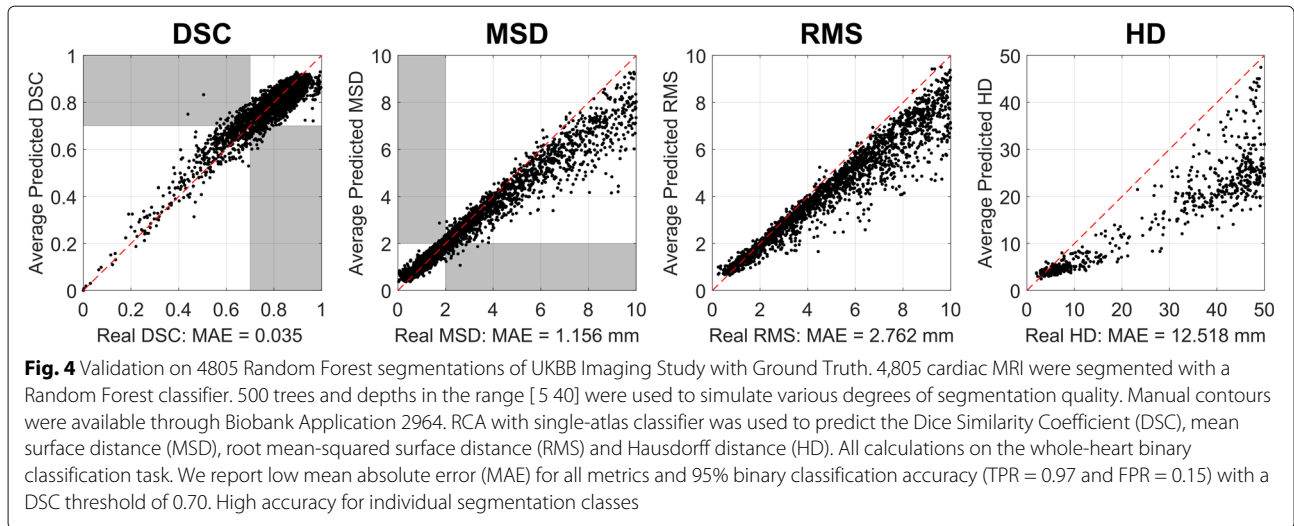
Table 3 Analysis of 4800 Random Forest segmentations with available ground truth

Class	Acc.	TPR	FPR	MAE			
				DSC	MSD <i>mm</i>	RMS <i>mm</i>	HD <i>mm</i>
LVC	0.968 0.975	0.997 0.962	0.330 0.011	0.042	0.906	2.514	11.09
LVM	0.454 0.972	0.956 0.962	0.571 0.012	0.125	0.963	2.141	11.83
RVC	0.868 0.969	0.957 0.977	0.352 0.040	0.057	1.140	2.790	15.23
Av.	0.763 0.972	0.970 0.967	0.418 0.032	0.075	1.003	2.482	12.72
WH	0.954 0.978	0.966 0.984	0.148 0.027	0.035	1.156	2.762	12.52

4800 RF segmentation at various depths [5 40] and 500 trees. Manual contours were available through Biobank Application 2964. Classes are LV Cavity (LVC), LV Myocardium (LVM), RV Cavity (RVC), an average over the classes (Av.) and a binary segmentation of the whole heart (WH). First row for each class shows the binary classification accuracy for ‘poor’ and ‘good’ segmentations in the Dice Similarity Coefficient (DSC) ranges [0.0 0.7] and [0.7 1.0] respectively. Second row for each class shows the binary classification accuracy for ‘poor’ and ‘good’ segmentations in the Mean Surface Distance (MSD) ranges [$> 2.0\text{mm}$] and [0.0mm 2.0mm] respectively. True-positive and false-positive rates are also shown. We report mean absolute errors (MAE) on the predictions of DSC and additional surface-distance metrics: root-mean-squared surface distance (RMS) and Hausdorff distance (HD)

(B) Large-scale validation with manual GT on UKBB

Results for the Random Forest segmentations are shown in Table 3. We report 95% binary classification accuracy with a DSC threshold of 0.7 and low MAE on the DSC. From 589 poor segmentations at this threshold, 443 could be correctly identified by RCA, i.e. 75.2%. Similarly, 4139 from 4216 (98.2%) good-quality segmentations were correctly labelled. Additionally, we find binary classification accuracy of 98% when applying a threshold of 2.0 mm on the MSD. From 2497 poor segmentations at this threshold, 2429 could be correctly identified by RCA, i.e. 97.3%. Similarly, 2270 from 2308 (98.3%) good-quality segmentations were correctly labelled. The TPR are high across the classes, this shows RCA is able to correctly and consistently identify ‘good’ quality segmentations. MSD-based FPR are shown to be lower than those based on DSC, this would indicate that MSD is more discriminative for ‘poor’ quality segmentations and does not misclassify them so much as DSC. We identify only two instances where RCA predictions do not conform to the overall trend and predict much higher than the real DSC. On inspection, we find that the GT of these segmentations were missing mid-slices causing the real DSC to drop. These points can be seen in the upper-left-hand quadrant on Fig. 4. The figure also shows that, over all metrics, there is high correlation between predicted and real quality metrics. This is very much comparable to the results from our initial validation study (A) in Fig. 3. The strong relationship between the predicted quality metrics from RCA and the equivalent scores calculated with respect to the manual segmentations demonstrates concretely that RCA is capable of correctly identifying, on a case-by-case basis, segmentations of poor quality in large-scale imaging studies.



On the CNN segmentations, we report 99.8% accuracy in binary classification for the whole-heart class. With a DSC threshold set at 0.7, RCA correctly identified 898 from 900 good-quality segmentations with 2 false-negatives. A visualization of this can be seen in the top panel of Fig. 5 where the predicted and real DSC can be seen clustered in the high-quality corner of each metric’s plot (upper-right for DSC and lower-left for surface-distance metrics). This reflects the high quality segmentations of the deep learning approach which

have been correctly identified as such using RCA. Table 4 shows the detailed statistics for this experiment.

We note that the individual class accuracy for the LV myocardium is lower in the CNN case when DSC is used at the quality metric. We show the results for this class in the bottom panel of Fig. 5. Segmentors can have difficulty with this class due to its more complex shape. From the plotted points we see all cases fall into a similar cluster to the average WH case, but the RCA score under-predicts the real DSC. This exemplifies a task-specific setting for

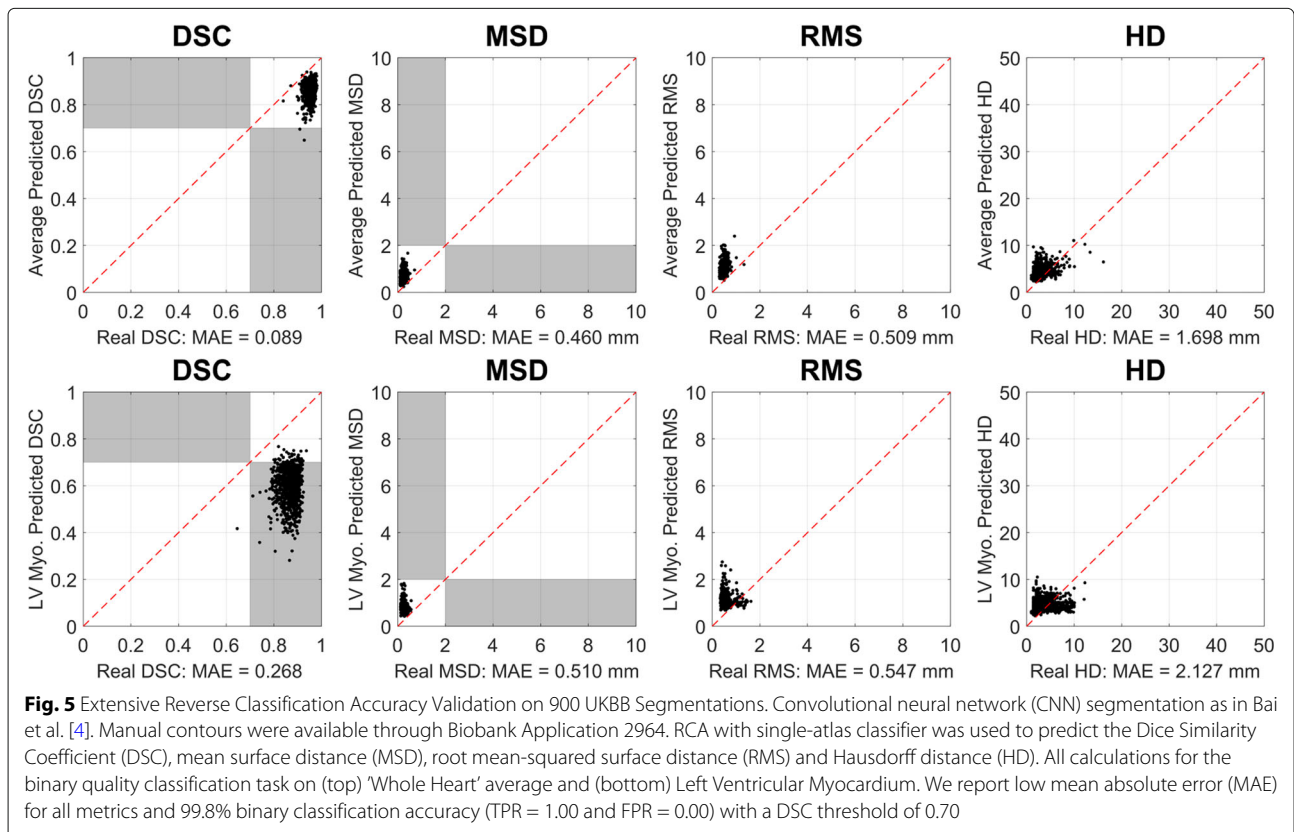


Table 4 Analysis of 900 CNN segmentations with available ground truth

Class	Acc.	TPR	FPR	MAE			
				DSC	MSD	RMS	HD
					<i>mm</i>	<i>mm</i>	<i>mm</i>
LVC	0.998	1.000	0.000	0.082	0.386	0.442	1.344
	1.000	1.000	0.000				
LVM	0.051	1.000	0.001	0.268	0.510	0.547	2.127
	1.000	1.000	0.000				
RVC	0.901	1.000	0.033	0.146	0.588	0.656	2.086
	0.997	0.997	0.000				
Av.	0.650	1.000	0.011	0.165	0.495	0.548	1.852
	0.999	0.999	0.000				
WH	0.998	1.000	0.000	0.089	0.460	0.509	1.698
	1.000	1.000	0.000				

CNN segmentations as in Bai et al. [4]. Manual contours were available through Biobank Application 2964. Classes are LV Cavity (LVC), LV Myocardium (LVM), RV Cavity (RVC), an average over the classes (Av.) and a binary segmentation of the whole heart (WH). First row for each class shows the binary classification accuracy for 'poor' and 'good' segmentations in the Dice Similarity Coefficient (DSC) ranges [0.0 0.7] and [0.7 1.0] respectively. Second row for each class shows the binary classification accuracy for 'poor' and 'good' segmentations in the Mean Surface Distance (MSD) ranges [$>$ 2.0mm] and [0.0mm 2.0mm] respectively. True-positive and false-positive rates are also shown. We report mean absolute errors (MAE) on the predictions of DSC and additional surface-distance metrics: root-mean-squared surface distance (RMS) and Hausdorff distance (HD)

how RCA would be used in practice. In this case one cannot rely only on DSC to predict the quality of the segmentation, so MSD could provide a more appropriate quality prediction.

(C) Quality control on 7250 UK Biobank images

Figure 6 shows the relationship between manual QC scores and the predicted DSC, MSD, RMS and HD obtained from RCA. Note, these predictions are for the LV myocardium and not the overall segmentation as this class was the focus of the manual QC procedure. Manual QC was not performed for the other classes.

Figure 6 also shows a sample of segmentations with manual QC scores of 0, 1, 5 and 6 for the LV myocardium. With a score of 0, 'A' must have a 'poor' quality segmentation of LV myocardium at the basal, apical and mid slices. Example 'B' shows relatively low surface-distance metrics and a low DSC, we see this visually as the boundary of the myocardium is in the expected region, but is incomplete in all slices. This segmentation has been given a score of 1 because the mid-slice is well segmented while the rest is not; which is correctly identified by RCA. In example 'C', the segmentation of the LV myocardium is clearly not good with respect to the image, yet it has been given a manual QC score of 5. Again, RCA is able to pick up such outliers by predicting a lower DSC. The final example 'D' displays an agreement between the high predicted

DSC from RCA and the high manual QC score. These examples demonstrate RCA's ability to correctly identify both good and poor quality segmentations when performing assessments over an entire 3D segmentation. It also demonstrates the limitations of manual QC and the success of RCA in identifying segmentation failure on a per-case basis.

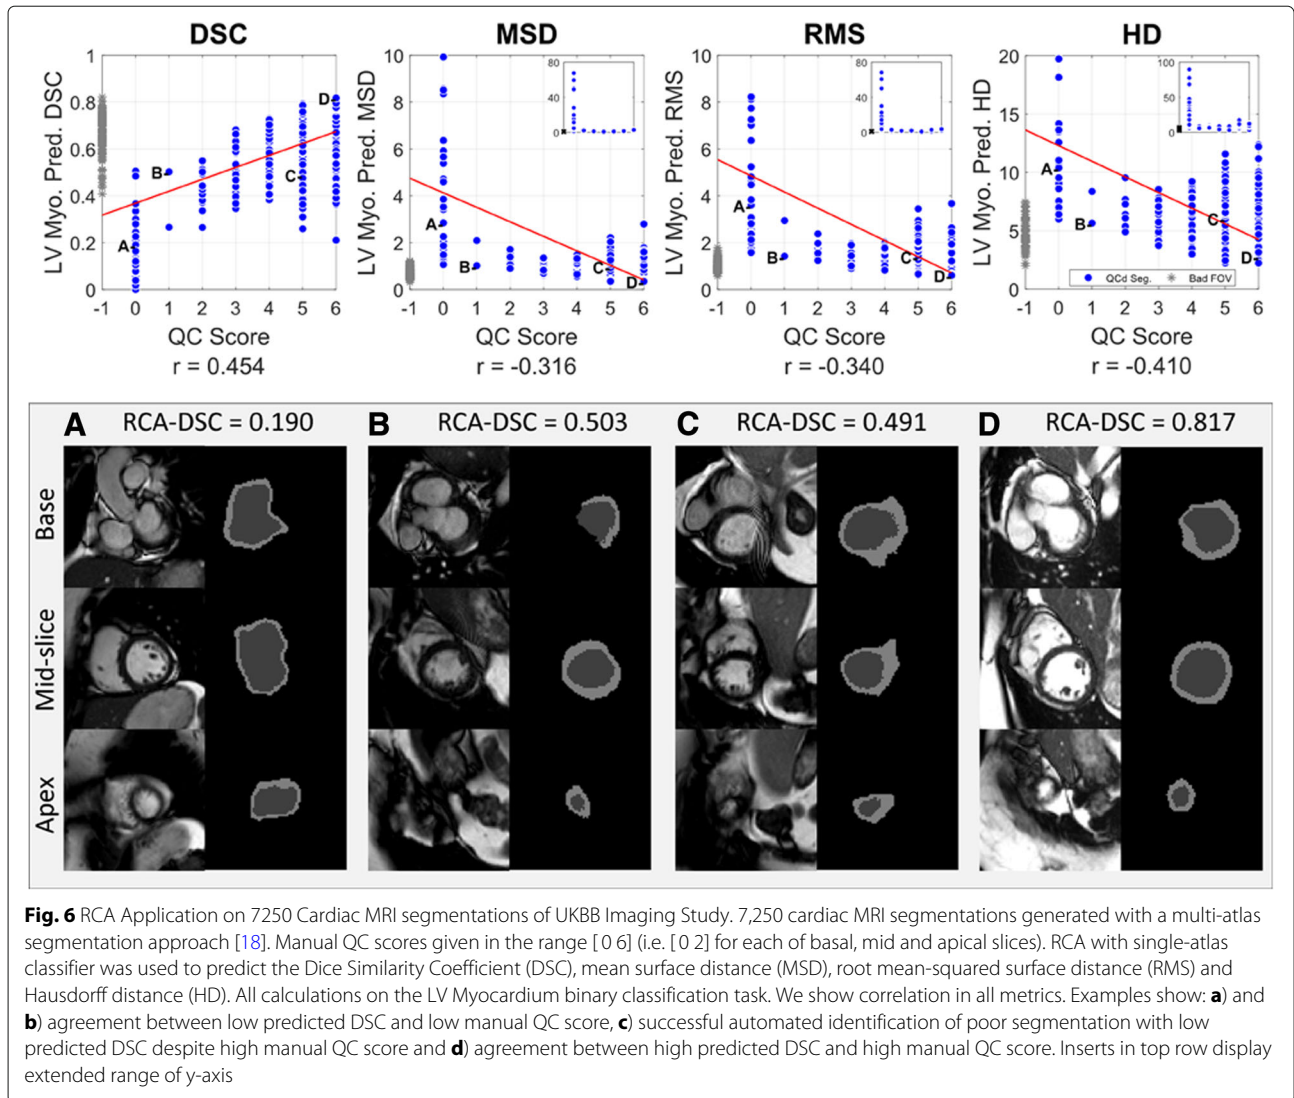
Creating a set of manual QC scores for over 7200 images is a laborious task but it has provided worthwhile evidence for the utility of RCA on large-scale studies. It is clear, however, that the 3-layer inspection approach with a single-rater has limitations. First, inspecting all layers would be preferable, but it highly time-consuming and, second, with multiple raters, averaging or majority voting could be employed to reduce human error.

We should note that RCA is unlikely to mimic the exact visual manual QC process, and it should not, as it naturally provides a different, more comprehensive, assessment of segmentation quality. The manual QC is a rather crude assessment of segmentation quality, as such, we did not perform a direct, quantitative comparison using the visual QC categories but rather wanted to demonstrate that there is a general correlation between manual QC and the predicted RCA score.

Discussion

We have shown that it is possible to assess the quality of individual CMR segmentations at large-scale and in the absence of ground truth. Previous approaches have primarily focused on evaluating overall, average performance of segmentation methods or required large sets of pre-annotated data of good and bad quality segmentations for training a classifier. Our method is well suited for use in image-analysis pipelines and clinical workflows where the quality of segmentations should be assessed on a per-case basis. We have also shown that RCA can provide predictions on a per-class basis. Note that our manually labelled dataset did not include the RV myocardium as a label and therefore has been omitted from our study.

The RCA validation process was carried out on 8-core Intel i7 3.6 GHz machines. The whole process for a single test segmentation - including 100 reference image registrations, warping 100 reference segmentations and metric evaluations - took on average 11 min, making it suitable for background processing in large-scale studies and clinical practice. However, this is a limitation as the runtime per case currently does not allow immediate feedback and prohibits applications with real-time constraints. For example, one could envision a process where CMR scans are immediately segmented after acquisition, and feedback on the quality would be required while the patient is still in the scanner. For this, the computation time of RCA would need to be reduced possibly through an automatic selection of a subset of reference images.



We report preliminary results for using a deep learning approach to speed up the process in [19]. With a real-time RCA framework, the method could be used to identify challenging cases for CNN-based segmentors where the RCA feedback could be used to improve the segmentation algorithm.

As noted earlier, using a subset of the reference set could help to optimize the run-time of RCA predictions. To better understand the effect of reference set size on prediction accuracy, we have performed an empirical evaluation using the data from Experiment B. We took the 4805 automated segmentations and their manual GT and performed RCA using randomly selected subsets of the 100 image-segmentation pairs from the full reference set. Five different randomly selected sets of sizes 10, 15, 25, 35, 50, 65 and 75 were created and used for obtaining RCA predictions on the 4805 images. Figure 7 shows the mean accuracy computed across the 5 runs for each

reference set size. Error bars indicate the highest and lowest accuracy achieved across the five runs. Accuracy is computed using the same DSC threshold of 0.7 as used in Experiment B. The figure shows that the mean accuracy increases with increasing number of reference images. The error bars in Fig. 7 show a decrease in size with increasing size of the reference set. As the reference set grows in size, a greater variability in the images is captured that allows the RCA process to become more accurate. Noteworthy, even with small reference sets of about 20 images high accuracy of more than 90% is obtained.

Although RCA can give a good indication of the real DSC score for an individual segmentation, an accurate one-to-one mapping between the predicted and real DSC has not been achieved. However, we have shown that the method will confidently differentiate between ‘good’ and ‘poor’ quality segmentations based on an application specific threshold. The threshold could be chosen depending

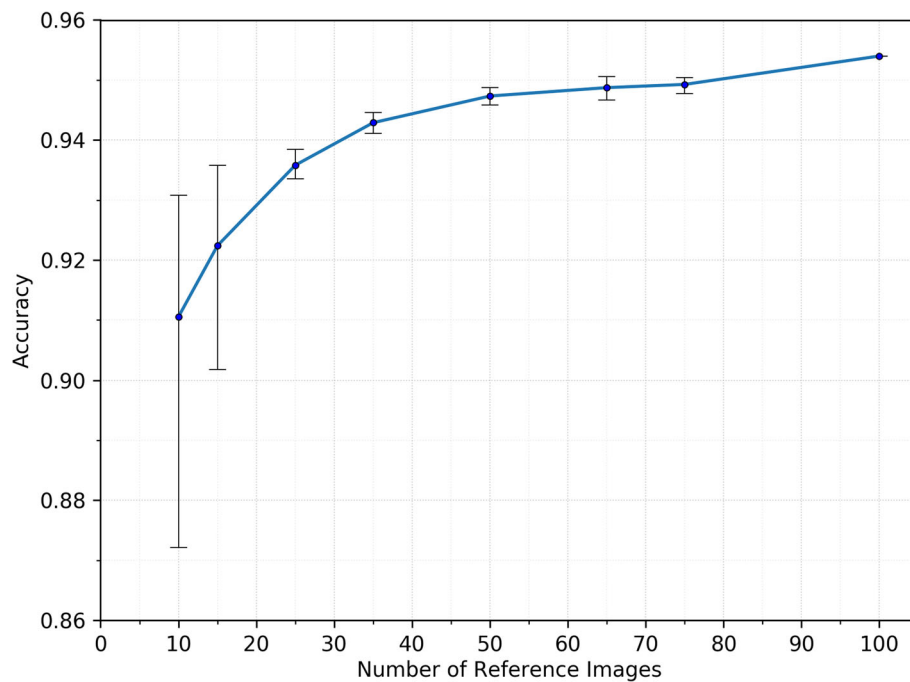


Fig. 7 Investigating the Effect of Reference Set Size on Prediction Accuracy. 4,805 automated segmentations from Experiment B were processed with Reverse Classification Accuracy (RCA) using differing numbers of reference images. Random subsets of 10, 15, 35, 50, 65 and 75 reference images were taken from the full set of 100 available reference images. Five random runs were performed to obtain error bars for each setting. Average prediction accuracy increases with increasing number of reference images and the variance between runs also decreases

on the application's requirements for what qualifies as a 'good' segmentation. Failed segmentations could be re-segmented with different parameters, regenerated with alternative methods, discarded from further analyses or, more likely, sent to a user for manual inspection. Additionally, whilst RCA has been shown to be robust to cardiovascular topology it would need to be re-evaluated for use in other anatomical regions.

Conclusion

RCA had previously been shown to effectively predict the quality of whole-body multi-organ segmentations. We have successfully validated the RCA framework on 3D CMR, demonstrating the robustness of the methodology to different anatomy. RCA has been successful in identifying poor-quality image segmentations with measurements of DSC, MSD, RMS and HD and has shown excellent MAE against all metrics. RCA has also been successful in producing a comparable outcome to a manual quality control procedure on a large database of 7250 images from the UKBB. We have shown further success in accurately predicting quality metrics on 4805 segmentations from Petersen et al., for which manual segmentations were available for evaluation. Predicting segmentation accuracy in the absence of ground truth is a step towards fully automated QC in image analysis pipelines.

Our contributions to the field are three-fold: 1) a thorough validation of RCA for the application of CMR segmentation QC. Our results indicate highly accurate predictions of segmentation quality across various metrics; 2) a feasibility study of using RCA for automatic QC in large-scale studies. RCA predictions correlate with a set of manual QC scores and enable outlier detection in a large set of 7250 CMR cases, and 3) a large-scale validation on 4800 CMR images from the UKBB. Furthermore, we have done this without the need for a large, labelled dataset and we can predict segmentation quality on a per-case basis.

Abbreviations

3D: 3-dimensional; AF: Atlas forest; CMR: Cardiovascular magnetic resonance; CNN: Convolutional neural network; CoM: Center of mass; CT: Computed tomography; DSC: Dice similarity coefficient; ED: End-diastole; FOV: Field of view; FPR: False-positive rate; GT: Ground truth; GWAS: Genome-wide association study; HD: Hausdorff distance; IQA: Image quality assessment; LV: Left ventricle; LVC: Left ventricular cavity; LVM: Left ventricular mass; MAE: Mean absolute error; MSD: Mean surface distance; QC: Quality control; RCA: Reverse classification accuracy; RF: Random forest; RMS: Root-mean-squared surface distance; RMSD: Root mean-square distance; RV: Right ventricle; RVC: Right ventricular cavity; SAX: Short axis; TPR: True-positive rate; UKBB: UK Biobank; WH: Whole heart

Acknowledgements

This work was carried out under UKBB Applications 18545 and 2964. The authors wish to thank all UKBB participants and staff.

Funding

RR is funded by both the King's College London & Imperial College London EPSRC Centre for Doctoral Training in Medical Imaging (EP/L015226/1) and

GlaxoSmithKline; VV by Indonesia Endowment for Education (LPDP) Indonesian Presidential PhD Scholarship and HS by Research Fellowship from Uehara Memorial Foundation. This work was also supported by the following institutions: KF is supported by The Medical College of Saint Bartholomew's Hospital Trust, an independent registered charity that promotes and advances medical and dental education and research at Barts and The London School of Medicine and Dentistry. AL and SEP acknowledge support from the NIHR Barts Biomedical Research Centre and from the "SmartHeart" EPSRC program grant (EP/P001009/1). SN and SKP are supported by the Oxford NIHR Biomedical Research Centre and the Oxford British Heart Foundation Centre of Research Excellence. This project was enabled through access to the MRC eMedLab Medical Bioinformatics infrastructure, supported by the Medical Research Council (grant number MR/L016311/1). NA is supported by a Wellcome Trust Research Training Fellowship (203553/Z/Z). The authors SEP, SN and SKP acknowledge the British Heart Foundation (BHF) for funding the manual analysis to create a cardiovascular magnetic resonance imaging reference standard for the UKBB imaging resource in 5000 CMR scans (PG/14/89/31194). PMM gratefully acknowledges support from the Edmond J. Safra Foundation and Lily Safra, the Imperial College Healthcare Trust Biomedical Research Centre, the EPSRC Centre for Mathematics in Precision Healthcare, the UK Dementia Research Institute and the MRC. BG received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 757173, project MIRA, ERC-2017-STG).

Availability of data and materials

The imaging data and manual annotations were provided by the UKBB Resource under Application Number 2946. Part (B) was conducted with data obtained through Application Number 18545. Researchers can apply to use the UKBB data resource for health-related research in the public interest. Python code for performing RCA as implemented in this work is in the public repository: <https://github.com/mlnotebook/RCA>. Example sets of reference and test-images are provided as links in the repository.

Authors' contributions

RR and BG conceived and designed the study; RR performed implementations, data analysis and write the manuscript. VV developed the original RCA framework. WB provided the automated segmentations that were given QC scores by HS. OO provided landmarks used in the original version of this work. SN, SKP and SEP are overall consortium leads for UKBB access application 2964 and responsible for the conceptualisation of creating a CMR segmentation reference standard. Data curation for application 2964 by SEP, NA, AML and VC and manual contours of 5,000 CMR produced by MMS, NA, JMP, FZ, KF, EL, VC and YJK. PM and DR provided advice in the stage of model development and clinical applications. All authors read and approved the final manuscript.

Ethics approval and consent to participate

The UKBB has approval from the North West Research Ethics Committee (REC reference: 11/NW/0382). All participants have given written informed consent.

Consent for publication

Not applicable.

Competing interests

Steffen E. Petersen provides consultancy to Circle Cardiovascular Imaging Inc. (Calgary, Alberta, Canada). Ben Glocker receives research funding from HeartFlow Inc. (Redwood City, CA, USA).

Publisher's Note

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

Author details

¹Biomedical Image Analysis Group, Department of Computing, Imperial College London, Queen's Gate, SW7 2AZ London, UK. ²Division of Brain Sciences, Dept. of Medicine, Imperial College London, Queen's Gate, SW7 2AZ London, UK. ³GlaxoSmithKline Research and Development, Stockley Park, UB11 1BT Uxbridge, UK. ⁴William Harvey Research Institute, NIHR Barts Biomedical Research Centre, Queen Mary University of London, Charterhouse Square, EC1M 6BQ London, UK. ⁵Barts Heart Centre, Barts Health NHS Trust, West Smithfield, EC1A 7BE London, UK. ⁶Division of Cardiovascular Medicine,

Radcliffe Department of Medicine, University of Oxford, OX3 9DU Oxford, UK. ⁷Department of Radiology, Severance Hospital, Yonsei University College of Medicine, Seoul, South Korea. ⁸UK Dementia Research Institute, Imperial College London, Queen's Drive, SW7 2AZ London, UK.

Received: 7 September 2018 Accepted: 4 February 2019

Published online: 14 March 2019

References

1. Sudlow C, Gallacher J, Allen N, Beral V, Burton P, Danesh J, Downey P, Elliott P, Green J, Landray M, Liu B, Matthews P, Ong G, Pell J, Silman A, Young A, Sprosen T, Peakman T, Collins R. UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. *PLoS Med.* 2015;12(3):1–10. <https://doi.org/10.1371/journal.pmed.1001779>.
2. Shariff A, Kangas J, Coelho LP, Quinn S, Murphy RF. Automated Image Analysis for High-Content Screening and Analysis. *J Biomol Screen.* 2010;15(7):726–34. <https://doi.org/10.1177/1087057110370894>.
3. de Bruijne M. Machine learning approaches in medical image analysis: From detection to diagnosis. *Med Image Anal.* 2016;33:94–97. <https://doi.org/10.1016/j.media.2016.06.032>.
4. Bai W, Sinclair M, Tarroni G, Oktay O, Rajchl M, Vaillant G, Lee AM, Aung N, Lukaschuk E, Sanghvi MM, Zemrak F, Funk K, Paiva JM, Carapella V, Kim YJ, Suzuki H, Kainz B, Matthews PM, Petersen SE, Piechnik SK, Neubauer S, Glocker B, Rueckert D. Human-level cmr image analysis with deep fully convolutional networks. *1710.09289v3*.
5. Crum WR, Camara O, Hill DLG. Generalized overlap measures for evaluation and validation in medical image analysis. *IEEE Trans Med Imaging.* 2006;25(11):1451–61. <https://doi.org/10.1109/TMI.2006.880587>.
6. Taha AA, Hanbury A. Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Med Imaging.* 2015;15:29. <https://doi.org/10.1186/s12880-015-0068-x>.
7. Carapella V, Jiménez-Ruiz E, Lukaschuk E, Aung N, Funk K, Paiva J, Sanghvi M, Neubauer S, Petersen S, Horrocks I, Piechnik S. Towards the Semantic Enrichment of Free-Text Annotation of Image Quality Assessment for UK Biobank Cardiac Cine MRI Scans. In: MICCAI Workshop on Large-scale Annotation of Biomedical Data and Expert Label Synthesis (LABELS). Cham: Springer; 2016. p. 238–48. https://doi.org/10.1007/978-3-319-46976-8_25.
8. Zhang L, Gooya A, Dong B, Hua R, Petersen SE, Medrano-Gracia P, Frangi AF. Automated Quality Assessment of Cardiac MR Images Using Convolutional Neural Networks. In: Tsaftaris SA, Gooya A, Frangi AF, Prince JL, editors. *Medical Image Computing and Computer-Assisted Intervention – SASHIMI 2016. Lecture Notes in Computer Science*, vol. 9968. Cham: Springer; 2016. p. 138–45. https://doi.org/10.1007/978-3-319-46630-9_14.
9. Zhong E, Fan W, Yang Q, Verscheure O, Ren J. Cross Validation Framework to Choose amongst Models and Datasets for Transfer Learning. In: *Lecture Notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol. 6323 LNAI. Springer; 2010. p. 547–62. https://doi.org/10.1007/978-3-642-15939-8_35. http://link.springer.com/10.1007/978-3-642-15939-8_35.
10. Fan W, Davidson I. Reverse testing. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '06*. New York: ACM Press; 2006. p. 147. <https://doi.org/10.1145/1150402.1150422>.
11. Valindria VV, Lavdas I, Bai W, Kamnitsas K, Aboagye EO, Rockall AG, Rueckert D, Glocker B. Reverse Classification Accuracy: Predicting Segmentation Performance in the Absence of Ground Truth. *Trans Med Imaging IEEE.* 2017. <https://doi.org/10.1109/TMI.2017.2665165>.
12. Zikic D, Glocker B, Criminisi A. Encoding atlases by randomized classification forests for efficient multi-atlas label propagation. *Med Image Anal.* 2014;18(8):1262–73. <https://doi.org/10.1016/j.media.2014.06.010>.
13. Robinson R, Valindria VV, Bai W, Suzuki H, Matthews PM, Page C, Rueckert D, Glocker B. Automatic quality control of cardiac mri segmentation in large-scale population imaging. In: *Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, Duchesne S, editors. Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*. Cham: Springer; 2017. p. 720–7.

14. Oktay O, Bai W, Guerrero R, Rajchl M, de Marvao A, O'Regan DP, Cook SA, Heinrich MP, Glocker B, Rueckert D. Stratified decision forests for accurate anatomical landmark localization in cardiac images. *IEEE Trans Med Imaging*. 2017;36(1):332–42. <https://doi.org/10.1109/tmi.2016.2597270>.
15. Petersen SE, Matthews PM, Francis JM, Robson MD, Zemrak F, Boubertakh R, Young AA, Hudson S, Weale P, Garratt S, Collins R, Piechnik S, Neubauer S. UK biobank's cardiovascular magnetic resonance protocol. *J Cardiovasc Magn Reson*. 2016;18(1):8. <https://doi.org/10.1186/s12968-016-0227-4>.
16. Petersen SE, Aung N, Sanghvi MM, Zemrak F, Fung K, Paiva JM, Francis JM, Khanji MY, Lukaschuk E, Lee AM, Carapella V, Kim YJ, Leeson P, Piechnik SK, Neubauer S. Reference ranges for cardiac structure and function using cardiovascular magnetic resonance (CMR) in caucasians from the UK biobank population cohort. *J Cardiovasc Magn Reson*. 2017;19(1):. <https://doi.org/10.1186/s12968-017-0327-9>.
17. Petersen SE, Sanghvi MM, Aung N, Cooper JA, Paiva JM, Zemrak F, Fung K, Lukaschuk E, Lee AM, Carapella V, Kim YJ, Piechnik SK, Neubauer S. The impact of cardiovascular risk factors on cardiac structure and function: Insights from the UK biobank imaging enhancement study. *PLoS ONE*. 2017;12(10):1–14. <https://doi.org/10.1371/journal.pone.0185114>.
18. Bai W, Shi W, O'Regan DP, Tong T, Wang H, Jamil-Copley S, Peters NS, Rueckert D. A Probabilistic Patch-Based Label Fusion Model for Multi-Atlas Segmentation With Registration Refinement: Application to Cardiac MR Images. *IEEE Trans Med Imaging*. 2013;32(7):1302–15. <https://doi.org/10.1109/TMI.2013.2256922>.
19. Robinson R, Oktay O, Bai W, Valindria V, Sanghvi M, Aung N, Paiva J, Zemrak F, Fung K, Lukaschuk E, Lee A, Carapella V, Kim YJ, Kainz B, Piechnik S, Neubauer S, Petersen S, Page C, Rueckert D, Glocker B. Real-time Prediction of Segmentation Quality. 2018. ArXiv e-prints. [1806.06244](https://arxiv.org/abs/1806.06244).

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

