

# Reverse Classification Accuracy: Predicting Segmentation Performance in the Absence of Ground Truth

Vanya V. Valindria,\* Ioannis Lavdas, Wenjia Bai, Konstantinos Kamnitsas, Eric O. Aboagye, Andrea G. Rockall, Daniel Rueckert, and Ben Glocker

**Abstract**—When integrating computational tools, such as automatic segmentation, into clinical practice, it is of utmost importance to be able to assess the level of accuracy on new data and, in particular, to detect when an automatic method fails. However, this is difficult to achieve due to the absence of ground truth. Segmentation accuracy on clinical data might be different from what is found through cross validation, because validation data are often used during incremental method development, which can lead to overfitting and unrealistic performance expectations. Before deployment, performance is quantified using different metrics, for which the predicted segmentation is compared with a reference segmentation, often obtained manually by an expert. But little is known about the real performance after deployment when a reference is unavailable. In this paper, we introduce the concept of *reverse classification accuracy* (RCA) as a framework for predicting the performance of a segmentation method on new data. In RCA, we take the predicted segmentation from a new image to train a reverse classifier, which is evaluated on a set of reference images with available ground truth. The hypothesis is that if the predicted segmentation is of good quality, then the reverse classifier will perform well on at least some of the reference images. We validate our approach on multi-organ segmentation with different classifiers and segmentation methods. Our results indicate that it is indeed possible to predict the quality of individual segmentations, in the absence of ground truth. Thus, RCA is ideal for integration into automatic processing pipelines in clinical routine and as a part of large-scale image analysis studies.

**Index Terms**—Abdominal, classification, image segmentation, machine learning, MRI, performance evaluation.

Manuscript received December 15, 2016; revised January 25, 2017; accepted January 28, 2017. Date of publication April 17, 2017; date of current version July 30, 2017. This work was supported in part by the National Institute for Health Research under EME Project 13/122/01 and in part by the EPSRC First Grant Scheme under Grant EP/N023668/1. The work of V. V. Valindria was supported by the Indonesia Endowment for Education (LPDP)—Indonesian Presidential Ph.D. Scholarship Programme. The work of K. Kamnitsas was supported by the Imperial College President's Ph.D. Scholarship Programme. *Asterisk indicates corresponding author.*

\*V. V. Valindria is with the BioMedIA Group, Department of Computing, Imperial College London, London SW7 2AZ, U.K. (e-mail: v.valindria15@imperial.ac.uk).

I. Lavdas and E. O. Aboagye are with the Comprehensive Cancer Imaging Centre, Hammersmith Hospital, Imperial College London, London W12 0NN, U.K.

W. Bai, K. Kamnitsas, D. Rueckert, and B. Glocker are with the BioMedIA Group, Department of Computing, Imperial College London, London SW7 2AZ, U.K.

A. G. Rockall is with The Royal Marsden NHS Foundation Trust, London SW3 6JJ, U.K.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2017.2665165

## I. INTRODUCTION

SEGMENTATION is an essential component in many image analysis pipelines that aim to extract clinically useful information from medical images to inform clinical decisions in diagnosis, treatment planning, or monitoring of disease progression. A multitude of approaches have been proposed for solving segmentation problems, with popular techniques based on graph cuts [1], multi-atlas label propagation [2], statistical models [3], and supervised classification [4]. Traditionally, performance of a segmentation method is evaluated on an annotated database using various evaluation metrics in a cross-validation setting. These metrics reflect the performance in terms of agreement [5] of a predicted segmentation compared to a reference ‘ground truth’ (GT).<sup>1</sup> Commonly used metrics include Dice’s similarity coefficient (DSC) [6] and other overlap based measures [7], but also metrics based on volume differences, surface distances, and others [8]–[10]. A detailed analysis of common metrics and their suitability for segmentation evaluation can be found in [11].

Once a segmentation method is deployed in routine little is known about its real performance on new data. Due to the absence of GT, it is not possible to assess performance using traditional evaluation measures. However, it is critical to be able to assess the level of accuracy on clinical data [12], and in particular, it is important to detect when an automatic segmentation method fails. Especially when the segmentation is an intermediate step within a larger automated processing pipeline where no visual quality control of the segmentation results is feasible. This is of high importance in large-scale studies such as the UK Biobank Imaging Study [13] where automated methods are applied to large cohorts of several thousand images, and the segmentation is to be used for further statistical population analysis. In this work, we are asking the question whether it is possible to assess segmentation performance and detect failure cases when there is no GT available to compare with. One possible approach to monitor the segmentation performance is to occasionally select a random dataset, obtain a manual expert segmentation and compare it to the automatic one. While this can merely provide a rough estimate about the average performance of the employed segmentation method, in clinical routine we are interested in the per case performance and want to detect

<sup>1</sup>For simplicity, we use the term ground truth to refer to the best known reference, which is typically a manual expert segmentation.

when the automated method fails. The problem is that the performance of a method might be substantially different on clinical data, and is usually lower than what is found through cross-validation on annotated data carried out beforehand due to several reasons. First, the annotated database is normally used during incremental method development for training, model selection and fine tuning of hyper-parameters. This can lead to overfitting [14] which is a potential cause for lower performance on new data. Second, the clinical data might be different due to varying imaging protocols or artefacts caused by pathology. To this end, we propose a general framework for predicting the real performance of deployed segmentation methods on a per case basis in the absence of GT.

### A. Related Work

Retrieving an objective performance evaluation without GT has been an issue in many domains, from remote sensing [15], graphics [16], to marketing strategies [17]. In computer vision, several works evaluate the segmentation performance by looking at contextual properties [18], by separating the perceptual salient structures [19], or by automatically generating semantic GT [20], [21]. However, these methods cannot be applied to a more general task, such as an image with many different class labels to be segmented. An attempt to compute objective metrics, such as precision and recall with missing GT is proposed by [22] but it cannot be used for data sets with partial GT since it applies a probabilistic model under the same assumptions. Another stand-alone method is to consider a meta-evaluation framework, where image features are used in a machine learning setting to provide a ranking of different methods [23]. But this does not allow the estimation of segmentation performance on an individual image level.

Meanwhile, unsupervised methods [24], [25] aim to estimate the segmentation accuracy directly from the images and labelmaps using, for example, information-theoretic and geometrical features. While unsupervised methods can be applied to scenarios where the main purpose of segmentation is to yield visually consistent results that are meaningful to a human observer, the application in medical settings is unclear.

When there are multiple reference segmentations available, a similarity measure index can be obtained by comparing an automatic segmentation with the set of references [26]. In medical imaging, the problem of performance analysis with multiple references which may suffer from intra-rater and inter-rater variability, has been addressed [27], [28]. The STAPLE approach [27] has led to the work of Bouix et al. [29] that proposed techniques for comparing the relative performance of different methods without the need of GT. Here, the different segmentation results are treated as plausible references, thus can be evaluated through STAPLE and the concept of common agreement. Another work by [30] has quantitatively evaluated the performance of several segmentation algorithms by region-correlation matrix. The limitation of this work is that it cannot evaluate the segmentation performance of a particular method on a particular image independently.

Recent work has explored the idea of learning a regressor to directly predict segmentation accuracy from a set of features that are related to various segmentation energy terms [31]. Here, the assumption is that those features are well suited to characterise segmentation quality. In an extension for a security application, the same features as in [31] are extracted and used to learn a generative model of good segmentations that can be used to detect outliers [32]. Similarly, the work of [33] considers training of a classifier that is able to discriminate between consistent and inconsistent segmentations. However, the approaches [31], [33] can only be applied when a training database with good and bad segmentations is available from which a mapping from features to segmentation accuracy can be learned. Examples of bad segmentations can be generated by altering parameters of automatic methods, but it is unclear whether those examples resemble realistic cases of segmentation failure. The generative model approach in [32] is appealing as it only requires a database of good segmentations. However, there is still the difficulty of choosing appropriate thresholds on the probabilities that indicate bad or failed segmentations. Such an approach cannot be used to directly predict segmentation scores such as DSC, but can be useful to inform automatic quality control or to automatically select the best segmentation from a set of candidates.

In the general machine learning domain, the lack-of-label problem has been tackled by exploiting transfer learning [34] using a reverse validation to perform cross-validation when the number of labeled data is limited. The basic idea of reverse validation [34] is based on reverse testing [35], where a new classifier is trained on predictions on the test data and evaluated again on the training data. This idea of reverse testing is closely related to our approach as we will discuss in the following.

### B. Contribution

The main contribution of this paper is the introduction of the concept of *reverse classification accuracy* (RCA) to assess the segmentation quality of an individual image in the absence of GT. RCA can be applied to evaluate the performance of any segmentation method on a per case basis. To this end, a classifier is trained using a single image with its predicted segmentation acting as *pseudo* GT. The resulting *reverse classifier* (or RCA classifier) is then evaluated on images from a reference database for which GT is available. It should be noted that the reference database can be (but does not have to be) the training database that has been used to train, cross-validate and fine-tune the original segmentation method. The assumption is that in machine learning approaches, such a database is usually already available, but it could also be specifically constructed for the purpose of RCA. Our hypothesis is that if the segmentation quality for a new image is high, then the RCA classifier trained on the predicted segmentation used as pseudo GT will perform well at least on some of the images in the reference database, and similarly, if the segmentation quality is poor, the classifier is likely to perform poorly on the reference images. For the segmentations obtained on the reference images through the RCA classifier,

we can quantify the accuracy, e.g., using DSC, since reference GT is available. It is expected that the maximum DSC score over all reference images correlates well with the real DSC that one would get on the new image if GT were available. While the idea of RCA is similar to reverse validation [34] and reverse testing [35], the important difference is that in our approach we train a reverse classifier on every single instance while the approaches in [34] and [35] train single classifiers over the whole test set and its predictions jointly to find out what the best original predictor is. RCA has the advantage of allowing to predict the accuracy for each individual case, while at the same time aggregating over such accuracy predictions allows drawing conclusions for the overall performance of a particular segmentation method.

In the following, we will first present the details of RCA and then evaluate its applicability to a multi-organ segmentation task by exploring the prediction quality of different segmentation metrics for different combinations of segmentation methods and RCA classifiers. Our results indicate that, at least to some extent, it is indeed possible to predict the performance level of a segmentation method on each individual case, in the absence of ground truth. Thus, RCA is ideal for integration into automatic processing pipelines in clinical routine and as part of large-scale image analysis studies.

## II. REVERSE CLASSIFICATION ACCURACY

The RCA framework is based on the idea of training reverse classifiers on individual images utilizing their predicted segmentation as pseudo GT. An overview of the RCA framework is shown in Fig. 1. In this work, we employ three different methods for realizing the RCA classifier and evaluate each in different combinations with three state-of-the-art image segmentation methods. Details about the different RCA classifiers are provided in the following.

### A. Learning Reverse Classifiers

Given an image  $I$  and its predicted segmentation  $S_I$ , we aim to learn a function  $f_{I,S_I}(\mathbf{x}) : \mathbb{R}^n \mapsto C$  that acts as a classifier by mapping feature vectors  $\mathbf{x} \in \mathbb{R}^n$  extracted for individual image points to class labels  $c \in C$ . In theory, any classification approach could be utilized within the RCA framework for learning the function  $f_{I,S_I}$ . We experiment with three different methods reflecting state-of-the-art machine learning approaches for voxel-wise classification and atlas-based label propagation.

a) *Atlas forests*: The first approach we consider for learning a RCA classifier is based on the recently introduced concept of Atlas Forests (AFs) [36] which demonstrates the feasibility of using Random Forests (RFs) [37] to encode individual atlases, i.e., images with corresponding segmentations. Random Forests have become popular for general image segmentation tasks as they naturally handle multi-class problems and are computationally efficient. Since they operate as voxel-wise classifiers, they do not (necessarily) require pre-registration of the images neither at training nor testing time. Although in [36] spatial priors have been incorporated by means of registering location probability maps to each atlas and new image, this is not a general requirement for

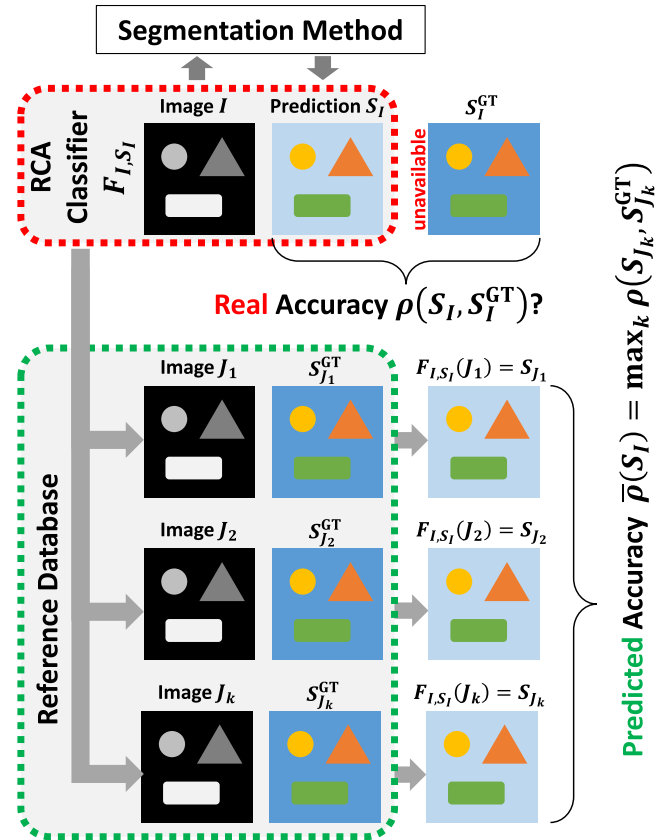


Fig. 1. Overview of the general framework for reverse classification accuracy. The accuracy of the predicted segmentation  $S_I$  of a new image  $I$  is estimated via a RCA classifier trained using  $S_I$  as pseudo ground truth, and applied to a set of reference images  $J_k$  for which ground truth reference segmentations  $S_{J_k}^{GT}$  are available. The best segmentation score computed on the reference images is used as prediction of the real accuracy of segmentation  $S_I$ .

using AFs to encode atlases. In fact, the way we employ AFs within our RCA framework does not require any image registration. The forest-based RCA classifiers in this work are trained all with the same set of parameters of maximum depth 30 and 50 trees. As we follow a very standard approach for RFs, we refer to [36] and [38] for more details. Worth to note that, similar to previous work, we employ simple box features which can be efficiently evaluated using integral images. This has the advantage that feature responses do not need to be precomputed. Instead, we randomly generate a large pool of potential features (typically around 10,000) by drawing randomly values for the feature parameters such as box sizes and offsets from predefined ranges. At each split node we then evaluate on-the-fly a few hundred box features with a brute force search for optimal thresholds over the range of feature responses to greedily find the most discriminative feature/threshold pair. This strategy has proven successful in a number of works using RFs for various tasks.

b) *Deep learning*: We also experiment with convolutional neural networks (CNNs) as RCA classifiers. Here, we utilize DeepMedic,<sup>2</sup> a 3D CNN architecture for automatic segmentation [39]. The architecture is computationally

<sup>2</sup>Code available at <https://github.com/Kamnitsask/deepmedic>



efficient as it can handle large image context by using a dual pathway for multi-scale processing. CNNs have been shown to be able to learn highly complex and discriminative data associations between input data and target output. The architecture of the network is defined by the number of layers and the number of activation functions in each layer. In CNNs, each activation function corresponds to a learned convolutional filter, and each filter produces a feature map (FM) by convolving the outputs of the previous layer. Through the sequential application of many convolutions, highly complex features are learned that are then used to produce voxel-wise predictions at the final, fully-connected layer. CNNs are a type of deep learning approach which normally requires large amounts of training data in order to perform well due to the thousands (or millions) of parameters corresponding to the weights of the filters. To be able to act as a RCA classifier that is trained on a single image, we require a specialised architecture. Here, we reduce the number of FMs in each layer by one third compared to the default setting of DeepMedic. We also cut the feature maps in the last fully connected layers, from 150 to 45. By reducing the feature maps without changing the architecture in terms of number of layers the network preserves its capability to see large image context as the size of the receptive field remains unchanged. With less number of filters, the number of parameters is substantially decreased, which leads to faster computations, but more importantly, reduces overfitting when trained on a single image. Training is performed in a patch-wise manner where the original input image is divided into 3D patches that are then sampled during training using backpropagation and batch normalization. For details about the training procedure and further analysis of DeepMedic we refer to [39].

c) *Atlas-based label propagation*: The third approach we consider is atlas-based label propagation. Label propagation using multiple atlases have been shown to yield state-of-the-art results on many segmentation problems [2]. A common procedure in multi-atlas methods is to use non-rigid registration to align the atlases with the image to be segmented and then perform label fusion strategies to obtain predictions for each image point. Although, multi-atlas methods based on registration are not strictly voxel-wise classifiers as they operate on the whole image during registration, the final stage of label fusion can be considered as a voxel-wise classification step. Here, we make use of an approach that has been originally developed in the context of segmentation of cardiac MRI [40].<sup>3</sup> For the purpose of RCA, however, there is only a single atlas and thus no label fusion is required. Using single atlas label propagation then boils down to making use of an efficient non-rigid registration technique as the one described in [40]. For RCA, the single atlas then corresponds to the image and its predicted segmentation for which we want to estimate the segmentation quality. We use the same configuration for image registration as in [40] and refer to this work for further details.

### B. Predicting Segmentation Accuracy

For the purpose of assessing the quality of an individual segmentation, we train a RCA classifier  $f_{I,S_I}$  on a single image

$I$  that has been segmented by any segmentation method, where  $S_I$  denotes the predicted segmentation that here acts as pseudo GT during classifier training. Our objective is to estimate the quality of  $S_I$  in the absence of GT. To this end, we define the segmentation function  $F_{I,S_I}(J) = S_I$  that applies the trained RCA classifier  $f_{I,S_I}$  to all voxels (or more precisely to the features extracted at each voxel) of another image  $J$  which produces a segmentation  $S_J$ . Assuming that for the image  $J$  a reference GT segmentation  $S_J^{\text{GT}}$  is available, we can now compute any segmentation evaluation metric on the pair  $(S_J, S_J^{\text{GT}})$ . The underlying hypothesis in our RCA framework is that there is a correlation between the values computed on  $(S_J, S_J^{\text{GT}})$  and the values one would get for the pair  $(S_I, S_I^{\text{GT}})$ , where  $S_I^{\text{GT}}$  is the reference GT of image  $I$  which in practice, however, is unavailable.

It is unlikely that this assumption of correlation holds for an arbitrary reference image  $J$ . In fact, the RCA classifier  $f_{I,S_I}$  is assumed to work best on images that are somewhat similar to  $I$ . Therefore, we further assume that a suitable reference database is available that contains multiple segmented images (or atlases)  $\mathcal{T} = \{(J_k, S_{J_k}^{\text{GT}})\}_{k=1}^m$  that capture the expected variability. Such a database is commonly available in the context of machine learning and multi-atlas based segmentation approaches, but could also be generated specifically for the purpose of RCA. If already available, we can re-use existing training databases that might have been previously used during method development and/or cross-validation and parameter tuning. When testing the RCA classifier on all of the available  $m$  reference images, we expect that the RCA classifier performs well on at least some of these, if and only if the predicted segmentation  $S_I$  is of good quality. If  $S_I$  is of bad quality, we expect the RCA classifier to perform poorly on all reference images. This leads to our definition of a proxy measure for predicting the segmentation accuracy as

$$\bar{\rho}(S_I) = \max_{1 \leq k \leq m} \rho(F_{I,S_I}(J_k), S_{J_k}^{\text{GT}}), \quad (1)$$

where  $\rho$  is any evaluation metric, such as DSC, assuming higher values correspond to higher quality segmentations.<sup>4</sup> Here, we only look for the maximum value that is found across all reference images, as this seems to be a good indicator of the quality of the segmentation  $S_I$ . Other statistics could be considered, such as the average of the top three scores, but we found that the maximum score works best as a proxy. Note, that the mean or median scores are not very useful measures as we do not expect the RCA classifier to work well on the majority of the reference images. Afterall, the RCA classifier does overfit to the single image and will not generalize to perform well on dissimilar images. Nonetheless, as we will demonstrate in the experiments,  $\bar{\rho}$  indeed provides accurate estimates for the segmentation quality in a wide range settings

### C. Summary

The following provides a summary of the required steps for using RCA in practice within a processing pipeline for automatic image segmentation. Given an image  $I$  to be segmented:

<sup>4</sup>For metrics where a lower value indicates better quality, such as surface distance, we can simply replace the max with a min operator.

<sup>3</sup>Code available at <https://github.com/baiwenjia/CIMAS>

- 1) Run the automated image segmentation method to obtain predicted segmentation  $S_I$ .
- 2) Train a RCA classifier on image  $I$  and its predicted segmentation  $S_I$  to obtain an image segmenter  $F_{I,S_I}$ .
- 3) Evaluate the RCA classifier on a reference database with images for which GT is available  $\mathcal{T} = \{(J_k, S_{J_k}^{GT})\}_{k=1}^m$  to obtain segmentations  $\forall k F_{S_I}(J_k) = S_{J_k}$ .
- 4) Compute the segmentation quality of  $S_I$  using a proxy measure  $\bar{\rho}(S_I)$  according to Eq. (1).

Depending on the application, a threshold may be defined on  $\bar{\rho}$  to flag up images with poor segmentation quality that need manual inspection, or to automatically identify high quality segmentations suitable for further analysis.

### III. EXPERIMENTAL VALIDATION

In order to test the effectiveness of the RCA framework, we explore a comprehensive multi-organ segmentation task on whole-body MRI. In this application, we evaluate the prediction accuracy of RCA in the context of three different state-of-the-art segmentation methods, a Random Forest approach [4], a deep learning approach using 3D CNNs [39], and a probabilistic multi-atlas label propagation approach [40]. The dataset used to validate our framework is from our MALIBO (MAchine Learning In whole Body Oncology) study. We collected whole-body, multi-sequence MRI (T1w Dixon and T2w images) of 35 healthy volunteers. Detailed manual segmentations of 15 anatomical structures, including abdominal organs (heart, left/right lung, liver, adrenal gland, gall bladder, left/right kidney, spleen, pancreas, bladder) and bones (spine, left/right clavicle, pelvis) have been generated by clinical experts as part of the study. These manual segmentations will serve as GT in the quantitative evaluation.

#### A. Experimental Setting

We use 3-fold cross validation to automatically segment all 525 structures (15 organs  $\times$  35 subjects) with each of the three different segmentation methods, namely Random Forests, CNNs, and Multi-Atlas. In each fold, we use the RCA framework with three different methods for realizing the RCA classifier, namely Atlas Forests, constrained CNNs, and Single-Atlas, as described above. Using the RCA classifiers that are trained on each image for which we want to assess segmentation quality, we obtain segmentations on all reference images which are then compared to their manual reference GT. Since the GT is available for all 35 cases, we can compare the predicted versus the real segmentation accuracy for all cases and all organs under various settings with nine different combinations of segmentation methods and RCA classifiers.

#### B. Quantifying Prediction Accuracy

The Dice's similarity coefficient is the most widely used measure for evaluating segmentation performance,<sup>5</sup> and in our main results we focus on evaluating how well DSC can be predicted using our RCA framework. In order to quantify prediction accuracy, we consider three different measures,

namely the correlation between predicted and real DSC, the mean absolute error (MAE), and a classification accuracy. Arguably, the most important measure for direct evaluation of how well RCA works is the MAE, as it directly tells us how close the predicted DSC is to the real one. Correlation is interesting, as it tells us something about the relation between predicted and real scores. We expect high correlation in order for RCA to be useful, but we might not always have an identity relation, as there could be a bias in the predictions. For example, if the predicted score is consistently lower than the real score, this can still be useful in practice, and will be indicated by high correlation but might not yield low MAEs. In such a case, a calibration might be considered as we will discuss later on. We also explore whether the predictions can be used to categorize segmentations according to their quality. We argue that for many clinical applications it is already of great value to be able to discriminate between good, bad, and possibly medium quality segmentations and that the absolute segmentation scores are of less importance. For proof-of-principle, we consider a three-category classification by grouping segmentations within DSC ranges [0.0, 0.6) for 'bad', [0.6, 0.8) for 'medium', and [0.8, 1.0] for 'good' cases. Note, that those ranges are somewhat arbitrary, in particular, as the quality of absolute DSC values is highly depending on the structure of interest. So in practice, those ranges would need to be adjusted specifically to the application at hand.

#### C. Results for Predicting Dice's Similarity Coefficients

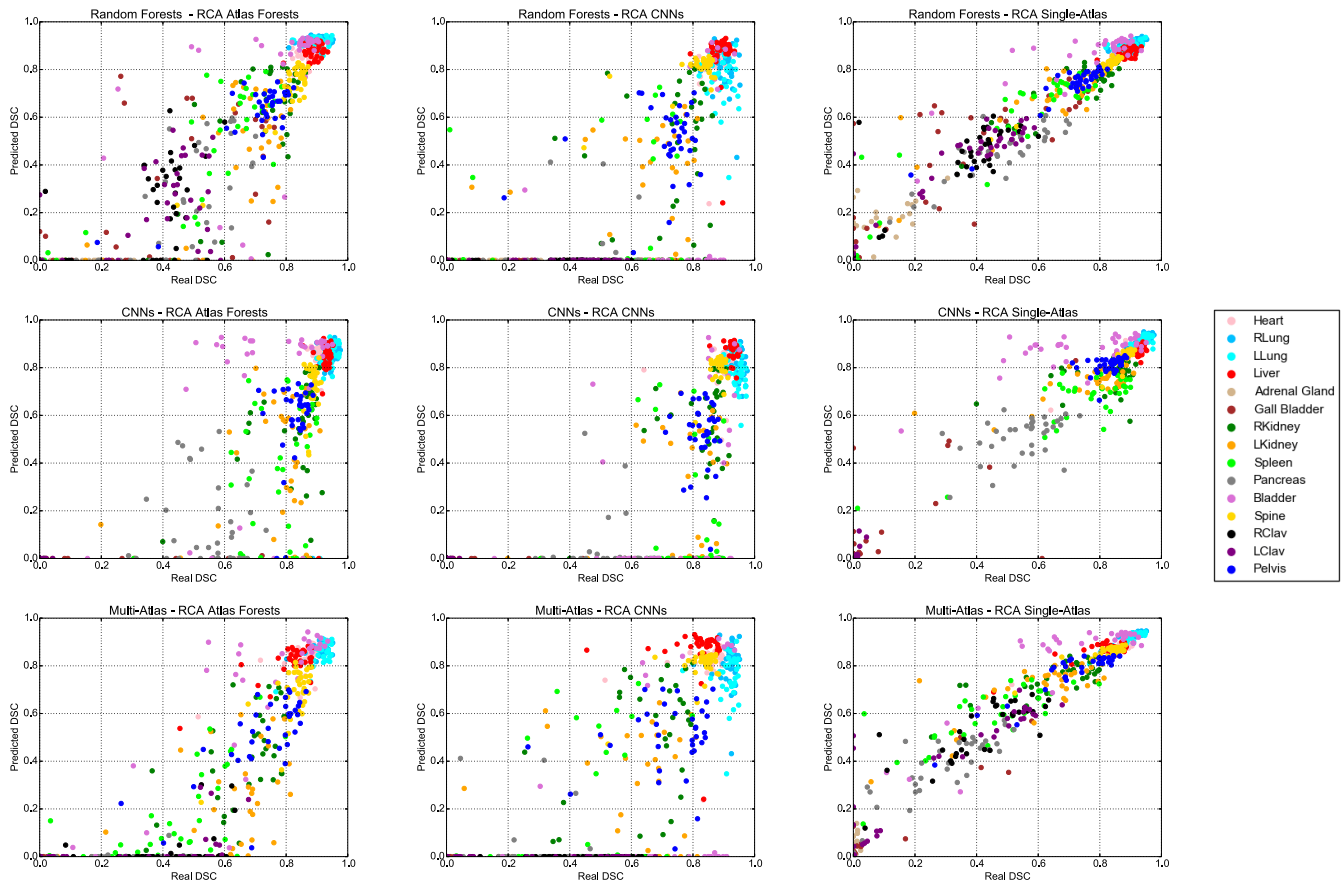
Our main results are summarized in Tab. I where we report the quantitative analysis of the predicted accuracy for nine different settings consisting of three different segmentation methods and three different ways of realizing the RCA classifier. In Fig. 2 we provide the scatter plots of real versus predicted DSC for all nine settings with 525 data points each.

Overall, we observe high correlation between predicted and real DSC for both Atlas Forests and Single-Atlas when used as RCA classifiers, with the Single-Atlas showing correlations above 0.95 for all three segmentation methods. The Single-Atlas approach also yields the lowest MAEs between 0.05 and 0.07, and good 3-category classification accuracies between 81% and 89%. This is visually confirmed by the scatter plots in the right column of Fig. 2 which show good linear relation close to the diagonal between predicted and real scores for most structures in the case where Random Forests or Multi-Atlas are used as the original segmentation method. When using Atlas Forests for RCA, we still observe good correlation but the relationship between predicted and real scores is off-diagonal with larger spread towards lower quality segmentation. The correlation is still good and above 0.82, MAEs are between 0.12 and 0.17 with classification accuracy going down to 0.62%, 0.75% and 0.78% depending on the original segmentation method. For the case of the constrained CNNs, we observe that the prediction quality is lowest confirmed by the scatter plots and all quantitative measures, with correlations below 0.78 and MAEs above 0.2. The constrained CNNs seem to only work for predicting segmentation accuracy in case of major organs such as liver, lungs, and the spine but

<sup>5</sup>Despite some well known shortcomings of DSC as discussed in [11].

**TABLE I**  
 PREDICTING DSC FOR DIFFERENT SEGMENTATION METHODS USING DIFFERENT RCA CLASSIFIERS. BEST RESULTS ARE OBTAINED WITH THE SINGLE-ATLAS APPROACH YIELDING THE LOWEST ERROR BETWEEN REAL AND PREDICTED DSC SCORES AND HIGH CLASSIFICATION ACCURACY

Segmentation Method	RCA Classifier	Correlation		MAE		Accuracy 3-Categories	
		All	No Zeros	All	No Zeros	All	No Zeros
Random Forests	Atlas Forests	0.881	0.867	0.120	0.130	0.783	0.776
CNNs	Atlas Forests	0.828	0.630	0.166	0.245	0.623	0.500
Multi-Atlas	Atlas Forests	0.863	0.877	0.168	0.177	0.749	0.726
Random Forests	Constrained CNNs	0.721	0.718	0.252	0.271	0.653	0.631
CNNs	Constrained CNNs	0.756	0.662	0.225	0.292	0.592	0.472
Multi-Atlas	Constrained CNNs	0.773	0.686	0.209	0.237	0.693	0.642
Random Forests	Single-Atlas	0.955	0.946	0.051	0.052	0.888	0.880
CNNs	Single-Atlas	0.973	0.892	0.052	0.065	0.811	0.756
Multi-Atlas	Single-Atlas	0.962	0.947	0.067	0.072	0.822	0.798

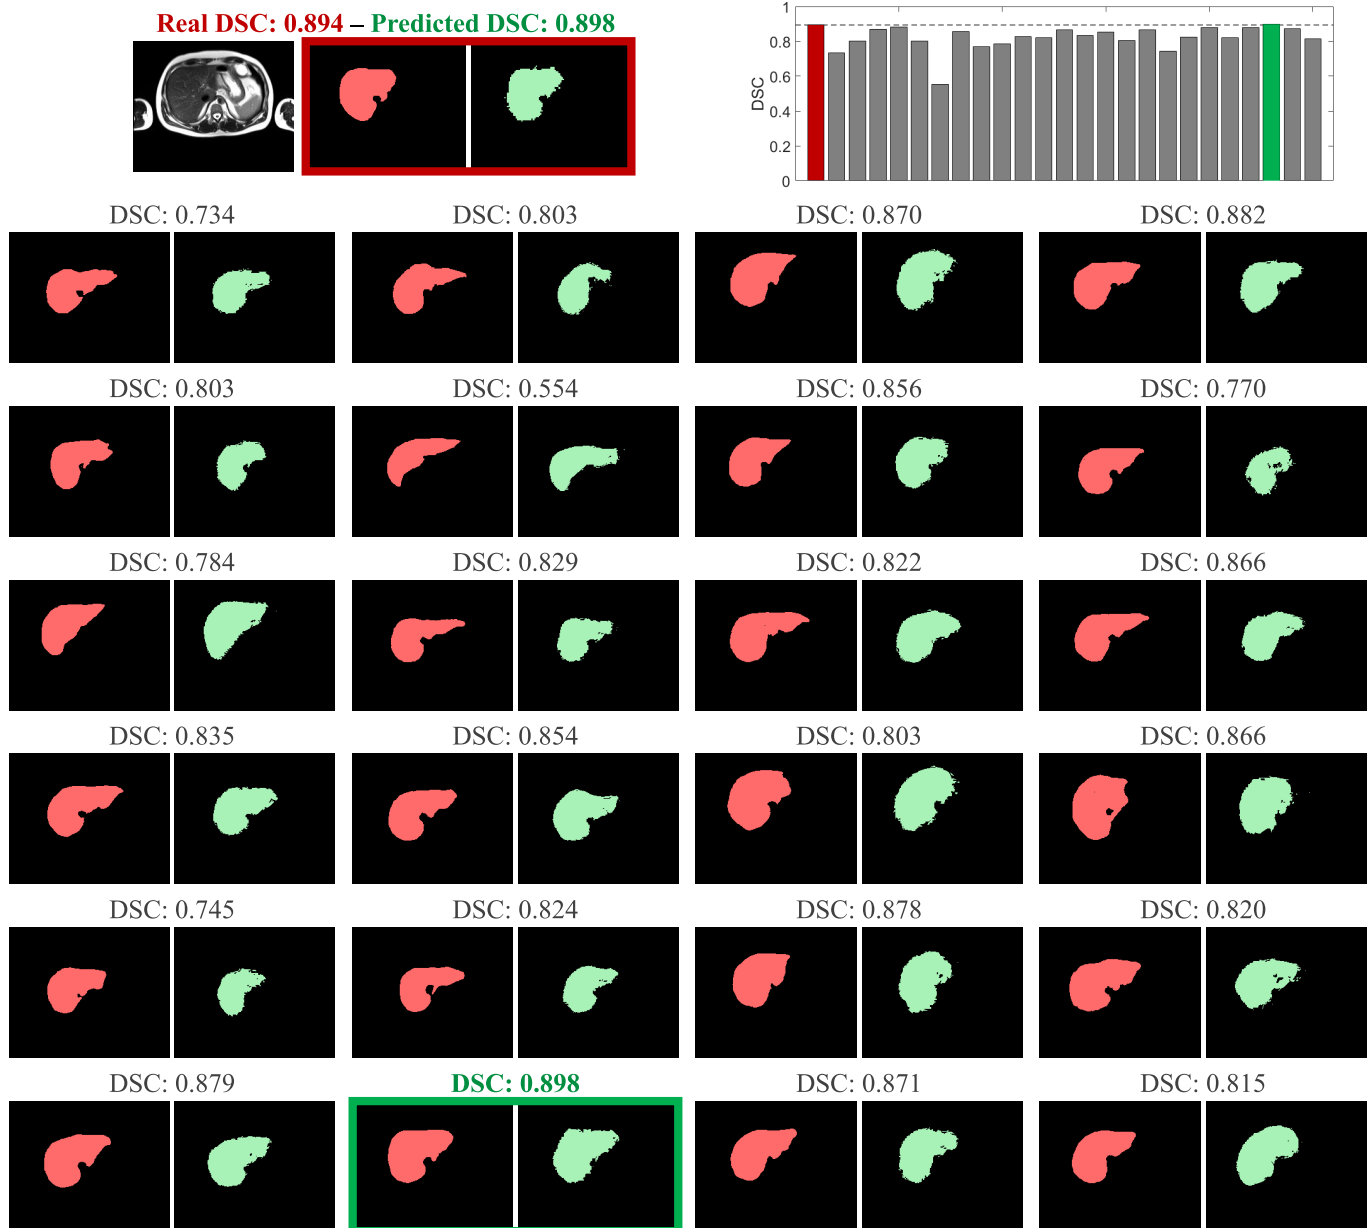


**Fig. 2.** Scatter plots of predicted and real DSC of multiple structures for three different segmentation methods (rows) using three different RCA classifiers (columns). High correlation and low prediction errors are obtained when employing the single-atlas label propagation as RCA classifier (right column). There is also good correlation with predictions in case of Atlas Forests (left) with larger spread towards lower quality segmentations. The constrained CNNs (middle column) are less suitable for RCA which is likely due to the difficulty of training on single images. Both, Atlas Forests and constrained CNNs work best for larger organs such as liver, lungs, and pelvis while leading to many zero predictions for smaller structures such as adrenal gland and clavicles. The single-atlas label propagation makes accurate predictions of segmentation quality across all 15 anatomical structures. A summary of the plots is given in [Tab. I](#).

clearly struggle with smaller structures leading to many zero predictions even when the real DSC is rather high. This is most likely caused by the difficulty of training the CNNs with single images and small structures which does not provide sufficient amounts of training data.

[Figure 3](#) shows an example for predicting the accuracy of a liver segmentation. Next to a slice from a T2w MRI volume we show the GT manual segmentation together with the result

from a Random Forest. Underneath, we show the 24 segmentations obtained on the reference database when using the Single-Atlas RCA approach. The bar plot in the same figure shows the variation of the 24 DSC scores. Similarly, the bar plots in [Fig. 4](#) of two more examples illustrate the distribution of DSC scores when predicting a good quality segmentation on the left, and a poor quality segmentation on the right. The three examples support the hypothesis that selecting the maximum



**Fig. 3.** Visual examples for predicting the segmentation quality of the liver for a new image shown on the top. Its GT segmentation (red colored) is unknown in practice, and we want to estimate the quality of the predicted, automatic segmentation shown on the most right (green colored). By taking the predicted segmentation as pseudo GT and training a RCA classifier we can obtain segmentations on a reference database with 24 images with available GT. The bar plot shows the real DSC in red and the different DSC values obtained for the reference images shown below. The green bar corresponds to the maximum DSC, and is selected as predicted DSC for the new image according to Eq. (1) which matches well the real DSC.

score across the reference database according to Eq. (1) is a good proxy for predicting segmentation quality.

Some of the original segmentation methods have problems segmenting structures such as the adrenal gland and clavicles. The CNNs, in particular, failed to segment adrenal glands in most cases. Because the real DSC for these is zero with no voxels labelled in the segmentation map, the RCA predictions are always correct as there are no labels for the RCA classifier for this structure. In order to investigate the effect of those zero predictions on the quantitative results, we also report in [Tab. I](#) under the columns ‘No Zeros’ the correlations, MAEs and classification accuracies when structures with a real DSC of zero are excluded. We observe that the zero predictions

have mostly an impact on CNNs, either employed as original segmentation method or as RCA classifier. For Atlas Forests and Single-Atlas the effect on the accuracies is very little, confirming that those both are well suited within the RCA framework, independent of the original segmentation method.

#### D. Detecting Segmentation Failure

In clinical routine it is of great importance to be able to detect when an automated method fails. We conducted a dedicated experiment to investigate how well RCA can predict segmentation failure. From the scatter plots in [Fig. 2](#) we can see that all three segmentation methods perform






Fig. 4. The bar plots show two examples for predicting the real DSC (red) in case of a good quality (left) and bad quality segmentation (right) using a database with 24 reference images with available GT. The predicted DSC (green bar) selected according to Eq. (1) matches well the real DSC.

TABLE II  
DETECTING SEGMENTATION FAILURE

RCA Classifier	Correlation	MAE	Accuracy 3-Cat.
Atlas Forests	0.853	0.096	0.884
Constrained CNNs	0.237	0.139	0.924
Single-Atlas	0.875	0.097	0.928

reasonably well on most major organs with no failure cases among structures such as liver, heart, and lungs. In order to further demonstrate that RCA can predict failure cases in these structures, we utilize degraded Random Forests by limiting the tree depth at test time to 8. This leads to much worse segmentation results for most structures which is confirmed in the corresponding scatter plots shown in Fig. 5. Again, we evaluate the performance of the three different RCA classifiers, Atlas Forests, constrained CNNs and Single-Atlas. The results are summarized in Tab. II. The constrained CNNs are again suffering from many zero predictions and less suitable for making accurate predictions. Atlas Forests and Single-Atlas, however, result in high correlations, low MAEs and very good classification accuracies. Low real DSC scores are correctly predicted and failed segmentations are identified. The only exception here is the bladder. This might be explained by the unique appearance of the bladder in the multi-spectral MRI with hyper-intensities in the T2w image, and its largely varying shape between subjects. It appears that even a badly segmented bladder can be sufficient for the RCA classifier to learn its appearance and segment the bladder well on at least one of the reference images. Overall, the experiment suggests that RCA with Atlas Forests and Single-Atlas can be employed in automatic quality control, for example, in large-scale studies where it is important to be able to detect failed segmentations which should be excluded from subsequent analyses.

#### E. Results for Predicting Different Segmentation Metrics

We further explore the ability to predict other evaluation metrics than DSC. We consider the following metrics: Jaccard index (JI), precision (PR), recall (RE), average surface distance (ASD), Hausdorff distance (HD) and relative volume difference (RVD). For this experiment, we use Random Forests as segmentation method, and Atlas Forests for RCA. The results are summarized in Tab. III.

Good correlation is obtained between predicted and real overlap based scores, with low MAEs, and high accuracies. Since Jaccard is directly related to DSC,<sup>6</sup> it is expected that

<sup>6</sup> $JI = DSC / (2 - DSC)$ .

TABLE III  
PREDICTING DIFFERENT SEGMENTATION METRICS

Metric	Correlation	MAE	Accuracy 3-Cat.
DSC [0 - 1]	0.881	0.120	0.783
JI [0 - 1]	0.899	0.110	0.749
PR [0 - 1]	0.572	0.242	0.638
RE [0 - 1]	0.833	0.170	0.663
HD [0-150]	0.189	41.38	0.387
ASD [0-10]	0.022	4.120	0.344
RVD [0-1]	0.474	0.051	0.678

the predictions are of similar quality. Prediction accuracy for precision is lower than for recall. The two metrics capture different parts of segmentation error; under-segmentation is not reflected in precision, while over-segmentation is not captured in recall.<sup>7</sup> Distance based errors are unbounded, so we define thresholds for HD and ASD, and errors above are clipped to the threshold value, which is set to 150mm for HD, and 10mm for ASD. This also allows us to define ranges for the error categorization. For HD, we use the ranges [0, 10], (10, 60], and (60, 150] for good, medium and bad quality segmentations. For ASD we divide the range into [0, 2] for good, (2, 5] for medium, and (5, 10] for bad segmentation quality. Compared to overlap based metrics, the RCA predictions for HD and ASD are not convincing with low correlation, high MAE, and low classification accuracy. RVD is the ratio of the absolute difference between the reference and predicted segmentation volume and the reference volume. Perfect segmentation will result in a value of zero. As RVD is also unbounded, we use a threshold of one to indicate maximum error. The predictions for RVD are good, with high classification accuracy of 0.68%, similar to the overlap based scores. In conclusion, it seems RCA works very well for overlap based measures and for RVD to some extent, while distance based metrics cannot be accurately predicted with the current setting and would require further investigation.

## IV. DISCUSSION AND CONCLUSION

The experimental validation of the RCA framework has shown that it is indeed possible to accurately predict the quality of segmentations in the absence of GT, with some limitations. We have explored different methods for realising the RCA classifier and could demonstrate that Atlas Forests and in particular, Single-Atlas label propagation yield accurate predictions for different segmentation methods. As the RCA

<sup>7</sup> $DSC = 2 \cdot PR \cdot RE / (PR + RE)$ .






Fig. 5. Scatter plots for the experiment for detecting segmentation failure when using degraded Random Forests with limited depth as the segmentation method. Atlas Forests (left) and single-atlas label propagation (right) make highly accurate predictions in the low DSC ranges and thus, are able to correctly detect such failed segmentations, with the exception of the bladder. Constrained CNNs are again less suitable for RCA with many zero predictions.

TABLE IV  
COMPARISON OF PREDICTING DSC WITH AND WITHOUT CALIBRATION VIA REGRESSION

Segmentation Method	RCA Classifier	Correlation		MAE		Accuracy 3-Categories	
		Direct	Calibrated	Direct	Calibrated	Direct	Calibrated
Random Forests	Atlas Forests	0.881	0.833	0.120	0.105	0.783	0.809
CNNs	Atlas Forests	0.828	0.929	0.166	0.079	0.623	0.783
Multi-Atlas	Atlas Forests	0.863	0.939	0.168	0.065	0.749	0.745
Random Forests	Constrained CNNs	0.721	0.826	0.252	0.104	0.653	0.768
CNNs	Constrained CNNs	0.756	0.961	0.225	0.049	0.592	0.830
Multi-Atlas	Constrained CNNs	0.773	0.874	0.209	0.100	0.693	0.787
Random Forests	Single-Atlas	0.955	0.872	0.051	0.089	0.888	0.815
CNNs	Single-Atlas	0.973	0.967	0.052	0.048	0.811	0.811
Multi-Atlas	Single-Atlas	0.962	0.918	0.067	0.080	0.822	0.825

framework is generic, other methods can be considered and it might be necessary to select the most appropriate one for the application at hand. We have also experimented with a constrained CNN trained on single images, which only works well for major organs such as liver, lungs and spine. There might be other more appropriate architectures for the purpose of RCA, which will be explored as part of future work.

An appealing property of the proposed framework is that unlike the supervised methods in [31] and [33] no training data is required that captures examples of good and bad segmentations. Instead, in RCA we simply rely on the availability of a reference database with available GT segmentations. The drawback, however, is that we assume a linear relationship between predicted and real scores which should be close to an identity mapping, something we only found in the case of using Single-Atlas label propagation (cf. right column of Fig. 2). In the case of off-diagonal correlation, as for example found for Atlas Forests, an extension to RCA could be considered where the predictions are calibrated. This, however, requires training data from which a regression function could be learned, similar to [31]. In order to demonstrate the potential of such an approach, we perform a simple experiment on the data that we used for conducting the main evaluation. After obtaining all predicted DSC scores, we run a leave-one-subject-out validation where in each fold we use Random Forest regression to calibrate the predictions. The results are summarized in Tab. IV where we compare the quantitative measures before and after calibration. Both the MAEs and classification accuracies improve significantly for the case of Atlas Forests and constrained CNNs.

For Single-Atlas, however, the results remain similar due to the already close to identity relationship between predicted and real scores before calibration. Calibration, however, comes with a risk of overfitting as the method will learn the relationship on the available training data but might not generalize to new data.

In our experiments we have found that best predictions are obtained for overlap based measures such as DSC and Jaccard Index. Whether those measures are sufficient to fully capture segmentation quality is debatable. Still, DSC is the most widely considered measure and being able to accurately predict DSC in the absence of GT has high practical value. Besides being useful for clinical applications where the goal is to identify failed segmentations after deployment of a segmentation method, we see an important application of RCA in large-scale imaging studies and analyses. In settings where thousands of images are automatically processed for the purpose of deriving population statistics, it is not feasible to employ manual quality control with visually inspection of the segmentation results. Here, RCA can be an effective tool to automatically extract the subset of high quality segmentations which can be used for subsequent analysis. We are currently exploring this in the context of population imaging on the UK Biobank imaging data where image data of more than 10,000 subjects is available which will be subsequently increased to 100,000 over the next couple of years. The UK Biobank data will enable the discovery of imaging biomarkers that correlate with non-imaging information such as lifestyle, demographics, and medical records. In the context of such large scale analysis, automatic quality control is a necessity and we

believe the RCA framework makes an important contribution in this emerging area of biomedical research. In future work, we will further explore the use of RCA for other image analysis and segmentation tasks. To facilitate the wide application of RCA and use by other researchers, the implementations of all employed methods are made publicly available on the website of the Biomedical Image Analysis group.<sup>8</sup>

### ACKNOWLEDGMENT

The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of two Titan X GPUs for our research. They would also like to thank Darko Zikic for inspiring discussions on early ideas that have led to this work.

### REFERENCES

- [1] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient N-D image segmentation," *Int. J. Comput. Vis.*, vol. 70, no. 2, pp. 109–131, 2006.
- [2] J. E. Iglesias and M. R. Sabuncu, "Multi-atlas segmentation of biomedical images: A survey," *Med. Image Anal.*, vol. 24, no. 1, pp. 205–219, 2015.
- [3] T. Heimann and H.-P. Meinzer, "Statistical shape models for 3D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, no. 4, pp. 543–563, 2009.
- [4] E. Geremia *et al.*, "Classification forests for semantic segmentation of brain lesions in multi-channel MRI," in *Decision Forests for Computer Vision and Medical Image Analysis*. London, U.K.: Springer, 2013, pp. 245–260.
- [5] K. H. Zou *et al.*, "Statistical validation of image segmentation quality based on a spatial overlap index: Scientific reports," *Acad. Radiol.*, vol. 11, no. 2, pp. 178–189, 2004.
- [6] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.
- [7] W. R. Crum, O. Camara, and D. L. G. Hill, "Generalized overlap measures for evaluation and validation in medical image analysis," *IEEE Trans. Med. Imag.*, vol. 25, no. 11, pp. 1451–1461, Nov. 2006.
- [8] A. A. Taha and A. Hanbury, "Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool," *BMC Med. Imag.*, vol. 15, no. 1, p. 29, 2015.
- [9] X. Deng *et al.*, "On simulating subjective evaluation using combined objective metrics for validation of 3D tumor segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2007, pp. 977–984.
- [10] C. Ledig, W. Shi, W. Bai, and D. Rueckert, "Patch-based evaluation of image segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 3065–3072.
- [11] E. Konukoglu, B. Glocker, D. H. Ye, A. Criminisi, and K. M. Pohl, "Discriminative segmentation-based evaluation through shape dissimilarity," *IEEE Trans. Med. Imag.*, vol. 31, no. 12, pp. 2278–2289, Dec. 2012.
- [12] B. van Ginneken, T. Heimann, and M. Styner, "3D segmentation in the clinic: A grand challenge," in *Proc. MICCAI Workshop 3D Segmentation Clin.*, 2007, pp. 7–15.
- [13] C. Sudlow *et al.*, "UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age," *PLoS Med.*, vol. 12, no. 3, p. e1001779, 2015.
- [14] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, "The reusable holdout: Preserving validity in adaptive data analysis," *Science*, vol. 349, no. 6248, pp. 636–638, 2015.
- [15] A. Baraldi, L. Bruzzone, and P. Blonda, "Quality assessment of classification and cluster maps without ground truth knowledge," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 4, pp. 857–873, Apr. 2005.
- [16] Y. Liu, J. Wang, S. Cho, A. Finkelstein, and S. Rusinkiewicz, "A no reference metric for evaluating the quality of motion deblurring," *ACM. Graph.*, vol. 32, no. 6, pp. 175-1–175-12, 2013.
- [17] D. Cerrato, R. Jones, and A. Gupta, "Classification of proxy labeled examples for marketing segment generation," in *Proc. 17th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2011, pp. 343–350.
- [18] P. L. Correia and F. Pereira, "Stand-alone objective segmentation quality evaluation," *EURASIP J. Appl. Signal Process.*, vol. 2002, pp. 390–402, 2002.
- [19] F. Ge, S. Wang, and T. Liu, "New benchmark for image segmentation evaluation," *J. Electron. Imag.*, vol. 16, no. 3, p. 033011, 2007.
- [20] L. Goldmann *et al.*, "Towards fully automatic image segmentation evaluation," in *Proc. Int. Conf. Adv. Concepts Intell. Vis. Syst.*, 2008, pp. 566–577.
- [21] H. Li, J. Cai, T. N. A. Nguyen, and J. Zheng, "A benchmark for semantic image segmentation," in *Proc. IEEE Int. Conf. Multimedia Expo*, Jul. 2013, pp. 1–6.
- [22] B. Lamiroy and T. Sun, "Computing precision and recall with missing or uncertain ground truth," in *Graphics Recognition. New Trends and Challenges*. Berlin, Germany: Springer-Verlag, 2013, pp. 149–162.
- [23] H. Zhang, S. Chollet, S. A. Goldman, and J. E. Fritts, "Meta-evaluation of image segmentation using machine learning," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 1, Jun. 2006, pp. 1138–1145.
- [24] S. Chabrier, B. Emile, C. Rosenberger, and H. Laurent, "Unsupervised performance evaluation of image segmentation," *EURASIP J. Appl. Signal Process.*, vol. 2006, pp. 217–229, Dec. 2006.
- [25] H. Zhang, J. E. Fritts, and S. A. Goldman, "Image segmentation evaluation: A survey of unsupervised methods," *Comput. Vis. Image Understand.*, vol. 110, no. 2, pp. 260–280, 2008.
- [26] R. Unnikrishnan, C. Pantofaru, and M. Hebert, "Toward objective evaluation of image segmentation algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 6, pp. 929–944, Jun. 2007.
- [27] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): An algorithm for the validation of image segmentation," *IEEE Trans. Med. Imag.*, vol. 23, no. 7, pp. 903–921, Jul. 2004.
- [28] X. Li, B. Aldridge, R. Fisher, and J. Rees, "Estimating the ground truth from multiple individual segmentations incorporating prior pattern analysis with application to skin lesion segmentation," in *Proc. IEEE Int. Symp. Biomed. Imaging, Nano Macro*, Chicago, IL, USA, 2011, pp. 1438–1441.
- [29] S. Bouix *et al.*, "On evaluating brain tissue classifiers without a ground truth," *NeuroImage*, vol. 36, no. 4, pp. 1207–1224, 2007.
- [30] K. Sikka and T. M. Deserno, "Comparison of algorithms for ultrasound image segmentation without ground truth," *Proc. SPIE*, vol. 7627, pp. 76271C1–76271C9, Feb. 2010.
- [31] T. Kohlberger, V. Singh, C. Alvino, C. Bahlmann, and L. Grady, "Evaluating segmentation error without ground truth," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2012, pp. 528–536.
- [32] L. Grady, V. Singh, T. Kohlberger, C. Alvino, and C. Bahlmann, "Automatic segmentation of unknown objects, with application to baggage security," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2012, pp. 430–444.
- [33] K. Frounchi, L. C. Briand, L. Grady, Y. Labiche, and R. Subramanian, "Automating image segmentation verification and validation by learning test oracles," *Inf. Softw. Technol.*, vol. 53, no. 12, pp. 1337–1348, 2011.
- [34] E. Zhong, W. Fan, Q. Yang, O. Verscheure, and J. Ren, "Cross validation framework to choose amongst models and datasets for transfer learning," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, 2010, pp. 547–562.
- [35] W. Fan and I. Davidson, "Reverse testing: An efficient framework to select amongst classifiers under sample selection bias," in *Proc. 12th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2006, pp. 147–156.
- [36] D. Zikic, B. Glocker, and A. Criminisi, "Encoding atlases by randomized classification forests for efficient multi-atlas label propagation," *Med. Image Anal.*, vol. 18, no. 8, pp. 1262–1273, 2014.
- [37] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [38] A. Criminisi, J. Shotton, and E. Konukoglu, "Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning," *Found. Trends Comput. Graph. Vis.*, vol. 7, nos. 2–3, pp. 81–227, Feb. 2012.
- [39] K. Kamnitsas *et al.*, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Med. Image Anal.*, vol. 36, pp. 61–78, Feb. 2017.
- [40] W. Bai *et al.*, "A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: Application to cardiac MR images," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1302–1315, Jul. 2013.

<sup>8</sup><https://biomedica.doc.ic.ac.uk/software/>