



ELSEVIER

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Explainable cardiac pathology classification on cine MRI with motion characterization by semi-supervised learning of apparent flow



Qiao Zheng*, Hervé Delingette, Nicholas Ayache

Université Côte d'Azur, Inria, 2004 Route des Lucioles, 06902 Sophia Antipolis, France

ARTICLE INFO

Article history:

Received 8 November 2018

Revised 27 March 2019

Accepted 4 June 2019

Available online 6 June 2019

Keywords:

Cardiac pathology

Classification

Cine MRI

Motion

Deep learning

Semi-supervised learning

Neural network

Apparent flow

ABSTRACT

We propose a method to classify cardiac pathology based on a novel approach to extract image derived features to characterize the shape and motion of the heart. An original semi-supervised learning procedure, which makes efficient use of a large amount of non-segmented images and a small amount of images segmented manually by experts, is developed to generate pixel-wise apparent flow between two time points of a 2D+t cine MRI image sequence. Combining the apparent flow maps and cardiac segmentation masks, we obtain a local apparent flow corresponding to the 2D motion of myocardium and ventricular cavities. This leads to the generation of time series of the radius and thickness of myocardial segments to represent cardiac motion. These time series of motion features are reliable and explainable characteristics of pathological cardiac motion. Furthermore, they are combined with shape-related features to classify cardiac pathologies. Using only nine feature values as input, we propose an explainable, simple and flexible model for pathology classification. On ACDC training set and testing set, the model achieves 95% and 94% respectively as classification accuracy. Its performance is hence comparable to that of the state-of-the-art. Comparison with various other models is performed to outline some advantages of our model.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Cine magnetic resonance imaging (cine MRI) is widely used in the clinic as an approach to identify cardiac pathology. For both the patients and the clinicians, there is hence a great need for automated accurate cardiac pathology identification and classification based on MRI images as mentioned in [Rueckert et al. \(2016\)](#) and [Comaniciu et al. \(2016\)](#), as well as in the myocardial infarct classification challenge run at the STACOM workshop in 2015 ([Suinesiaputra et al. \(2018\)](#)). Recently, the state-of-the-art cardiac pathology classification methods extract various features from MRI images and perform classification based on these features. Despite the great results achieved so far, there are still some aspects that need to be further explored.

First, most classification models, including the state-of-the-art models, take many feature values together as input to a single or a group of machine learning classifiers (e.g. [Khened et al. \(2017\)](#), [Khened et al. \(2018\)](#), [Wolterink et al. \(2017\)](#), [Cetin et al. \(2017\)](#),

[Isensee et al. \(2017\)](#)), and output the predicted probability distribution over several classes. Like many other machine learning methods, or more specifically like most deep learning methods, these classification models are not easy to interpret. On the one hand, most of the models contain at least hundreds of parameters and it is impractical to examine and explain the role of each parameter. On the other hand, as many features are used simultaneously, it is hard to tell in a straightforward manner which feature value contributes to the identification of which category. This drawback on explainability causes many problems as pointed out in [Holzinger et al. \(2017\)](#). For instance, the lack of explainability is a significant hurdle for their widespread adoption in the clinic despite their performance. Moreover, under the new European General Data Protection Regulation, it may also generate legal issues in business, as companies are required to be able to explain why decisions have been made by their models upon demand. Hence we propose a simple classification model with 9 input features and 14 parameters in total such that the role and contribution of each feature or parameter are clear and explainable.

Second, in terms of data availability in medical image analysis, we usually have access to a large amount of unlabeled data and a small amount of labeled data. How to make good use of

* Corresponding author.

E-mail address: qiao.zheng@inria.fr (Q. Zheng).

the available data to train automatic methods remains an open question (Weese and Lorenz (2016)). Semi-supervised learning appears to be a powerful approach to tackle this challenge in general (Bai et al. (2017), Gu et al. (2017), Cheplygina et al. (2018)). In this paper, while cardiac motion is estimated in a flow-based manner like in many other methods (Gao et al. (2016), Parajuli et al. (2017)), we extend it as a semi-supervised learning method to train a network for apparent flow generation, using the dataset of Automatic Cardiac Diagnosis Challenge (ACDC) of MIC-CAI 2017 (Bernard et al. (2018)), for which the ground-truth segmentation mask is only available for 2 time frames. Although the percentage of the segmented frames in the dataset is small, making efficient use of their segmentation masks in training is essential for the generated flow to have better consistency. In particular, with the supervision of the masks in training, we show that cardiac structures are better preserved after warping by the generated flow.

Third, the state-of-the-art classification methods most exclusively focus on features extracted at two instants: the instants of end-diastole (ED) and end-systole (ES). The other instants are often ignored in pathology classification. For example, in the ACDC challenge, 3 out of the 4 cardiac pathology classification methods, including Khened et al. (2017) (as well as its updated version Khened et al. (2018)), Wolterink et al. (2017) and Cetin et al. (2017), use only features based on ED and ES. The authors of Isensee et al. (2017) propose the only method in the ACDC challenge which explores the instants other than ED and ES by quantifying the volume change and by measuring the LV-RV dissynchrony. Yet much information about cardiac motion (e.g. how individual myocardial segments move) is still excluded from the extracted features. While more and more research efforts are put on cardiac motion estimation (e.g. Qin et al. (2018a), Qin et al. (2018b), Xue et al. (2018), Yang et al. (2017), Yan et al. (2018)) and cardiac disease assessment via motion analysis (e.g. Gilbert et al. (2017), Dawes et al. (2017), Lu et al. (2018)), we propose to explore the impact of specific motion features to learn the detection of cardiac pathologies by extracting some useful time series of simple and straightforward features from cine MRI image sequences. Ideally, the resulting time series should be both informative enough to be used for classification and intuitive to be understood by a physician.

In this paper, we propose a novel and explainable method to classify a subset of cardiac pathologies using deep learning of cardiac motion (in the form of apparent flows) and shape. Our main contribution is threefold:

- **Semi-supervised learning of flow:** a novel semi-supervised learning method is applied to train a neural network model, which outputs apparent flows given two MRI images from the same 2D+t cine MRI image sequence. This allows to learn the motion as apparent flows efficiently from both segmented and non-segmented image data.
- **Motion-characteristic features:** combining the apparent flows across time with cardiac segmentation, time series of the radius and thickness of myocardial segments are extracted to describe cardiac motion. As features, they are easy to interpret and allow to characterize different shapes and motions of cardiac pathologies.
- **Explainable classification model:** we train a set of 4 simple classifiers to perform binary classifications. Each classifier performs a logistic regression and takes no more than 3 feature values as input, which makes it very simple and easy to interpret. On the ACDC challenge training set and testing set, our model achieves 95% and 94% as classification accuracy respectively, which is comparable to the state-of-the-art.

2. Data

2.1. Dataset

The proposed method is trained and evaluated on the ACDC challenge dataset, which consists of a training set of 100 cases and a testing set of 50 cases. The cine MRIs were acquired with a conventional SSFP sequence (Bernard et al. (2018)). Most of the cases contain about 10 slices of short-axis MRIs. And the number of frames in the cases varies between 12 and 35. ACDC training set and testing set are respectively divided into 5 pathological groups of equal size (we cite below the properties of each group as provided on the website, though they are only roughly exact according to our measure and observation):

- dilated cardiomyopathy (DCM): left ventricle cavity (LVC) volume at ED larger than 100 mL/m^2 and LVC ejection fraction lower than 40%
- hypertrophic cardiomyopathy (HCM): left ventricle (LV) cardiac mass higher than 110 g/m^2 , several myocardial segments with a thickness higher than 15 mm at ED and a normal ejection fraction
- myocardial infarction (MINF): LVC ejection fraction lower than 40% and several myocardial segments with abnormal contraction
- RV abnormality (RVA): right ventricle cavity (RVC) volume higher than 110 mL/m^2 or RVC ejection fraction lower than 40%
- normal subjects (NOR)

Please note that the abnormal contraction mentioned in the characteristics of MINF is quite vague as a property. In addition, both MINF and DCM cases have low LVC ejection fractions. And sometimes, a myocardial infarction causes a dilated LVC (for which we should classify the case to MINF instead of DCM according to ACDC challenge). As we will present later, it is indeed a challenge to distinguish them.

For the cases of ACDC training set, expert manual segmentation for LVC, RVC and the left ventricular myocardium (LVM) is provided as ground-truth for all slices at ED and ES phases; all other structures in the image are considered as background. For the cases of ACDC testing set, no ground-truth information about classification or segmentation is available. For performance evaluation on the testing set, the predicted results of a model need to be submitted online.

2.2. Notation

In this paper, slices in image stacks are indexed in spatial order from the basal part to the apical part of the heart. Given an image stack S , we denote N_S the number of its slices. Given two values a and b between 0 and $N_S - 1$, we note $S[a, b]$ the sub-stack consisting of slices of indexes in the interval $[\text{round}(a), \text{round}(b)]$ ($\text{round}(a)$ is included while $\text{round}(b)$ is excluded).

3. Methods

Our method mainly consists of two parts: feature extraction (Fig. 1) and (classification based on features (Fig. 2)). But the region of interest (ROI) needs to be determined first.

3.1. Preprocessing: Region of interest (ROI) determination

As a preprocessing step, the ROI needs to be determined on the original MRI images. Short-axis MRI images usually cover a zone much larger than that of the heart. To save memory usage and to increase the speed of apparent flow and segmentation methods, it is better to work on an appropriate ROI instead. For this

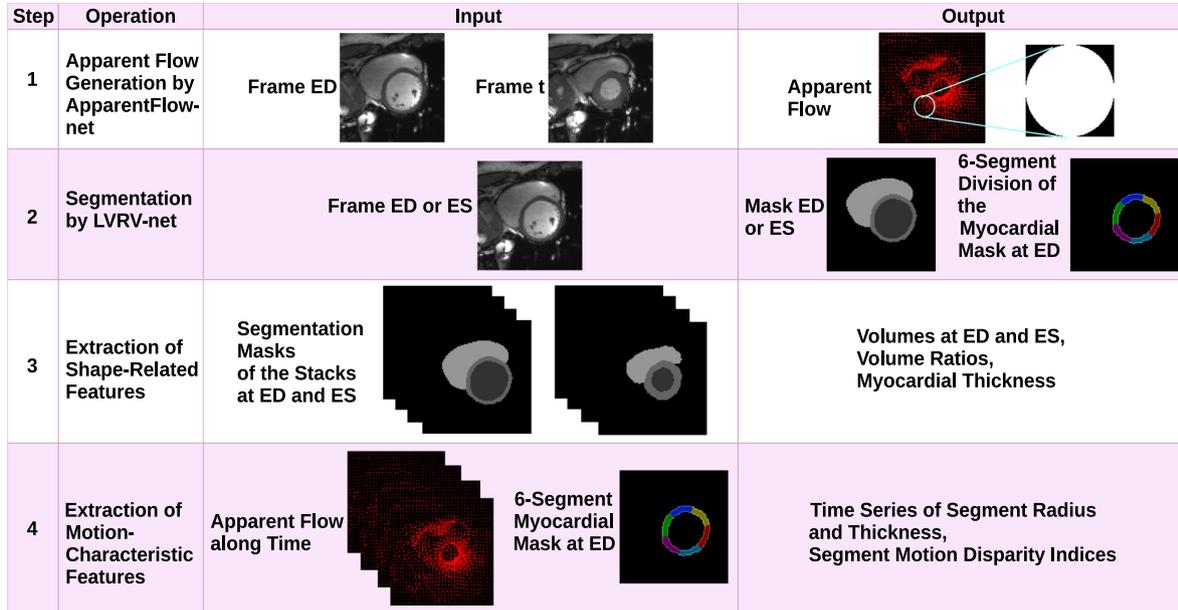


Fig. 1. Overview of the feature extraction method: 1. Apparent flow generation given the ED frame and another frame on the same slice; 2. Cardiac segmentation on the ED and ES frames and division of the ED myocardium mask to 6 segments; 3. Extraction of the shape-related features, including the calculation of the volumes, volume ratios and myocardial thickness of a heart given the segmentation masks; 4. Extraction of motion-characteristic features, including the creation of segment radius and thickness time series given a slice with the corresponding apparent flow maps and segmentation mask.

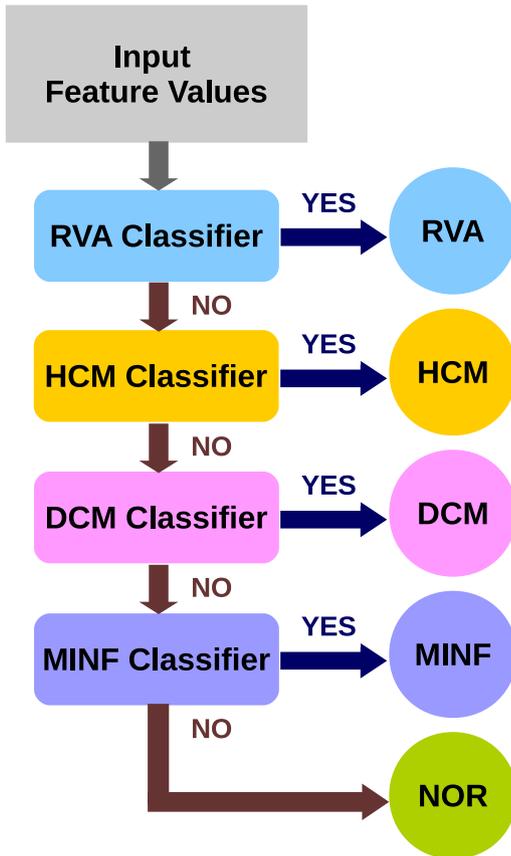


Fig. 2. Overview of the classification method: the 4 binary classifiers are applied in sequence to classify a case to RVA, HCM, DCM, MINF or NOR.

purpose, we directly apply an existing ROI method: we use the trained ROI-net exactly as described in Zheng et al. (2018) to define an ROI. Briefly speaking, the ROI-net is a variant of U-net (Ronneberger et al. (2015)) for heart/background binary segmen-

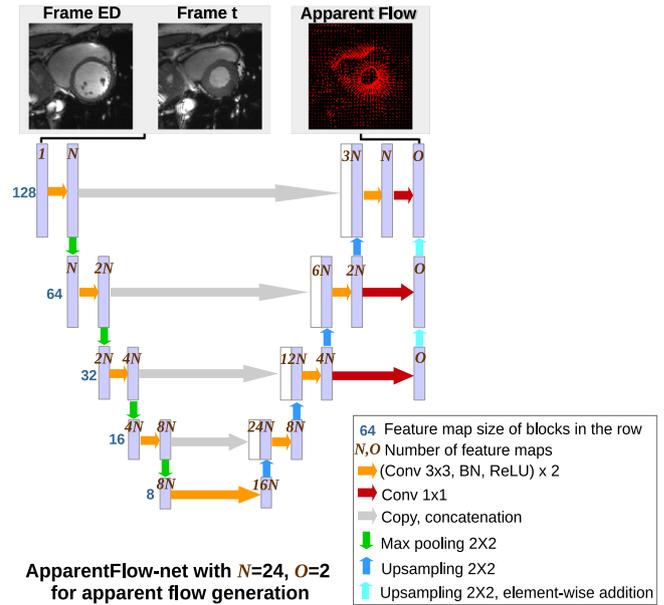


Fig. 3. ApparentFlow-net: for apparent flow generation. The output is a map of pixel-wise flow F_t .

tation. It is applied on several middle slices on the ED image stack. As shown in Zheng et al. (2018), this ROI determination method is very robust and succeeds in all cases of the ACDC dataset. In the remainder of this paper, we only refer to the automatically cropped ROI of the images.

3.2. Feature extraction step 1: Apparent flow generation

As shown in Fig. 1, there are four steps for feature extraction. In this first step, the ApparentFlow-net, which is a variant of U-net (Ronneberger et al. (2015)) as shown in Fig. 3, is proposed. U-net, with the encoder-decoder structure consisting of

Adjacent Pixel Pairs and Their Transformed Positions by Apparent Flow								
Crossing (and Hence Penalty) on the X-Components	NO	YES	NO	YES	-	-	-	-
Crossing (and Hence Penalty) on the Y-Components	-	-	-	-	NO	YES	NO	YES

Fig. 4. Examples of adjacent pixel pairs transformed by apparent flow for which the crossing penalty applies or not.

layers of various sizes of receptive fields, can effectively integrate local and global information, which is necessary for the analysis of the shape and motion of the heart on MRIs. Previously, we successfully used some variants of U-net for cardiac segmentation (Zheng et al. (2018)). So we expect a similar structure would also work for the estimation of cardiac motion. The ApparentFlow-net is applied to generate pixel-wise apparent flow given a pair of image frames on the same slice as input: the ED frame and another frame of index t on the same slice. In other words, the generated apparent flow map is a displacement field of the slice between ED and instant t . In a later step, combined with the segmentation mask, we will extract cardiac motion features from the sequences of apparent flow maps on a slice. The details of this extraction are available in the sub-section 3.5. While there exists some researches that explore image registration (or equivalently, apparent flow) using unsupervised learning (e.g. Balakrishnan et al. (2018), Krebs et al. (2018), de Vos et al. (2017), Li and Fan (2017)), we propose a semi-supervised learning approach to make efficient use of a large amount of non-segmented images and a small amount of images segmented manually by experts.

In general, the idea of representing motion by apparent flow is based on two assumptions. First, we assume that the pixel intensities of an object do not change much between the two frames. Second, it is assumed that neighboring pixels have similar motion. By observation, we find that these assumptions usually hold on the slices located below the base and above the apex with some margin. This is due to the limited out-of-plane motion on these slices (this is less the case for the slices around the base and the apex). Hence ApparentFlow-net is trained and applied on the middle slices only.

If we note $I_{ED}(\mathbf{P})$ and $I_t(\mathbf{P})$ the pixel intensity of the two input frames of ApparentFlow-net at position $\mathbf{P} = (x, y)$, according to the first assumption above, ApparentFlow-net should generate an apparent flow map \mathbf{F}_t with $\mathbf{F}_t(\mathbf{P}) = (F_t^x(\mathbf{P}), F_t^y(\mathbf{P}))$ between ED and t enabling image reconstruction such that the following intensity discrepancy is minimized:

$$L_{IMG}(\mathbf{F}_t) = \sum_{\mathbf{P}} \left(I_{ED}(\mathbf{P}) - I_t(\mathbf{P} + \mathbf{F}_t(\mathbf{P})) \right)^2 \quad (1)$$

Meanwhile, the flow should also preserve the regularity of the motion of neighboring pixels according to the second assumption above. While there are already some methods in the community to impose diffeomorphisms (e.g. demon's algorithm as in Pennec et al. (1999), LDDMM as in Hernandez et al. (2008)), we propose a simple one to only discourage the occurrence of the extreme situations such as the crossing between two adjacent pixels or rotations greater than 90° (Fig. 4). As long as these unrealistic motion patterns do not appear, there is no penalty on the regularity at all and the network is free to generate whatever flow without being influenced by the regularity constraint. More precisely, let us note $W_{\mathbf{F}_t}$ as the warping function such that $W_{\mathbf{F}_t}(\mathbf{P}) =$

$\mathbf{P} + \mathbf{F}_t(\mathbf{P})$. For two adjacent pixels $\mathbf{P} = (x, y)$ and $\mathbf{P}^{x+} = (x + 1, y)$ in a row, we want the warped pixel $W_{\mathbf{F}_t}(\mathbf{P}^{x+})$ to stay on the right of the warped pixel $W_{\mathbf{F}_t}(\mathbf{P})$ (similarly for the adjacent pixels \mathbf{P} and $\mathbf{P}^{y+} = (x, y + 1)$ in a column) (see Fig. 4). Otherwise, we say that a crossing on the x-components (y-components) of the flow pairs occurs and a penalty should apply. This translates as the following criterion to be minimized (more details about the derivation are available in Appendix A):

$$L_{CROSS}(\mathbf{F}_t) = \sum_{\mathbf{P}} \min \left(1 + \frac{\partial F_t^x(\mathbf{P})}{\partial x}, 0 \right)^2 + \min \left(1 + \frac{\partial F_t^y(\mathbf{P})}{\partial y}, 0 \right)^2 \quad (2)$$

Moreover, we further encourage the flow to preserve the segmentation masks of cardiac structures $S \in \{LVC, LVM, RVC\}$. The warped segmentation masks of these structures should approximately match the ground-truth masks on the corresponding frame. Let us note M_{ED}^S and M_{ES}^S the binary ground-truth segmentation mask (of pixel intensity value 0 or 1) of S at the instants of ED and ES (the only instants for which the ground-truth is available in the ACDC training set). This constraint on the flow between ED and ES is based on the Dice coefficient

$$L_{GT}(\mathbf{F}_{ES}) = \sum_{S \in \{LVC, LVM, RVC\}} Dice(M_{ED}^S, M_{ES}^S \circ W_{\mathbf{F}_{ES}}) \quad (3)$$

The formula of the Dice function is provided in Appendix A.

Finally, the overall loss function for training the ApparentFlow-net is a linear combination of the terms L_{IMG} , L_{CROSS} and potentially L_{GT} . We adopt a semi-supervised approach for which L_{GT} is applied when ground-truth segmentation is available:

$$L_{flow}(\mathbf{F}_t) = L_{IMG}(\mathbf{F}_t) + p_1 L_{CROSS}(\mathbf{F}_t) + p_2 \mathbf{1}_{t=ES} L_{GT}(\mathbf{F}_t) \quad (4)$$

where $\mathbf{1}_{t=ES}$ is the indicator function for the event $t = ES$. $\mathbf{1}_{t=ES}$ is necessary as for the instants t other than ED and ES, the ground-truth segmentation is not provided in ACDC. Please note that this is a typical method of semi-supervised learning. It makes use of a small amount of labeled data (the images with ground-truth segmentation) and a large amount of unlabeled data (the images without ground-truth).

3.3. Feature extraction step 2: Segmentation

In this step, an existing model for segmentation proposed in Zheng et al. (2018), the LVRV-net, is applied to segment MRI image stacks as presented in Zheng et al. (2018). With the concept of propagation along the long axis, this method was proven to be robust, as the results achieved on several different datasets are all comparable or even better than the state-of-the-art. For more details about the structure, training and application of the LVRV-net, please refer to Zheng et al. (2018). When we train and evaluate our method on the ACDC training set (100 cases), in each

fold of a 5-fold cross-validation, the trained LVRV-net as given by Zheng et al. (2018) is finetuned with the 80 cases used for training before being applied on the remaining 20 cases; for the evaluation of our method on ACDC testing set (50 cases), the trained LVRV-net is first finetuned with the 100 cases of ACDC training set.

In fact, in Zheng et al. (2018), LVRV-net was trained to start the segmentation propagation from a given slice on which the ventricle cavities are supposed to be present. In other words, it was only trained to identify LV and RV labels on the slices below the base. So it might not work well if the basal slice is not determined in a stack and if the top slice in the volumetric image is located above the base. In this case, if we apply the original LVRV-net starting from the top slice, it might make a false positive prediction. With finetuning on ACDC, we find that this issue is solved. In general, the finetuned LVRV-net successfully learns from the ground-truth segmentation masks of ACDC that no foreground pixel is present (i.e. predict everything to be background) on the slices above the base and start segmentation propagation only when the base is reached. So it is no longer necessary to determine the basal slice manually. On the resulting sets of segmentation masks, we can hence also determine the location of the base, which is necessary for the calculation of volumes and the determination of sub-stacks for motion extraction as we will present later.

With the segmentation mask, we determine \mathbf{B}_L and \mathbf{B}_R , the barycenters of LVC and RVC respectively. Then all the pixels \mathbf{P} labeled to LVM on the segmentation mask are divided into 6 segments, depending on in which interval $[k\pi/3, (k+1)\pi/3[$ for k in $[0,5]$ the angle between the vectors $\mathbf{B}_L\mathbf{P}$ and $\mathbf{B}_L\mathbf{B}_R$ is. An example of the resulting 6 segments are shown in Fig. 1. This division of segments is inspired by the 17-segment system recommended by the American Heart Association (AHA) in Cerqueira et al. (2002). Indeed, in the AHA system, on all short-axis slices around the base and at the level of mid-cavity, the myocardium is divided into 6 segments.

3.4. Feature extraction step 3: Shape-Related features

Based on the segmentation masks generated in the previous step, we estimate the volumes of LVC, LVM and RVC of a case at ED and ES. For each of the two phases, the volume of LVC is calculated by approximating the LVC between two adjacent slices as a truncated cone and summing up all the truncated cone volumes:

$$V_{LVC} = \sum_i \left(S_i + S_{i+1} + \sqrt{S_i S_{i+1}} \right) (L_{i+1} - L_i) / 3 \quad (5)$$

where S_i is the area of LVC on the slice i and L_i is the slice position along the long axis. The volume of LVM and RVC is calculated in a similar way. Then we normalize all the volumes by the corresponding body surface area (BSA) of the subject, which is a traditional practice based on the assumption that BSA is related to the metabolic rate. BSA can be computed from the height and the weight provided in ACDC (using the Mosteller formula $BSA = \sqrt{\text{height} * \text{weight} / 60}$).

With the segmentation masks and volumes at ED and ES, we then compute the 7 shape-related features as listed in the first 7 rows of Table 1.

3.5. Feature extraction step 4: Motion-Characteristic features

3.5.1. Slice selection

For each case, let S be the image stack at ED (following the Notation part in the previous section). Given the segmentation masks of each slice generated in Step 2, we note i_1 the index of the first slice on which RVC mask is present (roughly the first slice below the base), and i_2 the index of the last slice on which LVC mask is

Table 1

The extracted features used by our classification model.

Feature	Notion (and Definition)
RVC volume at ED	$V_{RVC,ED}$
LVC volume at ES	$V_{LVC,ES}$
RVC ejection fraction	$EF_{RVC} (= 1 - V_{RVC,ES}/V_{RVC,ED})$
LVC ejection fraction	$EF_{LVC} (= 1 - V_{LVC,ES}/V_{LVC,ED})$
Ratio between RVC and LV volumes at ED	$R_{RVC,ED}$ ($= V_{RVC,ED}/(V_{LVC,ED} + V_{LVM,ED})$)
Ratio between LVM and LVC volumes at ED	$R_{LVM,ED}$ ($= V_{LVM,ED}/V_{LVC,ED}$)
Maximal LVM thickness in all the slices at ED	$MT_{LVM,ED}$
Radius motion disparity	RMD
Thickness motion disparity	TMD

present (roughly the last slice above the apex), and $h = i_2 - i_1 + 1$. Then we focus on extracting motion information from the sub-stack $S_{mid} = S[i_1 + 0.1h, i_2 + 1 - 0.2h]$. Please note that among the slices between the base and the apex, we exclude the top 10% and the bottom 20% and consider the remaining 70% in the middle, since the out-of-plane motion is particularly large in the slices close to the base or the apex.

3.5.2. Frame sampling

As presented in Fig. 1, for each slice in S_{mid} , let us note f the number of frames available (all the frames together form a cardiac cycle). We sample 10 frames of instant t_i for i in $[0,9]$, such that t_0 is the instant of ED and $t_i = \text{round}(t_0 + i * f/10) \bmod f$ for i in $[1,9]$. The 10 sampled frames hence cover the whole cardiac cycle. Applying the ApparentFlow-net of Step 1 in all the 9 pairs of frame (t_0, t_i) , we obtain 9 apparent flow maps \mathbf{F}_{t_i} . Hence for each pixel \mathbf{P} , we get its warped position $W_{\mathbf{F}_{t_i}}(\mathbf{P})$ for i in $[1,9]$.

3.5.3. Time series extraction

Then, with the convention that F_{t_0} is the null apparent flow (and hence $W_{\mathbf{F}_{t_0}}$ is the identity function), the barycenter of LVC at t_i for $i \in [0, 9]$, \mathbf{B}_i , is defined as the average of $W_{\mathbf{F}_{t_i}}(\mathbf{P})$ for all the pixels \mathbf{P} labeled as LVC on the segmentation mask M_{ED} at t_0 (the 2nd column of Fig. 5):

$$\mathbf{B}_i = \text{average}(\{W_{\mathbf{F}_{t_i}}(\mathbf{P}) \mid \mathbf{P} \in \text{LVC on } M_{ED}\}) \quad (6)$$

In a similar way, for each myocardial segment S_k ($k \in [0, 5]$) and each instant t_i ($i \in [0, 9]$), we define $\mathbf{I}_{k,i}$, the barycenter of the inner boundary of the myocardial segment S_k at t_i (the 3rd column of Fig. 5):

$$\mathbf{I}_{k,i} = \text{average}(\{W_{\mathbf{F}_{t_i}}(\mathbf{P}) \mid \mathbf{P} \in \text{LVC on } M_{ED} \text{ \& } \mathbf{P} \text{ has neighboring pixel}(s) \in S_k\}) \quad (7)$$

and $\mathbf{O}_{k,i}$, the barycenter of the outer boundary of the myocardial segment S_k at t_i (the 4th column of Fig. 5):

$$\mathbf{O}_{k,i} = \text{average}(\{W_{\mathbf{F}_{t_i}}(\mathbf{P}) \mid \mathbf{P} \in S_k \text{ \& } \mathbf{P} \text{ has neighboring pixel}(s) \in \text{background on } M_{ED}\}) \quad (8)$$

Finally, as shown in the 4th column of Fig. 5, we define the radius of S_k at t_i normalized by BSA:

$$RA_{k,i} = |\mathbf{B}_i \mathbf{I}_{k,i}| / \text{BSA} \quad (9)$$

and the thickness of S_k at t_i normalized by BSA:

$$T_{k,i} = |\mathbf{B}_i \mathbf{O}_{k,i}| / \text{BSA} - RA_{k,i} \quad (10)$$

We hence generate two time series $\{RA_{k,i}; i \in [0, 9]\}$ and $\{T_{k,i}; i \in [0, 9]\}$ to represent the contraction and the thickening of S_k .

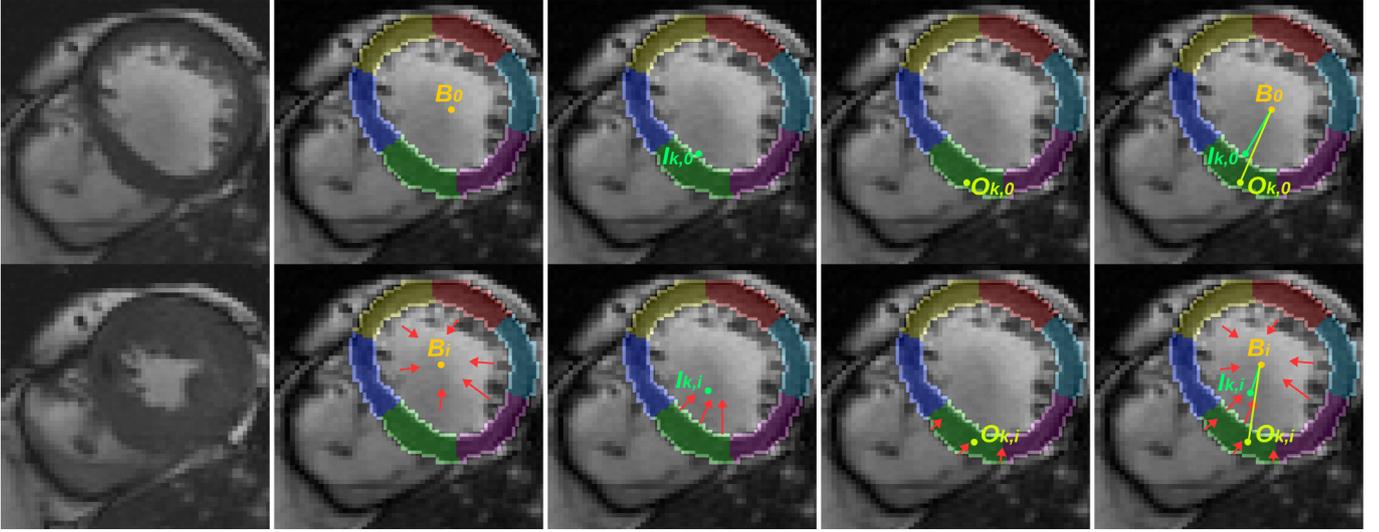


Fig. 5. Definitions of B_i , $I_{k,i}$, $O_{k,i}$, $RA_{k,i}$ and $T_{k,i}$, for the extraction of motion-characteristic time series. The first row shows the definitions at t_0 ; the second row presents the definitions at t_i for $i \in [1, 9]$. 1st column: Frames at t_0 and t_i , based on which the apparent flow is generated. 2nd column: B_i is the barycenter of warped LVC (segmented at t_0) at t_i . 3rd column: $I_{k,i}$ is the barycenter of the warped inner boundary of segment S_k at t_i . 4th column: $O_{k,i}$ is the barycenter of the warped outer boundary of segment S_k at t_i . 5th column: $RA_{k,i} = |B_i I_{k,i}| / BSA$, $T_{k,i} = |B_i O_{k,i}| / BSA - RA_{k,i}$.

3.5.4. Visual correspondence between time series and pathologies

We compute the two time series introduced above for all the slices in S_{mid} of all the cases in ACDC. From the majority of the cases, we manage to visually identify several typical slices with the time series characterizing the motion of the corresponding category. Examples of such typical slices are presented in Fig. 6. To sum up our observation on the typical slices of each category as shown in Fig. 6:

- NOR: all segments have similar radius and thickness at all instants; their contraction and thickening are synchronous and with comparable scales.
- HCM: the segments not only look proportionally thicker at ED, but also thicken more and contract stronger in the radial direction.
- DCM: the radiuses are quite large; the segments are moving so little that neither contraction nor thickening is obvious.
- MINF: the radiuses are quite large; some segments are clearly much more active than others.

3.5.5. Motion-Characteristic feature values

To better distinguish DCM and MINF cases, we define two additional feature values which often indicate the abnormal contraction described in the definition of MINF.

The first one is “radius motion disparity” (RMD). Given a case, we consider the set of radius series $\{RA_{k,i}; i \in [0, 9]\}$ of all the segments S_k on all the slices in the sub-stack S_{mid} (e.g. if there are 4 slices in S_{mid} , we consider a set of $6 \times 4 = 24$ time series). We first define the disparity of motion over all the segments in S_{mid} at the instant t_i as the difference between the maximum and minimum contraction at t_i :

$$RD_i = \max_{S_k \in S_{mid}} RA_{k,i} / RA_{k,0} - \min_{S_k \in S_{mid}} RA_{k,i} / RA_{k,0} \quad (11)$$

Then RMD is defined as the maximum disparity along the cardiac cycle:

$$RMD = \max_{i \in [0,9]} RD_i \quad (12)$$

The second motion-characteristic feature value is named “thickness motion disparity” (TMD). For each slice s in S_{mid} and each t_i ,

we define the thickness motion disparity of the slice s at t_i as

$$TD_{s,i} = \left(\max_{k \in [0,5]} T_{k,i} - \min_{k \in [0,5]} T_{k,i} \right) / \min_{k \in [0,5]} T_{k,0} \quad (13)$$

where we normalize the thicknesses by the minimum segment thickness at t_0 on slice s taking into account that myocardial thickness may vary across slice.

Finally, TMD is defined as

$$TMD = \max_{s \in S_{mid}, i \in [0,9]} TD_{s,i} \quad (14)$$

3.6. Classification

3.6.1. 4-Classifer classification model

Using the 7 shape-related features and the 2 motion-characteristic features as input, a classification model is proposed (Fig. 2) to classify the 5 pathological categories of ACDC. It consists of 4 binary classifiers:

- RVA classifier: RVA cases v.s. all the other cases.
- HCM classifier: HCM cases v.s. MINF, DCM and NOR cases.
- DCM classifier: DCM cases v.s. MINF and NOR cases.
- MINF classifier: MINF cases v.s. NOR cases.

The 4 binary classifications are arranged in increasing order of difficulty of the binary classification tasks. RVA and HCM cases can be identified based on the commonly used shape-related features. So they are classified first. DCM and MINF cases are somewhat similar in terms of sizes and ejection fractions. We use the novel motion-characteristic features to better distinguish them. Hence this more difficult classification is performed at the end.

3.6.2. Explainable manual feature selection

To keep the classifiers simple, limit their risk of overfitting and increase their explainability, we chose no more than 3 features for each classifier as shown in Table 2:

- For RVA classifier, according to the definition provided by ACDC, the RVC volume at ED and the RVC ejection fraction are the most relevant features. We add one more feature, the ratio

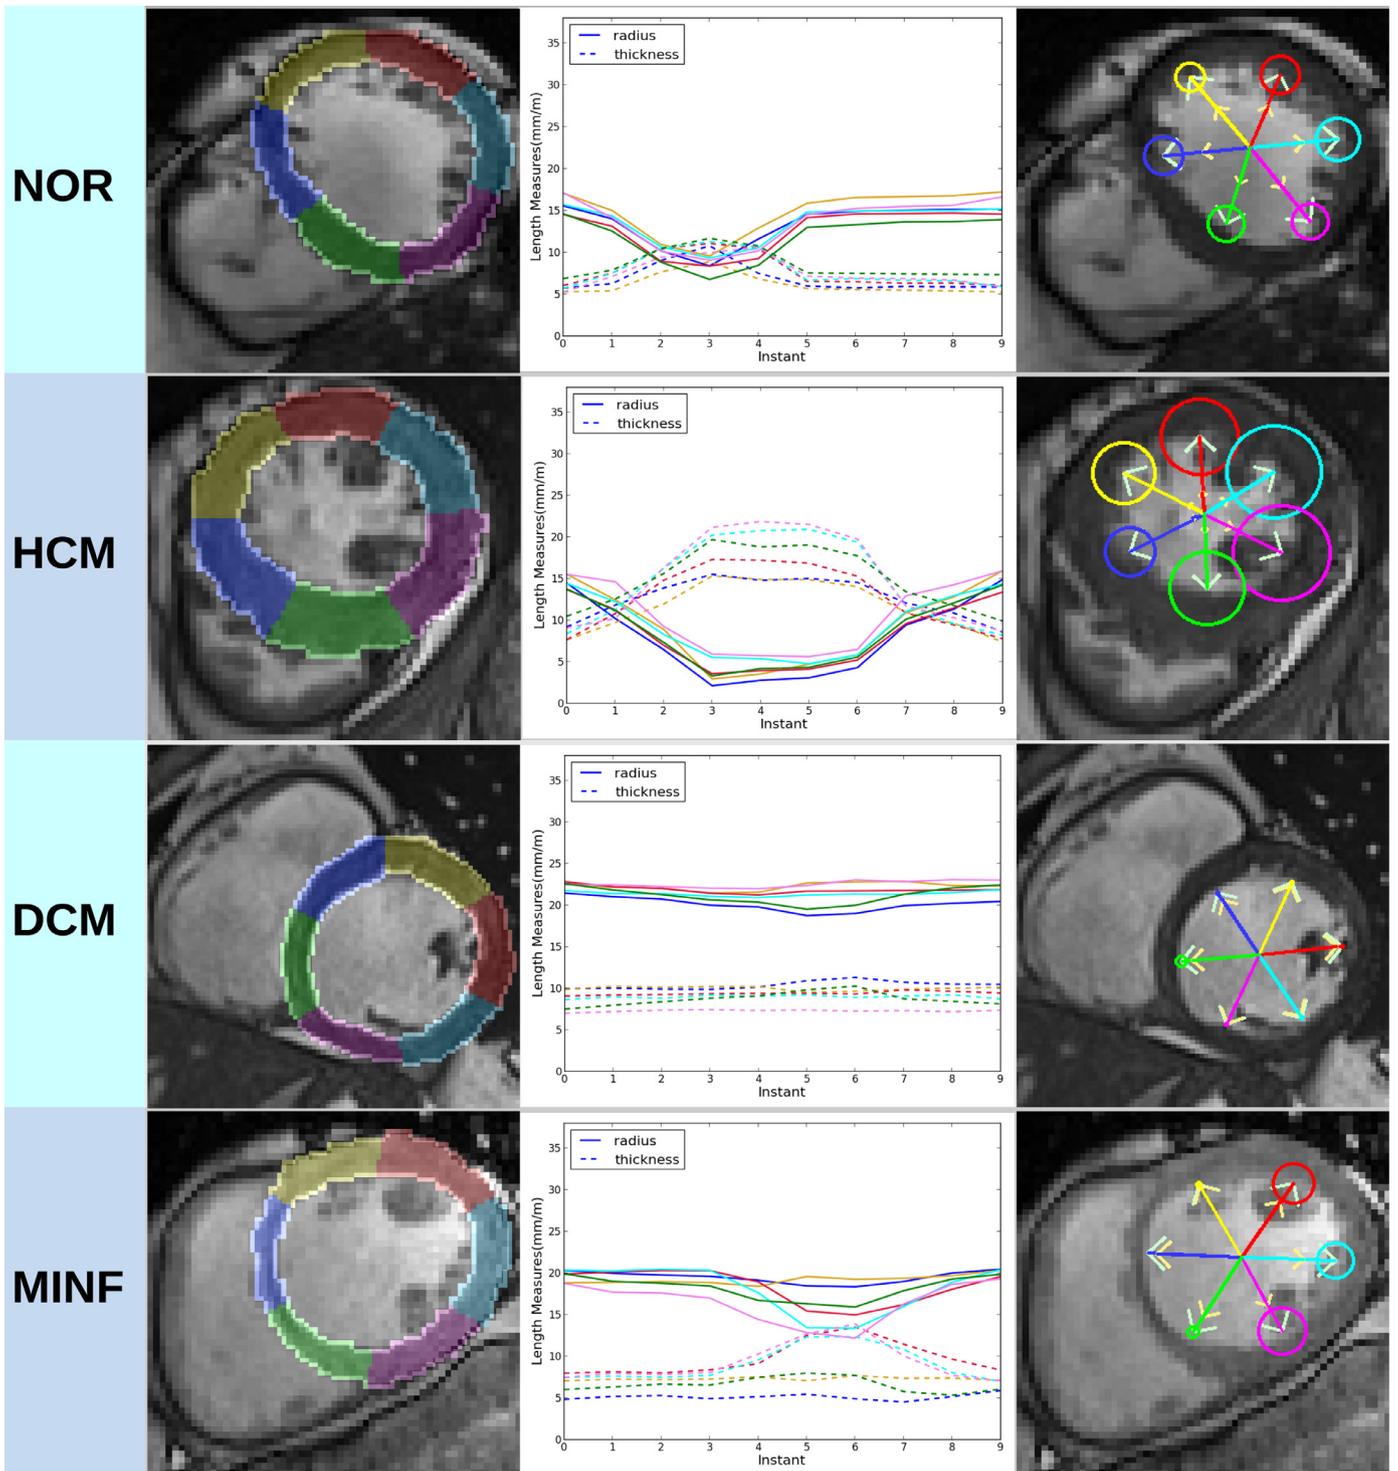


Fig. 6. Examples of typical slice from 4 of the 5 pathological categories in ACDC. 1st column: the segmentation of the 6 myocardial segments (the boundaries of the segmentation masks are marked by lighter colors). 2nd column: time series of the radius (solid lines) and the thickness (dotted lines) of the 6 segments. 3rd column: a visualization of the motion information. For each segment, the radius connecting the LVC barycenter \mathbf{B}_0 and the segment inner boundary barycenter $\mathbf{I}_{k,0}$ (marked by the light green arrow) at ED is plotted. The segment inner boundary barycenter at ES is marked by the light orange arrow. The radius of the circle is proportional to the difference of segment thicknesses at ES and ED ($\Delta T_{k,i}$). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- between RVC and LV volumes at ED, as we find that most RCA cases have disproportionately large RVC.
- For HCM classifier, LVC ejection fraction and maximal LVM thickness are selected according to the definition of HCM. The ratio between LVM and LVC volumes at ED is added because with most HCM cases this ratio is exceptionally high due to the exceptional myocardial thickness.

- For DCM classifier, as DCM cases are usually dilated at ED and inactive from ED to ES, their volumes of LVC at ES can be exceptionally large. So this feature is used. In addition, we also use radius motion disparity and thickness motion disparity.
- For MINF classifier, by definition, LVC ejection fraction is enough to distinguish MINF cases from NOR cases

Table 2
The input features of the 4 binary classifiers.

	Input Feature(s)
RVA Classifier	$V_{RVC,ED}, EF_{RVC}, R_{RVCLV,ED}$
HCM Classifier	$EF_{LVC}, R_{LVM/LVC,ED}, MT_{LVM,ED}$
DCM Classifier	$V_{LVC,ES}, RMD, TMD$
MINF Classifier	EF_{LVC}

3.6.3. Model of classifiers

Each of the 4 classifiers is just a ridge logistic regression model. For a training case of index m , if we note $f_{m,i}$ the i -th feature values used as input of the classifier and y_m (-1 or 1, corresponding to no or yes) the binary ground-truth of the case, then the classifier is trained by minimizing

$$L_{\text{classifier}}(\{p_i\}, b) = \frac{1}{2} \sum_i p_i^2 + C \sum_m \log \left(\exp \left(-y_m \left(\sum_i p_i f_{m,i} + b \right) \right) + 1 \right) \quad (15)$$

with respect to the parameters $\{p_i\}$ and b . C is the inverse of regularization strength. After the training is done, given a case of index l and feature values $f_{l,i}$, the prediction the sign of $\sum_i p_i f_{l,i} + b$. If it is non-negative, the prediction of the trained classifier is yes; otherwise it is no.

3.6.4. Flexibility and versatility of the model

Finally, we would also like to point out that the 4 classifiers function independently. While they are grouped together to form the proposed classification model in this paper, they can certainly be applied separately or grouped in a different manner in other situations if appropriate. This proposed classification model is hence very flexible and versatile.

4. Experiments and results

We evaluate our method in two different ways. On the one hand, the model is trained with ACDC training set and then tested on ACDC testing set. On the other hand, a 5-fold cross-validation is performed on ACDC training set. For the latter, the 100 cases of ACDC training set are partitioned into 5 subsets of 20 cases, such that in each subset there are exactly 4 cases of each of the 5 categories.

In addition, we also analyze the proposed model by comparing it with various other models. Since the ground-truth category is only available for the cases in the training set (and not for those in the testing set), this analysis is based on the results on the training set.

4.1. Training apparentflow-net

4.1.1. Parameters and data

In the training process with the whole ACDC training set, as well as in each of the 5 training processes of the 5-fold cross validation, the ApparentFlow-net is trained using the loss function $L_{\text{flow}}(\mathbf{F}_t)$ introduced in the Method section for 50 epochs with batch size 16. In terms of loss function parameter, we empirically find that $p_1 = 10^3$ and $p_2 = 10^5$ work well. These values are hence used for training. In terms of training data, for each case in the corresponding training set, we use the slices in the sub-stack $S[i_1 + 0.2h, i_2 + 1 - 0.2h]$ (with the notation introduced in the subsection 3.5.1). In other words, we approximately exclude the top 20% and the bottom 20% of all the slices covering the LV cavity, and select the remaining 60% in the middle. This slice selection for training (middle 60%) is slightly more conservative than that for the application of the method (middle 70%). This design is aimed

Table 3
The mean(standard deviation) of Dice coefficients achieved by comparing $M_{ES} \circ W_{FES}$ and M_{ED} for 3 cardiac structures in the 5-fold cross-validation on ACDC training set.

Training Method	Dice		
	LVM	LVC	RVC
semi-supervised (proposed)	0.84(0.07)	0.94(0.07)	0.87(0.19)
unsupervised	0.76(0.08)	0.93(0.06)	0.83(0.22)

to further reduce the impact of the out-of-plane motion in training. For each selected slice, the frame pairs of indices (ED, t) for all frame index t are used to train the ApparentFlow-net. Only when $t = ES$, the term $L_{GT}(\mathbf{F}_t)$ in $L_{\text{flow}}(\mathbf{F}_t)$ using the segmentation ground truth is applied. With our automatic slice selection approach, in total, there are 13,672 frame pairs used for training in the ACDC training set. Among the 13,672 frame pairs, only 515 pairs (3.77%) come with segmentation ground-truth such that the term $L_{GT}(\mathbf{F}_t)$ applies.

4.1.2. Performance

To measure its performance, in each evaluation of the 5-fold cross-validation, for all the slices in the sub-stack $S[i_1 + 0.2h, i_2 + 1 - 0.2h]$ of all the 20 cases for evaluation, we apply the trained ApparentFlow-net to generate \mathbf{F}_{ES} . Then we use it to warp the ground-truth segmentation mask at ES, noted as M_{ES} , to obtain $M_{ES} \circ W_{FES}$. $M_{ES} \circ W_{FES}$ is then compared with M_{ED} , the corresponding ground-truth masks at ED, using Dice coefficient (2D version) on LVM, LVC and RVC. Overall, the means(standard deviations) of Dice coefficients achieved on LVM, LVC and RVC in the 5-fold cross-validation are reported in Table 3.

Additionally, we also visually evaluate the apparent flow generated by the ApparentFlow-net. We find that the apparent flow is indeed good enough to characterize the cardiac motion of the typical cases in the pathological categories. Several examples are given in Appendix D.

4.1.3. Importance of supervision in training

In order to understand the importance of the small amount of segmentation ground-truth used in the proposed semi-supervised learning method, we also train a variant of ApparentFlow-net using only unsupervised learning. The only modification is the removal of the term $L_{GT}(\mathbf{F}_t)$ from $L_{\text{flow}}(\mathbf{F}_t)$ such that the variant is trained without any ground-truth for supervision. As reported in Table 3, the means of Dice coefficients on LVM, LVC and RVC are all lower than the corresponding results achieved by the semi-supervised learning method. In particular, there is a large drop from 0.84 to 0.76 on the mean of Dice coefficient on LVM. So the proposed semi-supervised learning method is indeed better than its unsupervised learning variant by making efficient use of the small amount of segmented images.

4.2. Finetuning LVRV-net

LVRV-net is already trained in Zheng et al. (2018) for 80 epochs on a subset of about 3000 cases of UK Biobank (Petersen et al. (2016)). In the training process with the whole ACDC training set, as well as in each of the 5 training processes of the 5-fold cross validation, LVRV-net is finetuned for 920 epochs on the corresponding training data, with exactly the same loss function and training parameters as given in Zheng et al. (2018). With the finetuning, the means (standard deviations) of 3D Dice coefficients achieved on LVC, LVM and RVC segmentation volumes in the 5-fold cross-validation are 0.94(0.06), 0.90(0.03) and 0.89(0.12).

Table 4

The classification performance on the testing set (50 cases) and training set (100 cases) of ACDC by different models.

Model	Testing Set Accuracy	Training Set Accuracy	Evaluation Method on Training Set
proposed model	94%	95%	5-fold cross-validation
Isensee et al. (2017)	92%	94%	5-fold cross-validation
Wolterink et al. (2017)	86%	91%	4-fold cross-validation
Cetin et al. (2017)	92%	100%	forward feature selection and leave-one-out cross-validation
Khened et al. (2017)	96%	90%	70 cases for training, 20 for validation, 10 for evaluation
Khened et al. (2018)	100%	N.A.	N.A.

4.3. Proposed classification model

Apparent flows and segmentation masks are generated by the ApparentFlow-net and the finetuned LVRV-net, from which the 7 shape-related features and the 2 motion-characteristic features are extracted. Then the 4 ridge logistic regression binary classifiers are implemented using Scikit-learn Pedregosa et al. (2011) and trained on the cases of the categories they are supposed to classify. For example, DCM classifier is trained on the cases of NOR, MINF and DCM; the cases of RVA or HCM are not used to train it. In terms of classifier parameter, we empirically find that $C = 50$ works well and use it in this paper. The performances of some variants with different values of C are provided in Appendix B.

4.3.1. Classification performance

As presented in Table 4, on the testing set, the accuracy of our model is 94%. In the 5-fold cross-validation on the training set, our method achieves an accuracy of 95%. Hence our model achieves performances that are comparable to those of the state-of-the-art on both the training set and the testing set. This is quite remarkable because, in contrast to the state-of-the-art, each classifier in our model uses only up to three features and has only up to 4 parameters. In total, our model uses 9 features and has 14 parameters. And each feature is selected in a clearly explainable manner. On the testing set, among the two methods with performances better than ours, Khened et al. (2017) uses a random forest of 100 trees and Khened et al. (2018) applies a more sophisticated ensemble system. Therefore, those classification models are less straightforward to interpret than ours. Furthermore, since our model has very similar performances on the training and testing sets, there seems to be little overfit.

Based on the confusion matrix of the prediction in the 5-fold cross-validation on the ACDC training set (Fig. 7), for the binary classification of NOR, RVA, HCM, DCM and MINF, we calculate and find that the precision values are 0.87, 1.00, 1.00, 0.91 and 1.00; the recall values are 1.00, 0.90, 0.95, 1.00 and 0.90.

4.3.2. Interpretation of mis-Classification

As our classifier can be interpreted easily, we figure out for each of the 5 misclassified cases (Fig. 7) why the prediction is different

		Predicted				
		NOR	RVA	HCM	DCM	MINF
Ground-Truth	NOR	20	0	0	0	0
	RVA	2	18	0	0	0
	HCM	1	0	19	0	0
	DCM	0	0	0	20	0
	MINF	0	0	0	2	18

Fig. 7. The confusion matrix of the predictions by the proposed classification model in the 5-fold cross-validation on the training set of ACDC.

from the ground-truth. In fact, they all seem to be somewhat ambiguous in terms of pathological category:

- Patients 082 and 088 are of ground-truth RVA but are classified as NOR. According to our segmentation, they have $V_{RVC,ED}$, EF_{RVC} and $R_{RVCLV,ED}$ values all very similar to that of the NOR cases. For instance, they have the third and the first lowest $R_{RVCLV,ED}$ values (0.755 and 0.691) among all the RVA cases, which are well in the range of that of the NOR cases.
- Patient 022 is of ground-truth HCM but is predicted as NOR. Unlike all the other HCM cases, patient 022 has both EF_{LVC} (0.622) and $MT_{LVM,ED}$ (14.7mm) in the normal ranges, which makes it look like a NOR case.
- Patients 050 and 060 are of ground-truth MINF but are predicted as DCM. Their values of $V_{LVC,ES}$ (118.0mL/m² and 83.5mL/m²) are the two highest among all the non-DCM cases and well in the range of that of the DCM cases. In terms of motion disparity, on RMD and TMD , unlike the majority of the MINF cases, their values (0.245 and 1.173 for patient 050, 0.316 and 1.246 for patient 060) are also in the ranges of that of the DCM cases. For these reasons, the DCM classifier predicts them to be DCM cases.

4.3.3. Explaining the classifiers

The 4 binary classifiers are just logistic regression models. As presented in the previous section, their prediction depends on the sign of the sum $\sum_i p_i f_{i,i} + b$. To understand what is learned from the data by the trained classifiers, in Table 5 we show the coefficients of the classifiers trained with all the relevant cases in ACDC. We find that the signs of the parameters p_i all correspond to the positive or negative correlation between the feature and the binary classification task. For instance, in the trained RVA classifier, the signs of the coefficients of $V_{RVC,ED}$ and $R_{RVCLV,ED}$ are both positive, as a large RVC volume and a high ratio between the RVC and LV volumes are both indicators of RV abnormality; on the other hand, since low RVC ejection fraction usually signifies RV abnormality, the coefficient of EF_{RVC} is negative. Similarly, such a correspondence applies to all the coefficients of the 3 other trained classifiers. In particular, for MINF classifier, the learned threshold on EF_{LVC} to distinguish MINF cases from NOR cases is 7.994/17.122 = 0.467, which can well separate them according to their definitions.

Table 5

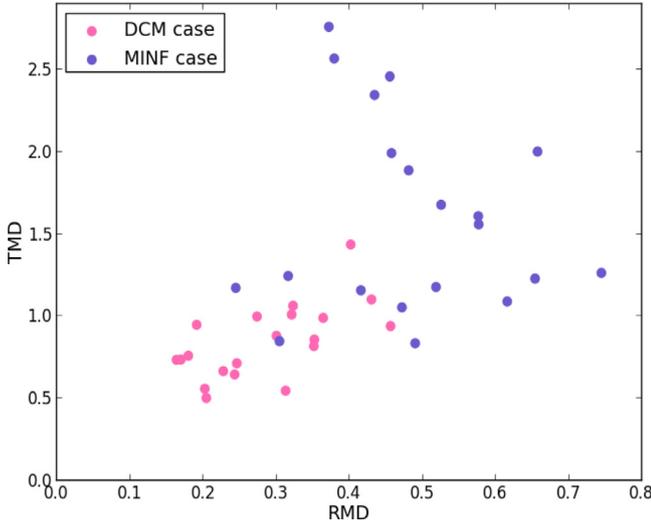
The parameters of the 4 logistic regression binary classifiers trained on the training set of ACDC.

	Parameters of the Trained Classifier
RVA Classifier	$0.010V_{RVC,ED} - 4.695EF_{RVC} + 14.012R_{RVCLV,ED} - 9.906$
HCM Classifier	$8.434EF_{LVC} + 4.614R_{LVM,ED} + 0.420MT_{LVM,ED} - 16.580$
DCM Classifier	$0.104V_{LVC,ES} - 0.918RMD - 7.758TMD - 0.321$
MINF Classifier	$-17.122EF_{LVC} + 7.994$

Table 6

The performance of the variants of DCM classifier on the training set of ACDC.

DCM Classifier Input	# of Mis-Classification on the 60 DCM, MINF and NOR cases
$V_{LVC,ES}$, RMD , TMD (proposed)	2
$V_{LVC,ES}$, TMD	2
$V_{LVC,ES}$, RMD	3
$V_{LVC,ES}$	4

**Fig. 8.** The motion-characteristic features (RMD and TMD) of the DCM and MINF cases in the training set of ACDC. The majority of the cases are well separable with these two features.

4.4. Variants of the proposed classification model

We compare the proposed classification model with its variants for a justification of our design and a more comprehensive understanding of the model.

4.4.1. Importance of motion-Characteristic features

To better understand the value of the two proposed motion-characteristic features, we further train three variants of DCM classifier which use zero or one motion-characteristic feature as input. And the set of input features is the only difference between these models. As shown in Table 6, on the 60 cases of NOR, MINF and DCM, while DCM classifier makes only two errors, the variant using only shape-related feature $V_{LVC,ES}$ misclassifies 4 cases. But improvements can be made by using at least one motion-characteristic feature. As can be visualized in Fig. 8, the motion characteristic features RMD and TMD allow the separation of the majority of the cases of DCM and MINF. Combining them with the shape-related feature $V_{LVC,ES}$ together as the input, the DCM classifier can make more accurate classification.

4.4.2. Proposed model on non-normalized features

We test whether BSA normalization is required for our model to achieve high performance. Among the 9 proposed features, only the values of $V_{RVC,ED}$ and $V_{LVC,ES}$ would be different without BSA-normalization. And only RVA and MINF classifiers which use these two features as input would be affected. As presented in Table 7, without BSA-normalization on the features, the 5-fold cross validation accuracy on ACDC training set only drops a little bit to 94%. The proposed model still remains accurate.

4.4.3. Proposed model with inversed classifier order

As presented previously, the 4 classifiers in the proposed model are arranged according to the estimated difficulties of the corre-

sponding classification task. To confirm the importance of this order, we create another model by inverting the order of the classifiers. So, unlike the proposed model shown in Fig. 2, in this variant, a case goes through successively MINF, DCM, HCM and RVA classifiers instead. As shown in Table 7, the accuracy of this variant is quite low. Hence the proposed order of the classifiers is indeed important.

4.4.4. Variants with other classifier models

We replace the proposed logistic regression classifiers with 3 other types of classifiers on the same sets of input features, including Lasso, LassoCV (Lasso with model selection by cross-validation) and random forest. Details of these models are available in Appendix C. As reported in Table 7, their performances are clearly below that of the proposed model. Our choice of logistic regression as the classifier model is hence justified.

4.4.5. Variants without manual input feature selection

To evaluate if the manual feature selection is useful for the model to be accurate, we train several modified versions of the proposed model without manual feature selection. They all consist of 4 classifiers to perform the same binary classification tasks as in the proposed model. But each of the 4 classifiers of these variants takes all 9 features together as input. In total, we implement 6 models with the following models as their classifiers respectively (details of these models are available in Appendix C): support vector machine (SVM), relevance vector machine (RVM), Lasso, LassoCV, random forest and high dimensional discriminant analysis (HDDA) model.

As reported in Table 7, on the BSA-normalized features as well as on the non-normalized features, they all have accuracy lower than that of the proposed model by at least 6%. This justifies the necessity of manual feature selection, at least on a relatively small dataset like ACDC. We are not yet able to examine the importance of manual feature selection on large datasets.

To better understand the roles of the features, we further examine the variant with random forest classifiers without manual feature selection trained on the 100 cases of ACDC training set. For each of the 4 classifiers, we compute the feature importance for each of the 9 features to determine the most important one. The importance of a feature is defined as the total reduction of the entropy brought by that feature in all the trees in the random forest. We find that for RVA, HCM and DCM classifiers, the most important features are $R_{RVCLV,ED}$, $R_{LVMLVC,ED}$ and $V_{LVC,ES}$ respectively, which are among the features manually selected for the corresponding proposed classifiers. For MINF classifier, the two most important features are $R_{RVCLV,ED}$ and EF_{LVC} , which have roughly the same importance (0.35 and 0.32). Only EF_{LVC} is used in the proposed model according to its direct relevance. These observations provide further support for our manual feature selection.

4.4.6. Variants without binary classification

The proposed model divides the 5-category classification task into 4 binary classification sub-tasks. In order to understand whether this special design contributes to the achieved high accuracy, we train and evaluate 2 variants on the same set of 9 features. A random forest and a multi-layer perceptron (MLP) are

Table 7

The 5-fold cross validation accuracy on the training set of ACDC of the variants of the proposed classification model .

Method	BSA-Normalized Features	Non-Normalized Features
logistic regression classifiers (proposed model)	95%	94%
logistic regression classifiers in inversed order	63%	64%
Lasso classifiers	89%	91%
LassoCV classifiers	80%	81%
random forest classifiers	85%	87%
logistic regression classifiers w/o manual feature selection	88%	88%
SVM classifiers	87%	84%
w/o manual feature selection		
RVM classifiers	88%	72%
w/o manual feature selection		
Lasso classifiers	85%	86%
w/o manual feature selection		
LassoCV classifiers	84%	87%
w/o manual feature selection		
random forest classifiers	86%	88%
w/o manual feature selection		
HDDA classifiers	49%	46%
w/o manual feature selection		
one single random forest	87%	88%
w/o binary classification		
one single MLP	84%	84%
w/o binary classification		

respectively trained to predict a case to be one of the five categories directly without binary classification (details of these models are available in [Appendix C](#)). As reported in [Table 7](#), their performances are not as good as that of the proposed model. Hence the strategy of performing a series of binary classification makes sense.

5. Conclusion and discussion

We propose a method of cardiac pathology classification based on originally designed and trained neural networks and classifiers¹. A novel semi-supervised training method is applied to train ApparentFlow-net which provides pixel-wise motion information. Combining the apparent flow generated by ApparentFlow-net and the segmentation masks predicted by LVRV-net, we introduce two novel features that characterize the motion of myocardial segments. These motion-characteristic features are not only intuitive for visualization but also very valuable in classification. The proposed classification model consists of 4 small binary classifiers. Each classifier works independently and takes up to 3 features with clearly explainable relevance as input. On ACDC training set and testing set, the proposed model achieves 95% and 94% respectively as classification accuracy. Its performance is hence comparable to that of the state-of-the-art. To justify our design of the proposed classification model, we also quantitatively compare it with other models.

The apparent flow generated by ApparentFlow-net and the originally designed time series of myocardial segment motion are not only straightforward to understand but also useful for classification. We believe that making the automatic methods more understandable and explainable is important, as it is not only helpful to facilitate the implementation and application of the research of medical image analysis in clinics but also useful to improve transparency and to gain trust in medical practice ([Holzinger et al. \(2017\)](#), [Rueckert et al. \(2016\)](#)).

Furthermore, the motion information we extract from the apparent flow is fairly rich. We believe that ApparentFlow-net may be a powerful tool of motion extraction for the community. The way we extract the time series and the motion-characteristic features from the flow maps is just one of the so many possible applications. Also, ApparentFlow-net is trained in a semi-supervised manner. This training approach is highly relevant to the current situation of data availability in medical image analysis, as we usually have access to a relatively large amount of unlabeled data and a relatively small amount of labeled data. In a word, much more potential applications in various circumstances of apparent flow are yet to be explored.

Regarding the extraction of the motion-characteristic features, one could use the segmentation network to segment all frames and then derive the motion-characteristic features from the segmentation masks. However, we find that the resulting time series characterizing the cardiac motion (e.g. the time series of the radius and thickness as shown in the second column of [Fig. 6](#)) by this approach are not as temporally consistent as we would expect. In fact, the segmentation network was trained to segment the frames at ED and ES only. And no constraint has ever been imposed to make the segmentation masks temporally consistent. The problem would be clearer if we look at the two frames in the first column in [Fig. 5](#). While the ED frame (upper image) is easy to segment, the other frame (lower image) appears to be more challenging due to the presence of massive trabeculations. Moreover, as the segmentation network segments the two frames independently, it is not obvious how to ensure the consistency of the segmentation masks. This problem can be solved using the ApparentFlow-net instead to extract motion. As shown in the second column of [Fig. 6](#), with the ApparentFlow-net, the extracted time series of the radius and thickness of the segments are reasonably smooth, which reflects the enforced temporal consistency.

We could have used existing traditional registration models to supervise the training of ApparentFlow-net or even replace ApparentFlow-net by a deformable registration algorithm (e.g. LD-DMM, LCC-Demons, etc.). However, we notice that in order to make the traditional registration models work reasonably well on an unseen dataset like ACDC, the estimation and finetuning of key

¹ The code and the model will be available in this repository: <https://github.com/julien-zheng/CardiacMotionFlow>

parameters in these models are necessary. For instance, the authors of Krebs et al. (2019) empirically estimate the key parameters of LCC-Demons (Lorenzi et al. (2013)) and SyN (Avants et al. (2008)) before applying them on ACDC. Our method is simpler in the sense that it learns everything from data and requires no manual model/parameter estimation/adjustment. Hence, on the one hand, our method is easier and more convenient to be applied to various datasets that are reasonably similar to the training set. On the other hand, it allows us to take advantage of the increasing number of data available in the community. We believe a method with these advantages is very interesting and worth trying. Moreover, as far as we observe, our registration method is accurate enough to characterize the cardiac motion. Some examples are provided in Fig. D.1 to show that the generated apparent flow characterizes the motion patterns of typical cases in several pathological categories. And as shown in the second column of Fig. 6, with the apparent flow generated by the ApparentFlow-net, the extracted time series of the radius and thickness of the segments enable us to easily distinguish the typical cases of different cardiac pathologies.

A straightforward comparison with prior works on 2D registration methods on the ACDC dataset shows that the ApparentFlow-net performs rather well. Indeed when looking at the Dice coefficients achieved on LVC and RVC, our approach leads to 0.94 and 0.87 respectively (see Table 3). In Hering et al. (2019), the authors describe a learning-based method leading to Dice coefficients at best equal to 0.90 on the same structures (based on Fig. 3 of Hering et al. (2019)) and also performances of a non-learning-based method similar to Rühhaak et al. (2013) with at most 0.80 of Dice. Note however that in this comparison, differences on cross-validation (5-fold v.s. 10-fold), slice selection and ROI determination may hinder the analysis.

While analyzing and extracting the cardiac motion, we adopt a 2D slice-by-slice processing method, without taking the motion on neighboring slices into account. The reason behind this choice is the fact that the inter-slice distance in the short-axis MRIs in ACDC is quite large. Usually, the inter-slice distance between two adjacent slices in MRI stacks is 5 to 10mm. The heart may hence have obviously different shape and motion even on two adjacent slices. In this case, ignoring the neighboring slices for motion estimation might be reasonable. However, if our method is to be applied on some volumetric images with small inter-slice distance, a modification of the approach by taking neighboring slices into account might be beneficial.

An issue that would hinder the generalization of pathology classification models like ours is the lack of a standard and universal definition of pathological category Suinesiaputra et al. (2016). For instance, there is another public dataset made available for the MICCAI 2009 challenge on automated LV segmentation Radau et al. (2009) (the dataset is also known as the Sunnybrook dataset (SD)) containing pathological cases. The 4 categories of SD are heart failure with infarction, heart failure without infarction, LV hypertrophy and healthy. While a hypertrophic case in ACDC has a LV cardiac mass over $110\text{g}/\text{m}^2$ and several myocardial segments of thickness over 15mm at ED by definition, the hypertrophic cases according to SD definition only need to have a LV cardiac mass over $83\text{g}/\text{m}^2$. And no threshold is proposed for the myocardial segment thickness by the SD definition. In fact, we find multiple cases in SD which are of LV cardiac mass between $83\text{g}/\text{m}^2$ and $110\text{g}/\text{m}^2$ and maximal segment thicknesses well below 15mm. They are identified as hypertrophic cases in SD. But they would not be considered as hypertrophic at all according to ACDC. Similarly, the category of infarction is defined differently in SD and ACDC. In SD, the infarction is determined by the evidence of late gadolinium enhancement; abnormal cardiac motion might not be observable. Yet in ACDC, the infarction category is defined by the presence of abnormal motion. With such discrepancies between the definitions

in different datasets, it is difficult for the community to train a classification model on a dataset such that it generalizes well to the others. We hence appeal for more attention on this issue.

Another issue that may limit the generalization of our classification model is the small size of the ACDC dataset used for training. ACDC training set has only 100 cases of 5 pathological categories. Moreover, in each pathological category, there are only 20 cases. Consequently, on the one hand, many pathological categories are not included in ACDC. On the other hand, for each of the 5 pathological categories in ACDC, we would expect that the 20 cases might not be enough to represent all cases in the category. In order to achieve good generalization, we may need larger datasets with more pathological categories to train the model.

Also, notice that the proposed simple classification model of only 14 parameters is somewhat specific to the ACDC dataset. If we need to adapt our model to perform classification on a larger dataset with more pathological categories, it may be necessary to increase the size and hence the number of parameters of the model.

Finally, we would like to point out that although some single-value hand-crafted motion-characteristic features (e.g. *RMD* and *TMD*) are used in this paper, we believe that for some pathology it would be better to use the whole time series of segment radius or thickness as input to a classification model. For instance, if we aim to discover subtler characteristics related to the motion (e.g. dyssynchrony, septal flash) from a larger dataset, doing so might become appropriate and necessary. We expect to carry out research on this topic in the future.

Conflict of Interest Statement

We wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed. We further confirm that the order of authors listed in the manuscript has been approved by all of us.

We confirm that we have given due consideration to the protection of intellectual property associated with this work and that there are no impediments to publication, including the timing of publication, with respect to intellectual property. In so doing we confirm that we have followed the regulations of our institutions concerning intellectual property.

We understand that the Corresponding Author is the sole contact for the Editorial process (including Editorial Manager and direct communications with the office). He is responsible for communicating with the other authors about progress, submissions of revisions and final approval of proofs. We confirm that we have provided a current, correct email address which is accessible by the Corresponding Author and which has been configured to accept email from qiao.zheng@inria.fr.

Acknowledgments

The authors acknowledge the partial support from the European Research Council (MedYMA ERC-AdG-2011-291080).

Appendix A. Loss Function for Training ApparentFlow-Net

To penalize the crossing or large rotations of flows, we compute the difference between of the warped x -components (resp. y -components) of each pair of horizontally (resp. vertically) adjacent pixels \mathbf{P}^{x+} and \mathbf{P} (resp. \mathbf{P}^{y+} and \mathbf{P}). There is a crossing if and only if

this difference is smaller than 0, for which a penalty which is equal to the square of this difference applies. Otherwise no penalty applies. Hence we come up with the term $L_{\text{CROSS}}(\mathbf{F}_t)$ to penalize the crossing of flows:

$$\begin{aligned} L_{\text{CROSS}}(\mathbf{F}_t) &= \sum_{\mathbf{P}} \min((x + 1 + F_t^x(\mathbf{P}^{x+})) - (x + F_t^x(\mathbf{P})), 0)^2 \\ &\quad + \sum_{\mathbf{P}} \min((y + 1 + F_t^y(\mathbf{P}^{y+})) - (y + F_t^y(\mathbf{P})), 0)^2 \\ &= \sum_{\mathbf{P}} \min\left(1 + \frac{\partial F_t^x(\mathbf{P})}{\partial x}, 0\right)^2 + \min\left(1 + \frac{\partial F_t^y(\mathbf{P})}{\partial y}, 0\right)^2 \quad (\text{A.1}) \end{aligned}$$

in which $\partial F_t^x(\mathbf{P})/\partial x$ is computed with finite difference as $F_t^x(\mathbf{P}^{x+}) - F_t^x(\mathbf{P})$ (and similarly for $\partial F_t^y(\mathbf{P})/\partial y$).

The Dice function in the term $L_{\text{CT}}(\mathbf{F}_{\text{ES}})$ is defined on two images U and V as described in Zheng et al. (2018)

$$\text{Dice}(U, V) = -\frac{2 \sum_{\mathbf{P}} U(\mathbf{P})V(\mathbf{P}) + \epsilon}{\sum_{\mathbf{P}} U(\mathbf{P}) + \sum_{\mathbf{P}} V(\mathbf{P}) + \epsilon} \quad (\text{A.2})$$

with $\epsilon = 1$ a term for better numerical stability in training.

Appendix B. Variants of the Proposed Classification Model with Different Values of Parameter C

We also perform 5-fold cross-validation on the ACDC training set for the variants of the proposed classification model by varying the parameter C in the 4 logistic regression classifiers. Their performances are reported in Table B.1.

Table B1

The 5-fold cross-validation performance on the ACDC training set of some variants of the proposed classification model with various values of parameter C .

C	Training Set Accuracy
1	76%
5	88%
10	92%
50	95%
100	95%
500	93%
1000	93%
5000	93%

Appendix C. Variants of the Proposed Classification Model with Different Classifiers and Input Features

C1. Variants with other classifier models

We replace the proposed ridge logistic regression classifiers by other types of classifiers on the same sets of input features:

- Lasso classifiers: in this variant, each of the 4 classifiers is a least absolute shrinkage and selection operator (Lasso). The constant α that multiplies the L_1 term is empirically chosen to be 10^{-4} .
- LassoCV classifiers: each of the 4 classifiers is a Lasso model with model selection by cross-validation (LassoCV). The optimal constant α is searched in the range $[10^{-4}, 10^{-0.5}]$ in a 4-fold cross-validation on the training data.
- random forest classifiers: each of the 4 classifiers is a random forest of 1000 trees which expand their nodes in training until all leaves are pure or all leaves contain less than 2 samples. Entropy is used to measure the quality of a split in training.

C2. Variants without manual input feature selection

We train several variants of the proposed model without manual feature selection. They all consist of 4 classifiers arranged in the same order to perform the same binary classification tasks as in the proposed model. But each of the 4 classifiers in these variants takes all the 9 features together as input. In total, we implement and examine 6 variants with the following models as their classifiers respectively:

- Variant with SVM classifiers: each of the 4 binary classifiers is a support vector machine (SVM) with linear kernel and penalty parameter $C=50$.
- Variant with RVM classifiers: each of the 4 binary classifiers is a relevance vector machine (RVM) as introduced in Tipping and Faul (2003) with linear kernel.
- Variant with Lasso classifiers: each of the 4 binary classifiers is a Lasso. Lasso is known as a model capable of performing both variable selection and regularization. α , the constant that multiplies the L_1 term, is empirically set to 10^{-4} .
- Variant with LassoCV classifiers: each of the 4 binary classifiers is a Lasso with model selection in a 4-fold cross-validation on the training data. The optimal constant α is searched in the range $[10^{-4}, 10^{-0.5}]$.
- Variant with random forest classifiers: in this variant, each of the 4 binary classifiers is a random forest of 1000 trees which expand their nodes in training until all leaves are pure or all leaves contain less than 2 samples. Entropy is used to measure the quality of a split in training.
- Variant with HDDA classifiers: each of the 4 binary classifiers is a high dimensional discriminant analysis (HDDA) model, which is an expectation-maximization algorithm designed for high-dimensional data clustering based on the ideas of dimension reduction and parsimonious modeling (Bouveyron et al. (2007), Orlhac et al. (2018)). Though the 9-feature space in this paper is not high dimensional, we show the performance of such a sophisticated method for comparison.

C3. Variants without binary classification

We train and evaluate the following 2 variants on all the 9 input features. These variants are obtained by replacing the 4 binary classifiers with a single multi-class one:

- Variant using random forest: it is a random forest of 1000 trees which expand their nodes in training until all leaves are pure or all leaves contain less than 2 samples. Entropy is used to measure the quality of a split in training.
- Variant using MLP: it is a multi-layer perceptron (MLP). It has 2 hidden layers of 32 neurons with tanh activation function. Adam optimizer is used to train it for 10^5 epochs with learning rate 0.001.

C4. Implementation of the variants

Among the above variants of the proposed classification model with different classifiers and input features, the HDDA classifiers are implemented using the HDDA python toolbox downloaded from the GitHub page <https://github.com/mfauvel/HDDA>, the RVM classifiers are implemented in Python according to the method described in Tipping and Faul (2003), and all the other variants are implemented with Scikit-learn.

Appendix D. Examples of Apparent Flow Generated by the ApparentFlow-net

We provide 4 examples of the apparent flow generated by the ApparentFlow-net of 4 ACDC training set cases in different

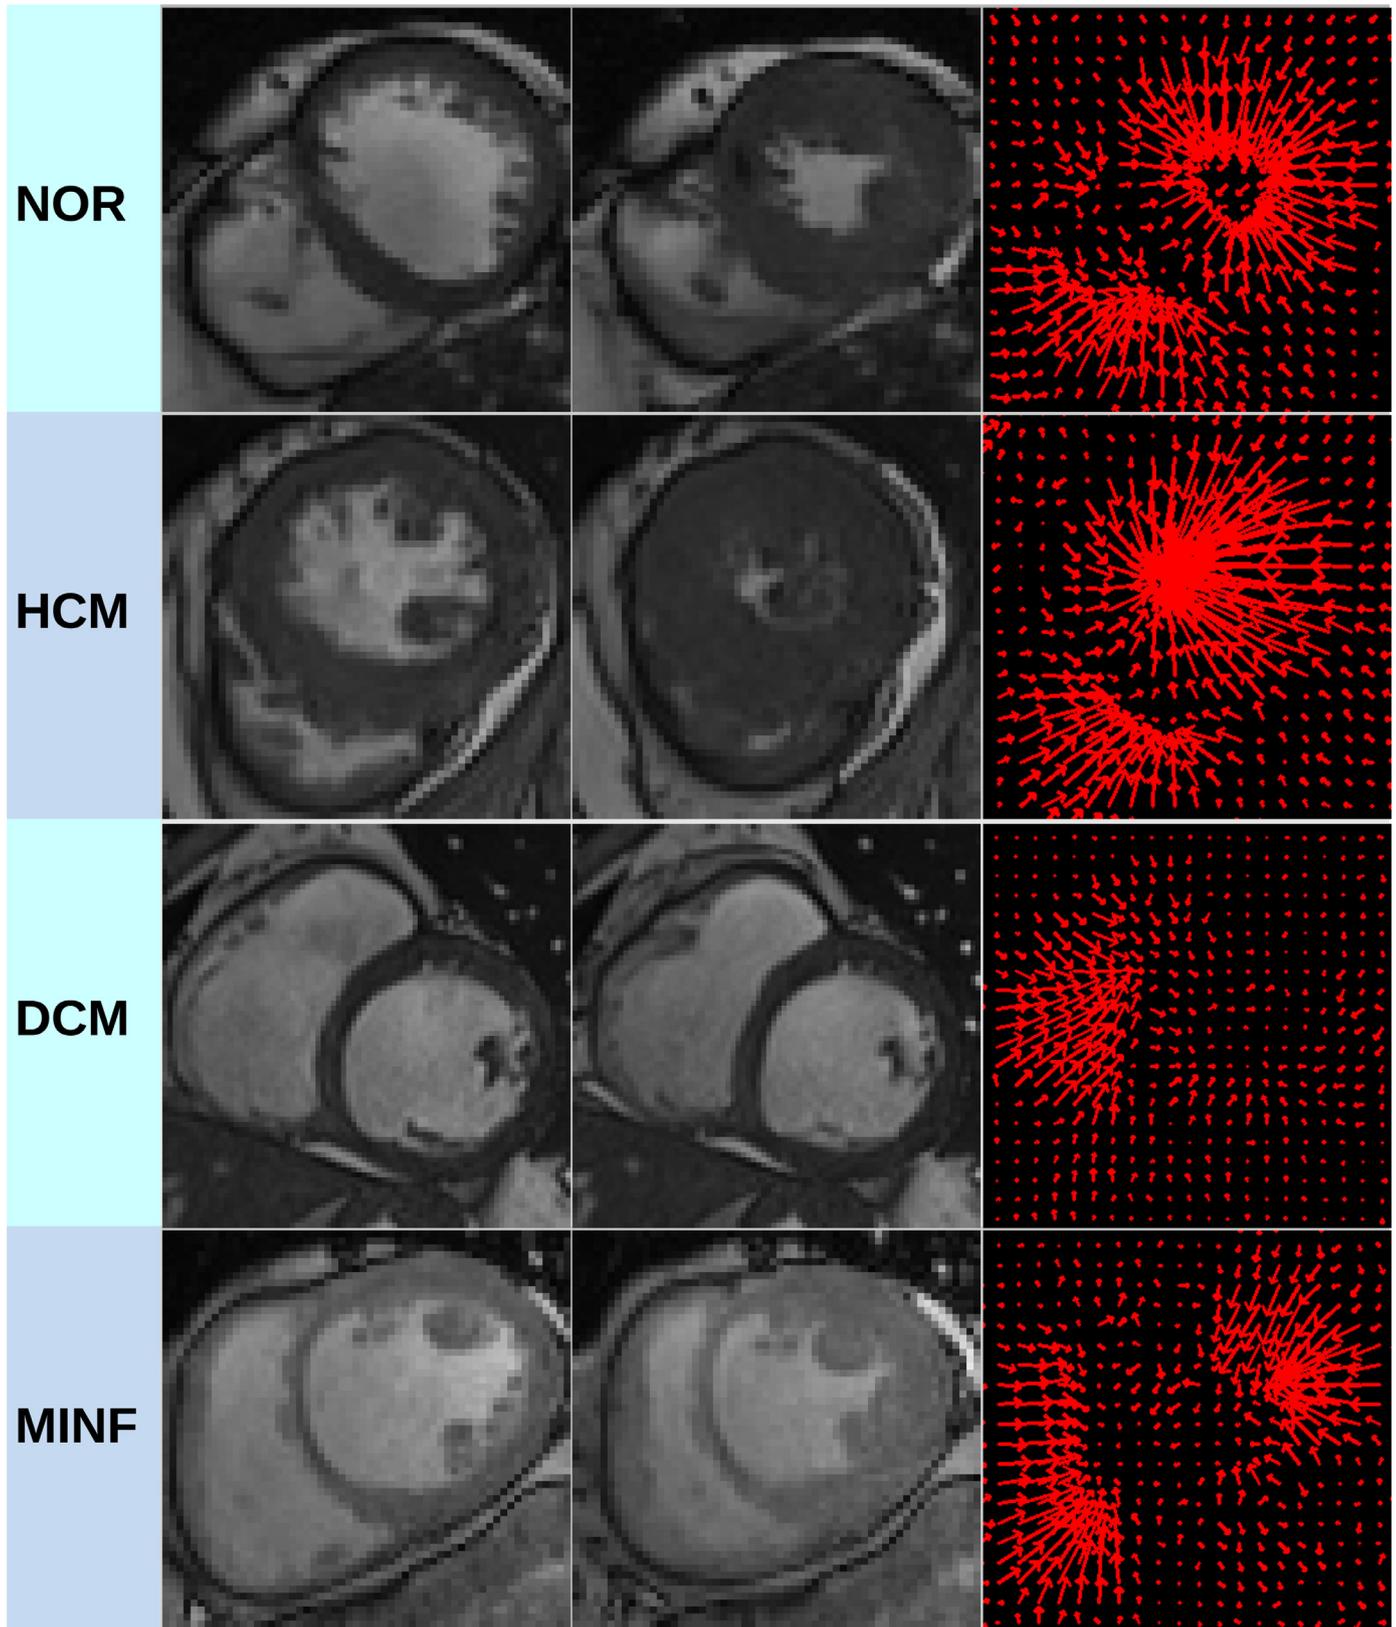


Fig. D1. Four examples of the apparent flow generated by the ApparentFlow-net of 4 ACDC training set cases in different pathological categories. The images in the first column are the frames at ED. The images in the second column are the frames around ES. The apparent flow maps corresponding to the pairs of frames in the first and second columns are shown in the third column. The apparent flow can indeed characterize the cardiac motion of the typical cases in the pathological categories. NOR: synchronous and uniform flow on LVM; HCM: excessively large flow on LVM; DCM: very small flow on LVM; MINF: asynchronous and ununiform flow on LVM.

pathological categories. In Fig. D.1, given the frames at ED (first column) and the frames around ES (second column), we apply the trained ApparentFlow-net to generate the apparent flow maps (third column).

Visually, we find that the apparent flow can indeed characterize the cardiac motion of the typical cases in the pathological categories. As expected, the apparent flow on the LVM of a NOR case is oriented along the gradient of the image intensity and has roughly the same amplitude throughout the left ventricle, signifying the synchronous and uniform contraction and thickening of the LVM for the NOR case. For a HCM cases, we can see that the flow on LVM is excessively large, which means that the contraction and thickening is excessive, a typical phenomenon we find on HCM cases. Conversely, the flow on the LVM of a DCM case is small since the hearts of DCM cases usually do not contract or thicken enough. Finally, the flow on the LVM of a MINF case is not uniform: some myocardial segments contract and thicken much less than the others. This is a typical symptom that we can find on MINF cases.

References

- Avants, B., Epstein, C., Grossman, M., Gee, J., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* 12(1), 26–41.
- Bai, W., Oktay, O., Sinclair, M., Suzuki, M., Rajchl, M., Tarroni, G., Glocker, B., King, A., Matthews, P., Rueckert, D., 2017. Semi-supervised learning for network-based cardiac MR image segmentation. *MICCAI* 253–260.
- Balakrishnan, G., Zhao, A., Sabuncu, M., Guttag, J., Dalca, A., 2018. VoxelMorph: a learning framework for deformable medical image registration. *arXiv preprint arXiv:1809.05231*.
- Bernard, O., Lalonde, A., Zotti, C., Cervensky, F., Yang, X., Heng, P., Cetin, I., Lekadir, K., Camara, O., Ballester, M., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V., Krishnamurthi, G., Rohé, M., Pennec, X., Sermesant, M., Isensee, F., Jäger, P., Maier-Hein, K., Baumgartner, C., Koch, L., Wolterink, J., Isgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P., 2018. Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE Trans. Med. Imaging* 37(11), 2514–2525.
- Bouveyron, C., Girard, S., Schmid, C., 2007. High dimensional data clustering. *Comput. Stat. Data Anal.* 52, 502–519.
- Cerqueira, M., Weissman, N., Dilsizian, V., Jacobs, A., Kaul, S., Laskey, W., Pennell, D., Rumberger, J., Ryan, T., Verani, M., 2002. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: a statement for healthcare professionals from the cardiac imaging committee of the council on clinical cardiology of the american heart association. *Circulation*.
- Cetin, I., Sanroma, G., Petersen, S., Napel, S., Camara, O., Ballester, M., Lekadir, K., 2017. A radiomics approach to computer-aided diagnosis with cardiac cine-MRI. In: *Proc. Statistical Atlases and Computational Models of the Heart (STACOM), ACDC challenge, MICCAI'17 Workshop*.
- Cheplygina, V., de Bruijne, M., Pluim, J., 2018. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *arXiv preprint arXiv:1804.06353*.
- Comaniciu, D., Engel, K., Georgescu, B., Mansi, T., 2016. Shaping the future through innovations: from medical imaging to precision medicine. *Med. Image Anal.* 33, 19–26.
- Dawes, T., de Marvao, A., Shi, W., Fletcher, T., Watson, G., Wharton, W., Rhodes, C., Howard, L., Gibbs, J., Rueckert, D., Cook, S., Wilkins, M., D. O., 2017. Machine learning of three-dimensional right ventricular motion enables outcome prediction in pulmonary hypertension: a cardiac MR imaging study. *Radiology* 283, 381–390.
- Gao, B., Liu, W., Wang, L., Liu, P., Croisille, P., Delachartre, P., Clarysse, P., 2016. Estimation of cardiac motion in cine-MRI sequences by correlation transform optical flow of monogenic features distance. *Phys. Med. Biol.* 61, 8640–8663.
- Gilbert, K., Pontre, B., Occlshaw, C., Cowan, B., Suinesiaputra, A., Young, A., 2017. 4D modelling for rapid assessment of biventricular function in congenital heart disease. *Int. J. Cardiovasc. Imaging* 34(3), 407–417.
- Gu, L., Zheng, Y., Bise, R., Sato, I., Imanishi, N., Aiso, S., 2017. Semi-supervised learning for biomedical image segmentation via forest oriented super pixels (voxels). *MICCAI* 702–710.
- Hering, A., Kuckertz, S., Heldmann, S., Heinrich, M., 2019. Enhancing label-driven deep deformable image registration with local distance metrics for state-of-the-art cardiac motion tracking. *Bildverarbeitung für die Medizin* 2019 309–314.
- Hernandez, M., Olmos, S., Pennec, X., 2008. Comparing algorithms for diffeomorphic registration: stationary LDDMM and diffeomorphic demons. *Proc. MFCA* 2008 24–35.
- Holzinger, A., Biemann, C., Pattichis, C., Kell, D., 2017. What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923*.
- Isensee, F., Jaeger, P., Full, P., Wolf, I., Engelhardt, S., Maier-Hein, K., 2017. Automatic cardiac disease assessment on cine-MRI via time-series segmentation and domain specific features. In: *Proc. Statistical Atlases and Computational Models of the Heart (STACOM), ACDC challenge, MICCAI'17 Workshop*.
- Khened, M., Alex, V., Krishnamurthi, G., 2017. Densely connected fully convolutional network for short-axis cardiac cine MR image segmentation and heart diagnosis using random forest. In: *Proc. Statistical Atlases and Computational Models of the Heart (STACOM), ACDC challenge, MICCAI'17 Workshop*.
- Khened, M., Alex, V., Krishnamurthi, G., 2018. Fully convolutional multi-scale residual densenets for cardiac segmentation and automated cardiac diagnosis using ensemble of classifiers. *arXiv preprint arXiv:1801.05173*.
- Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T., 2019. Learning a probabilistic model for diffeomorphic registration. *IEEE Trans. Med. Imaging*.
- Krebs, J., Mansi, T., Mailhé, B., Ayache, N., Delingette, H., 2018. Unsupervised probabilistic deformation modeling for robust diffeomorphic registration. In: *Proc. Deep Learning in Medical Image Analysis (DLMIA), MICCAI'18 Workshop*.
- Li, H., Fan, Y., 2017. Non-rigid image registration using fully convolutional networks with deep self-supervision. *arXiv preprint arXiv:1709.00799*.
- Lorenzi, M., Ayache, N., Frisoni, G., Pennec, X., 2013. LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage* 81, 470–483.
- Lu, A., Parajuli, N., Zontak, M., Stendahl, J., Ta, K., Liu, Z., Boutagy, N., Jeng, G., Alkhalil, I., Staib, L., O'Donnell, M., Sinusas, A., Duncan, J., 2018. Learning-based regularization for cardiac strain analysis with ability for domain adaptation. *arXiv preprint arXiv:1807.04807*.
- Orlhac, F., Mattei, P., Bouveyron, C., Ayache, N., 2018. Class-specific variable selection in high-dimensional discriminant analysis through bayesian sparsity. *preprint HAL* 01811514.
- Parajuli, N., Lu, A., Stendahl, J., Zontak, M., Boutagy, N., Alkhalil, I., Eberle, M., Lin, B., O'Donnell, M., Sinusas, A., Duncan, J., 2017. Flow network based cardiac motion tracking leveraging learned feature matching. *MICCAI* 279–286.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Pennec, X., Cachier, P., Ayache, N., 1999. Understanding the “demon’s algorithm”: 3d non-rigid registration by gradient descent. *MICCAI'99* 597–605.
- Petersen, S., Matthews, P., Francis, J., Robson, M., Zembrak, F., Boubertakh, R., Young, A., Hudson, S., Weale, P., Garratt, S., Collins, R., Piechnik, S., Neubauer, S., 2016. UK Biobank’s cardiovascular magnetic resonance protocol. *J. Cardiovasc. Magn. Reson.* 18:8, 8+.
- Qin, C., Bai, W., Schlemper, J., Petersen, S., Piechnik, S., Neubauer, S., Rueckert, D., 2018. Joint learning of motion estimation and segmentation for cardiac MR image sequences. *MICCAI* 472–480.
- Qin, C., Bai, W., Schlemper, J., Petersen, S., Piechnik, S., Neubauer, S., Rueckert, D., 2018. Joint motion estimation and segmentation from undersampled cardiac MR image. *MLMIR* 2018 55–63.
- Radau, P., Lu, Y., Connelly, K., Paul, G., Dick, A., Wright, G., 2009. Evaluation framework for algorithms segmenting short axis cardiac MRI. *The MIDAS Journal - Cardiac MR Left Ventricle Segmentation Challenge* <http://hdl.handle.net/10380/3070>.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation. *MICCAI* 9351, 234–241.
- Rueckert, D., Glocker, B., Kainz, B., 2016. Learning clinically useful information from images: past, present and future. *Med. Image Anal.* 33, 13–18.
- Rühaak, J., Heldmann, S., Kipshagen, T., Fischer, B., 2013. Highly accurate fast lung CT registration. *Med. Imag.* 2013 8669.
- Suinesiaputra, A., Ablin, P., Alba, X., Alessandrini, M., Allen, J., Bai, W., Cimen, S., Claes, P., Cowan, B., D’hooge, J., Duchateau, N., Ehrhardt, J., Frangi, A., Gooya, A., Grau, V., Lekadir, K., Lu, A., Mukhopadhyay, A., Oksuz, I., Parajuli, N., Pennec, X., Pereanez, M., Pinto, C., Piras, P., Rohe, M., Rueckert, D., Saring, D., Sermesant, M., Siddiqi, K., Tabassian, M., Teresi, L., Tsaftaris, S., Wilms, M., Young, A., Zhang, X., Medrano-Gracia, P., 2018. Statistical shape modeling of the left ventricle: myocardial infarct classification challenge. *IEEE J. Biomed. Health Inform.* 22(2), 503–515.
- Suinesiaputra, A., McCulloch, A., Nash, M., Pontre, B., Young, A., 2016. Cardiac image modelling: breadth and depth in heart disease. *Med. Image Anal.* 33, 38–43.
- Tipping, M., Faul, A., 2003. Fast marginal likelihood maximisation for sparse bayesian models. *AISTATS*.
- de Vos, B., Berendsen, F., Viergever, M., Staring, M., Isgum, I., 2017. End-to-end unsupervised deformable image registration with a convolutional neural network. *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, September 14, Proceedings* 204–212.
- Weese, J., Lorenz, C., 2016. Four challenges in medical image analysis from an industrial perspective. *Med. Image Anal.* 33, 44–49.
- Wolterink, J., Leiner, T., Viergever, M., Isgum, I., 2017. Automatic segmentation and disease classification using cardiac cine MRI images. In: *Proc. Statistical Atlases and Computational Models of the Heart (STACOM), ACDC challenge, MICCAI'17 Workshop*.
- Xue, W., Brahm, G., Pandey, S., Leung, S., Li, S., 2018. Full left ventricle quantification via deep multitask relationships learning. *Med. Image Anal.* 43, 54–65.
- Yan, W., Wang, Y., Li, Z., van der Geest, R., Tao, Q., 2018. Left ventricle segmentation via optical-flow-net from short-axis cine MRI: preserving the temporal coherence of cardiac motion. *MICCAI* 613–621.

Yang, D., Wu, P., Tan, C., Pohl, K., Axel, L., Metaxas, D., 2017. 3D motion modeling and reconstruction of left ventricle wall in cardiac MRI. : International Conference on Functional Imaging and Modeling of the Heart 481–492.

Zheng, Q., Delingette, H., Duchateau, N., Ayache, N., 2018. 3D consistent and robust segmentation of cardiac images by deep learning with spatial propagation. *IEEE Trans. Med. Imaging* 37(9), 2137–2148.